Java Platform, Standard Edition

Security Developer's Guide

Release 11
E94828-04
April 2019

Java Platform, Standard Edition Security Developer's Guide, Release 11

E94828-04

Copyright © 1993, 2019, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience xix
Documentation Accessibility xix
Related Documents xix
Conventions xix

1 General Security

Terms and Definitions 1-1
Java Security Overview 1-4
- Introduction to Java Security 1-4
- Java Language Security and Bytecode Verification 1-5
- Basic Security Architecture 1-6
  - Security Providers 1-6
  - File Locations 1-8
- Java Cryptography 1-9
- Public Key Infrastructure 1-10
  - Key and Certificate Storage 1-10
  - Public Key Infrastructure Tools 1-11
- Authentication 1-12
- Secure Communication 1-13
  - TLS and DTLS Protocols 1-14
  - Simple Authentication and Security Layer (SASL) 1-14
  - Generic Security Service API and Kerberos 1-15
- Access Control 1-16
  - Permissions 1-16
  - Security Policy 1-16
  - Access Control Enforcement 1-17
- XML Signature 1-19
- Additional Information about Java Security 1-19
- Java Security Classes Summary 1-20
- Deprecated Security APIs Marked for Removal 1-22
- Security Tools Summary 1-22

Built-In Providers 1-23
Java SE Platform Security Architecture 1-23
Introduction 1-23
The Original Sandbox Model 1-24
Evolving the Sandbox Model 1-25
Protection Mechanisms – Overview of Basic Concepts 1-26
Permissions and Security Policy 1-28
The Permission Classes 1-28
java.security.CodeSource 1-38
java.security.Policy 1-38
java.security.GeneralSecurityException 1-51
Access Control Mechanisms and Algorithms 1-52
java.security.ProtectionDomain 1-52
java.security.AccessController 1-52
Inheritance of Access Control Context 1-57
java.security.AccessControlContext 1-58
Secure Class Loading 1-59
Class Loader Class Hierarchies 1-60
The Primordial Class Loader 1-60
Class Loader Delegation 1-60
Class Resolution Algorithm 1-61
Security Management 1-62
Managing Applets and Applications 1-62
SecurityManager versus AccessController 1-63
Auxiliary Tools 1-64
GuardedObject and SignedObject 1-65
java.security.GuardedObject and java.security.Guard 1-65
java.security.SignedObject 1-66
Discussion and Future Directions 1-67
Resource Consumption Management 1-67
Arbitrary Grouping of Permissions 1-67
Object-Level Protection 1-68
Subdividing Protection Domains 1-68
Running Applets with Signed Content 1-68
Appendix A: API for Privileged Blocks 1-69
Using the doPrivileged API 1-69
What It Means to Have Privileged Code 1-74
Reflection 1-76
Appendix B: Acknowledgments 1-76
Appendix C: References 1-76
Standard Algorithm Names 1-77

Permissions in the JDK 1-77
Permission Descriptions and Risks 1-78
Methods and the Permissions They Require 1-79
java.lang.SecurityManager Method Permission Checks 1-105
JDK Supported Permissions 1-110
Default Policy Implementation and Policy File Syntax 1-110
Default Policy Implementation 1-111
Default Policy File Locations 1-111
Modifying the Policy Implementation 1-113
Policy File Syntax 1-113
Policy File Examples 1-118
Property Expansion in Policy Files 1-121
Windows Systems, File Paths, and Property Expansion 1-122
General Expansion in Policy Files 1-123
Appendix A: FilePermission Path Name Canonicalization Disabled By Default 1-125
Troubleshooting Security 1-127

2 Java Cryptography Architecture (JCA) Reference Guide

Introduction to Java Cryptography Architecture 2-1
JCA Design Principles 2-2
Provider Architecture 2-3
Cryptographic Service Providers 2-3
How Providers Are Actually Implemented 2-5
Keystores 2-7
Engine Classes and Algorithms 2-7
Core Classes and Interfaces 2-8
The Provider Class 2-9
How Provider Implementations Are Requested and Supplied 2-10
Installing Providers 2-12
Provider Class Methods 2-13
The Security Class 2-13
Managing Providers 2-14
Security Properties 2-15
The SecureRandom Class 2-16
Creating a SecureRandom Object 2-16
Seeding or Re-Seeding the SecureRandom Object 2-17
Using a SecureRandom Object 2-17
Generating Seed Bytes 2-17
The MessageDigest Class 2-17
Creating a MessageDigest Object 2-18

Updating a Message Digest Object 2-18
Computing the Digest 2-18
The Signature Class 2-19
Signature Object States 2-20
Creating a Signature Object 2-20
Initializing a Signature Object 2-20
Signing with a Signature Object 2-21
Verifying with a Signature Object 2-21
The Cipher Class 2-22
Other Cipher-based Classes 2-31
The Cipher Stream Classes 2-31
The SealedObject Class 2-34
The Mac Class 2-35
Key Interfaces 2-37
The KeyPair Class 2-38
Key Specification Interfaces and Classes 2-39
The KeySpec Interface 2-39
The KeySpec Subinterfaces 2-39
The EncodedKeySpec Class 2-40
Generators and Factories 2-40
The KeyFactory Class 2-41
The SecretKeyFactory Class 2-42
The KeyPairGenerator Class 2-44
The KeyGenerator Class 2-46
The KeyAgreement Class 2-47
Key Management 2-49
The KeyStore Class 2-51
Algorithm Parameters Classes 2-55
The AlgorithmParameterSpec Interface 2-55
The AlgorithmParameters Class 2-56
The AlgorithmParameterGenerator Class 2-58
The CertificateFactory Class 2-59
How the JCA Might Be Used in a SSL/TLS Implementation 2-60
Cryptographic Strength Configuration 2-62
Jurisdiction Policy File Format 2-65
How to Make Applications Exempt from Cryptographic Restrictions 2-67
Standard Names 2-71
Packaging Your Application 2-72
Additional JCA Code Samples 2-72
Computing a MessageDigest Object 2-73
Generating a Pair of Keys 2-74

Generating and Verifying a Signature Using Generated Keys 2-76
Generating/Verifying Signatures Using Key Specifications and KeyFactory 2-77
Generating Random Numbers 2-79
Determining If Two Keys Are Equal 2-80
Reading Base64-Encoded Certificates 2-80
Parsing a Certificate Reply 2-81
Using Encryption 2-81
Using Password-Based Encryption 2-82
Sample Programs for Diffie-Hellman Key Exchange, AES/GCM, and HMAC-SHA256 2-84
Diffie-Hellman Key Exchange between Two Parties 2-84
Diffie-Hellman Key Exchange between Three Parties 2-88
AES/GCM Example 2-91
HMAC-SHA256 Example 2-92

3 How to Implement a Provider in the Java Cryptography Architecture

Who Should Read This Document 3-1
Notes on Terminology 3-1
Introduction to Implementing Providers 3-1
Engine Classes and Corresponding Service Provider Interface Classes 3-2
Steps to Implement and Integrate a Provider 3-5
Step 1: Write your Service Implementation Code 3-5
Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations 3-6
Step 2: Give your Provider a Name 3-7
Step 3: Write Your Master Class, a Subclass of Provider 3-7
Step 3.1: Create a Provider That Uses String Objects to Register Its Services 3-8
Step 3.2: Create a Provider That Uses Provider.Service 3-10
Step 3.3: Specify Additional Information for Cipher Implementations 3-12
Step 4: Create a Module Declaration for Your Provider 3-14
Step 5: Compile Your Code 3-15
Step 6: Place Your Provider in a JAR File 3-15
Step 7: Sign Your JAR File, If Necessary 3-16
Step 7.1: Get a Code-Signing Certificate 3-16
Step 7.2: Sign Your Provider 3-18
Step 8: Prepare for Testing 3-19
Step 8.1: Configure the Provider 3-19
Step 8.2: Set Provider Permissions 3-21
Step 9: Write and Compile Your Test Programs 3-22
Step 10: Run Your Test Programs 3-22

Step 11: Apply for U.S. Government Export Approval If Required 3-24
Step 12: Document Your Provider and Its Supported Services 3-25
Step 12.1: Indicate Whether Your Implementation is Cloneable for Message Digests and MACs 3-25
Step 13: Make Your Class Files and Documentation Available to Clients 3-27
Further Implementation Details and Requirements 3-27
Alias Names 3-27
Service Interdependencies 3-28
Default Initialization 3-30
Default Key Pair Generator Parameter Requirements 3-30
The Provider.Service Class 3-31
Signature Formats 3-32
DSA Interfaces and their Required Implementations 3-33
RSA Interfaces and their Required Implementations 3-35
Diffie-Hellman Interfaces and their Required Implementations 3-37
Interfaces for Other Algorithm Types 3-39
Algorithm Parameter Specification Interfaces and Classes 3-39
Key Specification Interfaces and Classes Required by Key Factories 3-42
Secret-Key Generation 3-48
Adding New Object Identifiers 3-48
Ensuring Exportability 3-50
Sample Code for MyProvider 3-50

4 JDK Providers Documentation

Introduction to JDK Providers 4-2
Import Limits on Cryptographic Algorithms 4-3
Cipher Transformations 4-3
SecureRandom Implementations 4-3
The SunPKCS11 Provider 4-4
The SUN Provider 4-5
The SunRsaSign Provider 4-8
The SunJSSE Provider 4-9
The SunJCE Provider 4-15
The SunJGSS Provider 4-21
The SunSASL Provider 4-21
The XMLDSig Provider 4-21
The SunPCSC Provider 4-22
The SunMSCAPI Provider 4-23
The SunEC Provider 4-24
The OracleUcrypto Provider 4-27
The Apple Provider 4-28

The JdkLDAP Provider 4-29
The JdkSASL Provider 4-29

5 PKCS#11 Reference Guide

SunPKCS11 Provider 5-1
SunPKCS11 Requirements 5-2
SunPKCS11 Configuration 5-2
Accessing Network Security Services (NSS) 5-8
Troubleshooting PKCS#11 5-11
Disabling PKCS#11 Providers and/or Individual PKCS#11 Mechanisms 5-11
Application Developers 5-13
- Token Login 5-13
- Token Keys 5-14
- Delayed Provider Selection 5-14
- JAAS KeyStoreLoginModule 5-15
- Tokens as JSSE Keystore and Trust Stores 5-16

Using keytool and jarsigner with PKCS#11 Tokens 5-17
Keystore Entry Syntax in Policy File 5-18
Provider Developers 5-19
- Provider Services 5-19
- Parameter Support 5-20

SunPKCS11 Provider Supported Algorithms 5-21
SunPKCS11 Provider KeyStore Requirements 5-24
Example Provider 5-26

6 Java Authentication and Authorization Service (JAAS)

Java Authentication and Authorization Service (JAAS) Reference Guide 6-1
- Who Should Read This Document 6-2
- Related Documentation 6-2
- Core Classes and Interfaces 6-2
  - Common Classes 6-2
  - Authentication Classes and Interfaces 6-8
  - Authorization Classes 6-11
- JAAS Tutorials and Sample Programs 6-12
- Appendix A: JAAS Settings in the java.security Security Properties File 6-13
  - Login Configuration Provider 6-14
  - Login Configuration URLs 6-15
  - Policy Provider 6-15
  - Policy File URLs 6-15

Appendix B: JAAS Login Configuration File 6-16
Login Configuration File Structure and Contents 6-16
Where to Specify Which Login Configuration File Should Be Used 6-18
JAAS Tutorials 6-19
JAAS Authentication Tutorial 6-19
The Authentication Tutorial Code 6-20
The Login Configuration 6-35
Running the Code 6-36
Running the Code with a Security Manager 6-37
JAAS Authorization Tutorial 6-40
What is JAAS Authorization? 6-41
How is JAAS Authorization Performed? 6-42
The Authorization Tutorial Code 6-43
The Login Configuration File for the JAAS Authorization Tutorial 6-49
The Policy File 6-49
Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide 6-53
Introduction to LoginModule 6-54
Steps to Implement a LoginModule 6-55
Step 1: Understand the Authentication Technology 6-55
Step 2: Name the LoginModule Implementation 6-56
Step 3: Implement the LoginModule Interface 6-56
Step 4: Choose or Write a Sample Application 6-61
Step 5: Compile the LoginModule and Application 6-61
Step 6: Prepare for Testing 6-61
Step 7: Test Use of the LoginModule 6-63
Step 8: Document Your LoginModule Implementation 6-64
Step 9: Make LoginModule JAR File and Documents Available 6-64

7 Java Generic Security Services (Java GSS-API)

Introduction to JAAS and Java GSS-API Tutorials 7-1
When to Use Java GSS-API Versus JSSE 7-2
Use of Java GSS-API for Secure Message Exchanges Without JAAS Programming 7-3
Overview of the Client and Server Applications 7-4
The SampleClient and SampleServer Code 7-5
Kerberos User and Service Principal Names 7-19
The Login Configuration File 7-20
The useSubjectCredsOnly System Property 7-21
Running the SampleClient and SampleServer Programs 7-22
JAAS Authentication 7-25

The Authentication Tutorial Code 7-25
The Login Configuration 7-27
Running the Code 7-28
Running the Code with a Security Manager 7-29
JAAS Authorization 7-31
What is JAAS Authorization? 7-32
How Is JAAS Authorization Performed? 7-33
The Authorization Tutorial Code 7-34
The Login Configuration File 7-36
The Policy File 7-36
Running the Authorization Tutorial Code 7-37
Use of JAAS Login Utility 7-39
What You Need to Know About the Login Utility 7-40
Application and Other File Requirements 7-40
The Sample Application Program 7-41
The Login Configuration File 7-41
The Policy File 7-42
Running the Sample Program with the Login Utility 7-43
Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges 7-45
Before You Start: Recommended Reading 7-46
Overview of the Client and Server Applications 7-46
Kerberos User and Service Principal Names 7-47
The Login Configuration File 7-48
The Policy Files 7-50
Running the SampleClient and SampleServer Programs 7-53
More Things You Can Do with Java GSS-API and JAAS 7-57
Executing Code on Behalf of the Client User 7-57
Using Credentials Delegated from the Client 7-63
Permission Required In Order to Delegate Credentials 7-65
Kerberos Requirements 7-65
Setting Properties to Indicate the Default Realm and KDC 7-66
Locating the krb5.conf Configuration File 7-66
Naming Conventions for Realm Names and Hostnames 7-67
Cross-Realm Authentication 7-67
Troubleshooting 7-68
Source Code for JAAS and Java GSS-API Tutorials 7-71
Related Documentation 7-93
Accessing Native GSS-API 7-93
Single Sign-on Using Kerberos in Java 7-97
Abstract 7-97
Introduction 7-97

Kerberos V5 7-98
Java Authentication and Authorization Service (JAAS) 7-98
Pluggable and Stackable Framework 7-98
Authentication and Authorization 7-98
Subject 7-99
doAs and doAsPrivileged 7-99
LoginContext 7-100
Callbacks 7-100
LoginModules 7-101
The Kerberos Login Module 7-101
Kerberos Classes 7-102
Authorization 7-103
Java Generic Security Service Application Program Interface (Java GSS-API) 7-103
Generic Security Service API (GSS-API) 7-103
Java GSS-API 7-104
The GSSName Interface 7-105
The GSSCredential Interface 7-105
The GSSContext Interface 7-106
Message Protection 7-110
Credential Delegation 7-111
Default Credential Acquisition Model 7-113
Exceptions to the Model 7-114
Security Risks 7-114
Credential Acquisition 7-115
Context Establishment 7-116
Credential Delegation 7-116
Conclusions 7-117
Acknowledgements 7-117
References 7-117
Advanced Security Programming in Java SE Authentication, Secure Communication and Single Sign-On 7-118
Part I : Secure Authentication using the Java Authentication and Authorization Service (JAAS) 7-119
Exercise 1: Using the JAAS API 7-119
Exercise 2: Configuring JAAS for Kerberos Authentication 7-120
Part II : Secure Communications using the Java SE Security API 7-122
Exercise 3: Using the Java Generic Security Service (GSS) API 7-122
Exercise 4: Using the Java SASL API 7-125
Exercise 5: Using the Java Secure Socket Extension with Kerberos 7-127
Part III : Deploying for Single Sign-On in a Kerberos Environment 7-130
Exercise 6: Deploying for Single Sign-On 7-130
Part IV : Secure Communications Using Stronger Encryption Algorithms 7-131

Exercise 7: Configuring to Use Stronger Encryption Algorithms in a Kerberos Environment, to Secure the Communication 7-131
Part V : Secure Authentication Using SPNEGO Java GSS Mechanism 7-133
Exercise 8: Using the Java Generic Security Services (GSS) API with SPNEGO 7-133
Part VI: HTTP/SPNEGO Authentication 7-135
Exercise 9: Using HTTP/SPNEGO Authentication 7-135
Source Code for Advanced Security Programming in Java SE Authentication, Secure Communication and Single Sign-On 7-140
Appendix A: Setting up Kerberos Accounts 7-174
The Kerberos 5 GSS-API Mechanism 7-175

8 Java Secure Socket Extension (JSSE) Reference Guide

Introduction to JSSE 8-1
JSSE Features and Benefits 8-2
JSSE Standard API 8-3
SunJSSE Provider 8-3
JSSE Related Documentation 8-3
JSSE Classes and Interfaces 8-4
JSSE Core Classes and Interfaces 8-5
SocketFactory and ServerSocketFactory Classes 8-5
SSLSocketFactory and SSLServerSocketFactory Classes 8-6
Obtaining an SSLSocketFactory 8-6
SSLSocket and SSLServerSocket Classes 8-7
Obtaining an SSLSocket 8-7
Cipher Suite Choice and Remote Entity Verification 8-7
SSLEngine Class 8-8
Understanding SSLEngine Operation Statuses 8-10
SSLEngine for TLS Protocols 8-15
SSLEngine for DTLS Protocols 8-19
Dealing With Blocking Tasks 8-28
Shutting Down a TLS/DTLS Connection 8-28
SSLSession and ExtendedSSLSession 8-29
HttpsURLConnection Class 8-30
Setting the Assigned SSLSocketFactory 8-30
Setting the Assigned HostnameVerifier 8-31
Support Classes and Interfaces 8-31
SSLContext Class 8-32
TrustManager Interface 8-34
TrustManagerFactory Class 8-34
X509TrustManager Interface 8-38

X509ExtendedTrustManager Class 8-41
KeyManager Interface 8-44
KeyManagerFactory Class 8-45
X509KeyManager Interface 8-46
X509ExtendedKeyManager Class 8-46
Relationship Between a TrustManager and a KeyManager 8-47
Secondary Support Classes and Interfaces 8-47
SSLParameters Class 8-47
SSLSessionContext Interface 8-48
SSLSessionBindingListener Interface 8-49
SSLSessionBindingEvent Class 8-49
HandShakeCompletedListener Interface 8-49
HandShakeCompletedEvent Class 8-49
HostnameVerifier Interface 8-49
X509Certificate Class 8-50
AlgorithmConstraints Interface 8-50
StandardConstants Class 8-51
SNIServerName Class 8-51
SNIMatcher Class 8-51
SNIHostName Class 8-51
Customizing JSSE 8-52
How to Specify a java.lang.System Property 8-60
How to Specify a java.security.Security Property 8-61
Customizing the X509Certificate Implementation 8-61
Specifying Default Enabled Cipher Suites 8-62
Specifying an Alternative HTTPS Protocol Implementation 8-63
Customizing the Provider Implementation 8-63
Registering the Cryptographic Provider Statically 8-64
Registering the Cryptographic Service Provider Dynamically 8-64
Provider Configuration 8-64
Configuring the Preferred Provider for Specific Algorithms 8-64
Customizing the Default Keystores and Truststores, Store Types, and Store Passwords 8-66
Customizing the Default Key Managers and Trust Managers 8-67
Disabled and Restricted Cryptographic Algorithms 8-68
Legacy Cryptographic Algorithms 8-70
Customizing the Encryption Algorithm Providers 8-70
Customizing Size of Ephemeral Diffie-Hellman Keys 8-70
Customizing Maximum Fragment Length Negotiation (MFLN) Extension 8-72
Configuring the Maximum and Minimum Packet Size 8-73
Limiting Amount of Data Algorithms May Encrypt with a Set of Keys 8-73

Client-Driven OCSP and OCSP Stapling 8-73
Client-Driven OCSP and Certificate Revocation 8-74
OCSP Stapling and Certificate Revocation 8-75
OCSP Stapling Configuration Properties 8-77
Hardware Acceleration and Smartcard Support 8-79
Configuring JSSE to Use Smartcards as Keystores and Truststores 8-79
Multiple and Dynamic Keystores 8-80
Additional Keystore Formats (PKCS12) 8-80
Server Name Indication (SNI) Extension 8-81
TLS Application Layer Protocol Negotiation 8-83
Setting up ALPN on the Client 8-84
Setting up Default ALPN on the Server 8-85
Setting up Custom ALPN on the Server 8-86
Determining Negotiated ALPN Value during Handshaking 8-88
ALPN Related Classes and Methods 8-91
Troubleshooting JSSE 8-92
Configuration Problems 8-92
SSLHandshakeException: No Available Authentication Scheme, Handshake Failure 8-92
CertificateException While Handshaking 8-93
Runtime Exception: SSL Service Not Available 8-94
Runtime Exception: "No available certificate corresponding to the SSL cipher suites which are enabled" 8-94
Runtime Exception: No Cipher Suites in Common 8-95
Socket Disconnected After Sending ClientHello Message 8-95
SunJSSE Cannot Find a JCA Provider That Supports a Required Algorithm and Causes a NoSuchAlgorithmException 8-96
Exception Thrown When Obtaining Application Resources from a Virtual Host Web Server that Requires an SNI Extension 8-97
IllegalArgumentException When RC4 Cipher Suites are Configured for DTLS 8-98
Debugging Utilities 8-98
Debugging TLS Connections 8-100
Compatibility Risks and Known Issues 8-117
Code Examples 8-118
Converting an Unsecure Socket to a Secure Socket 8-118
Running the JSSE Sample Code 8-121
Creating a Keystore to Use with JSSE 8-128
Using the Server Name Indication (SNI) Extension 8-132
Typical Client-Side Usage Examples 8-132
Typical Server-Side Usage Examples 8-133
Working with Virtual Infrastructures 8-134

Standard Names 8-139
Provider Pluggability 8-139
JAXP Security Processing 8-139

9 Running the JSSE Sample Code

Sample Truststores 9-1
Sample Code Illustrating a Secure Socket Connection Between a Client and a Server 9-3
Configuration Requirements for SSL Socket Samples 9-3
Running SSLSocketClient 9-4
Running SSLSocketClientWithTunnelling 9-6
Running SSLSocketClientWithClientAuth 9-10
Running ClassFileServer 9-13
Running SSLSocketClientWithClientAuth with ClassFileServer 9-19
Sample Code Illustrating HTTPS Connections 9-20
Running URLReader 9-20
Running URLReaderWithOptions 9-21
Sample Code Illustrating a Secure RMI Connection 9-23
Sample Code Illustrating the Use of an SSLEngine 9-28
Running SSLEngineSimpleDemo 9-28
Troubleshooting JSSE Sample Code 9-35

10 Java PKI Programmer's Guide

PKI Programmer's Guide Overview 10-1
Introduction to Public Key Certificates 10-2
X.509 Certificates and Certificate Revocation Lists (CRLs) 10-3
Core Classes and Interfaces 10-7
Basic Certification Path Classes 10-8
The CertPath Class 10-8
The CertificateFactory Class 10-9
The CertPathParameters Interface 10-11
Certification Path Validation Classes 10-11
The CertPathValidator Class 10-12
The CertPathValidatorResult Interface 10-13
Certification Path Building Classes 10-14
The CertPathBuilder Class 10-14
The CertPathBuilderResult Interface 10-15
Certificate/CRL Storage Classes 10-15
The CertStore Class 10-16
The CertStoreParameters Interface 10-17

The CertSelector and CRLSelector Interfaces 10-18
PKIX Classes 10-24
The TrustAnchor Class 10-24
The PKIXParameters Class 10-25
The PKIXCertPathValidatorResult Class 10-28
The PolicyNode Interface and PolicyQualifierInfo Class 10-28
The PKIXBuilderParameters Class 10-30
The PKIXCertPathBuilderResult Class 10-31
The PKIXCertPathChecker Class 10-32
Using PKIXCertPathChecker in Certificate Path Validation 10-38
Implementing a Service Provider 10-42
Steps to Implement and Integrate a Provider 10-43
Service Interdependencies 10-45
Certification Path Parameter Specification Interfaces 10-45
Certification Path Result Specification Interfaces 10-46
Certification Path Exception Classes 10-46
Appendix A: Standard Names 10-46
Appendix B: CertPath Implementation in SUN Provider 10-47
Appendix C: OCSP Support 10-50
Appendix D: CertPath Implementation in JdkLDAP Provider 10-53
Appendix E: Disabling Cryptographic Algorithms 10-54

11 Java SASL API Programming and Deployment Guide

Java SASL API Overview 11-2
Creating the Mechanisms 11-2
Passing Input to the Mechanisms 11-3
Using the Mechanisms 11-3
Using the Negotiated Security Layer 11-5
How SASL Mechanisms are Installed and Selected 11-6
The SunSASL Provider 11-7
The SunSASL Provider Client Mechanisms 11-8
The SunSASL Provider Server Mechanisms 11-13
The JdkSASL Provider 11-15
The JdkSASL Provider Client Mechanism 11-15
The JdkSASL Provider Server Mechanism 11-16
Debugging and Monitoring 11-17
Implementing a SASL Security Provider 11-18

12 XML Digital Signature API Overview and Tutorial

Package Hierarchy 12-1
Service Providers 12-2
Introduction to XML Signatures 12-3
Example of an XML Signature 12-3
XML Signature Secure Validation Mode 12-5
XML Digital Signature API Examples 12-5
Validate Example 12-6
Validating an XML Signature 12-10
Instantiating the Document that Contains the Signature 12-10
Specifying the Signature Element to be Validated 12-10
Creating a Validation Context 12-10
Unmarshalling the XML Signature 12-11
Validating the XML Signature 12-11
Using KeySelectors 12-12
GenEnveloped Example 12-13
Generating an XML Signature 12-17
Instantiating the Document to be Signed 12-17
Creating a Public Key Pair 12-17
Creating a Signing Context 12-18
Assembling the XML Signature 12-18
Generating the XML Signature 12-19
Printing or Displaying the Resulting Document 12-19

13 Security API Specification

14 Deprecated Security APIs Marked for Removal

15 Security Tools

16 Related Security Topics

Preface

This guide provides information about the Java security technology, tools, and implementations of commonly used security algorithms, mechanisms, and protocols on the Java Platform, Standard Edition (Java SE).

Audience

This document is intended for experienced developers who build applications using the comprehensive Java security framework. It is also intended for the user or administrator with a a set of tools to securely manage applications.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

See JDK 12 Documentation.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1
General Security

Terms and Definitions list commonly used cryptography terms and their definitions.

Java Security Overview provides an overview of the motivation of major security features, an introduction to security classes and their usage, a discussion of the impact of the security architecture on code, and thoughts on writing security-sensitive code.

Java SE Platform Security Architecture gives an overview of the motivation of the major security features implemented for the JDK.

Java Security Standard Algorithm Names Specification describes the set of standard names for algorithms, certificate and keystore types that Java SE requires and uses.

Permissions in the JDK describes the built-in JDK permission types and discusses the risks of granting each permission.

Troubleshooting Security lists options for the `java.security.debug` system property that enable you to monitor security access.

Terms and Definitions

The following are commonly used cryptography terms and their definitions.

authentication
The process of confirming the identity of a party with whom one is communicating.

certificate
A digitally signed statement vouching for the identity and public key of an entity (person, company, and so on). Certificates can either be self-signed or issued by a Certificate Authority (CA) an entity that is trusted to issue valid certificates for other entities. Well-known CAs include Comodo, Entrust, and GoDaddy. X509 is a common certificate format that can be managed by the JDK's `keytool`.

cipher suite
A combination of cryptographic parameters that define the security algorithms and key sizes used for authentication, key agreement, encryption, and integrity protection.

cryptographic hash function
An algorithm that is used to produce a relatively small fixed-size string of bits (called a hash) from an arbitrary block of data. A cryptographic hash function is similar to a checksum and has three primary characteristics: it’s a one-way function, meaning that it is not possible to produce the original data from the hash; a small change in the original data produces a large change in the resulting hash; and it doesn’t require a cryptographic key.

Cryptographic Service Provider (CSP)
Sometimes referred to simply as providers for short, the Java Cryptography Architecture (JCA) defines it as a package (or set of packages) that implements one

or more engine classes for specific cryptographic algorithms. An engine class defines a cryptographic service in an abstract fashion without a concrete implementation.

Datagram Transport Layer Security (DTLS) Protocol
A protocol that manages client and server authentication, data integrity, and encrypted communication between the client and server based on an unreliable transport channel such as UDP.

decryption
See encryption/decryption.

digital signature
A digital equivalent of a handwritten signature. It is used to ensure that data transmitted over a network was sent by whoever claims to have sent it and that the data has not been modified in transit. For example, an RSA-based digital signature is calculated by first computing a cryptographic hash of the data and then encrypting the hash with the sender's private key.

encryption/decryption
Encryption is the process of using a complex algorithm to convert an original message (cleartext) to an encoded message (ciphertext) that is unintelligible unless it is decrypted. Decryption is the inverse process of producing cleartext from ciphertext. The algorithms used to encrypt and decrypt data typically come in two categories: secret key (symmetric) cryptography and public key (asymmetric) cryptography.

endpoint identification
An IPv4 or IPv6 address used to identify an endpoint on the network.
Endpoint identification procedures are handled during SSL/TLS handshake.

handshake protocol
The negotiation phase during which the two socket peers agree to use a new or existing session. The handshake protocol is a series of messages exchanged over the record protocol. At the end of the handshake, new connection-specific encryption and integrity protection keys are generated based on the key agreement secrets in the session.

java-home
Variable placeholder used throughout this document to refer to the directory where the Java Development Kit (JDK) is installed.

key agreement
A method by which two parties cooperate to establish a common key. Each side generates some data, which is exchanged. These two pieces of data are then combined to generate a key. Only those holding the proper private initialization data can obtain the final key. Diffie-Hellman (DH) is the most common example of a key agreement algorithm.

key exchange
A method by which keys are exchanged. One side generates a private key and encrypts it using the peer's public key (typically RSA). The data is transmitted to the peer, who decrypts the key using the corresponding private key.

key manager/trust manager
Key managers and trust managers use keystores for their key material. A key manager manages a keystore and supplies public keys to others as needed (for

example, for use in authenticating the user to others). A trust manager decides who to trust based on information in the truststore it manages.

Keyed-Hash Message Code (HMAC)
A specific type of message authentication code that involves a cryptographic hash function and a secret cryptographic key.

Keyed-Hash Message Code (HMAC)-based Extract-and-Expand Key Derivation Function (HKDF)
A function used for key generation and validation.

keystore/truststore
A keystore is a database of key material. Key material is used for a variety of purposes, including authentication and data integrity. Various types of keystores are available, including PKCS12 and Oracle's JKS.
Generally speaking, keystore information can be grouped into two categories: key entries and trusted certificate entries. A key entry consists of an entity's identity and its private key, and can be used for a variety of cryptographic purposes. In contrast, a trusted certificate entry contains only a public key in addition to the entity's identity. Thus, a trusted certificate entry can't be used where a private key is required, such as in a `javax.net.ssl.KeyManager`. In the JDK implementation of JKS, a keystore may contain both key entries and trusted certificate entries.
A truststore is a keystore that is used when making decisions about what to trust. If you receive data from an entity that you already trust, and if you can verify that the entity is the one that it claims to be, then you can assume that the data really came from that entity.
An entry should only be added to a truststore if the user trusts that entity. By either generating a key pair or by importing a certificate, the user gives trust to that entry. Any entry in the truststore is considered a trusted entry.
It may be useful to have two different keystore files: one containing just your key entries, and the other containing your trusted certificate entries, including CA certificates. The former contains private information, whereas the latter does not. Using two files instead of a single keystore file provides a cleaner separation of the logical distinction between your own certificates (and corresponding private keys) and others' certificates. To provide more protection for your private keys, store them in a keystore with restricted access, and provide the trusted certificates in a more publicly accessible keystore if needed.

message authentication code (MAC)
Provides a way to check the integrity of information transmitted over or stored in an unreliable medium, based on a secret key. Typically, MACs are used between two parties that share a secret key in order to validate information transmitted between these parties.
A MAC mechanism that is based on cryptographic hash functions is referred to as HMAC. HMAC can be used with any cryptographic hash function, such as Message Digest 5 (MD5) and the Secure Hash Algorithm (SHA-256), in combination with a secret shared key. HMAC is specified in RFC 2104.

public-key cryptography
A cryptographic system that uses an encryption algorithm in which two keys are produced. One key is made public, whereas the other is kept private. The public key and the private key are cryptographic inverses; what one key encrypts only the other key can decrypt. Public-key cryptography is also called asymmetric cryptography.

Record Protocol
A protocol that packages all data (whether application-level or as part of the handshake process) into discrete records of data much like a TCP stream socket converts an application byte stream into network packets. The individual records are then protected by the current encryption and integrity protection keys.

secret-key cryptography
A cryptographic system that uses an encryption algorithm in which the same key is used both to encrypt and decrypt the data. Secret-key cryptography is also called symmetric cryptography.

Secure Sockets Layer (SSL) Protocol
A protocol that manages client and server authentication, data integrity, and encrypted communication between the client and server. SSL has been renamed to Transport Layer Security (TLS).

session
A named collection of state information including authenticated peer identity, cipher suite, and key agreement secrets that are negotiated through a secure socket handshake and that can be shared among multiple secure socket instances.

Transport Layer Security (TLS) Protocol
A protocol that manages client and server authentication, data integrity, and encrypted communication between the client and server based on a reliable transport channel such as TCP.

trust manager
See key manager/trust manager.

truststore
See keystore/truststore.

Java Security Overview

Java security includes a large set of APIs, tools, and implementations of commonly-used security algorithms, mechanisms, and protocols. The Java security APIs span a wide range of areas, including cryptography, public key infrastructure, secure communication, authentication, and access control. Java security technology provides the developer with a comprehensive security framework for writing applications, and also provides the user or administrator with a set of tools to securely manage applications.

Introduction to Java Security

The JDK is designed with a strong emphasis on security. At its core, the Java language itself is type-safe and provides automatic garbage collection, enhancing the robustness of application code. A secure class loading and verification mechanism ensures that only legitimate Java code is executed. The Java security architecture includes a large set of application programming interfaces (APIs), tools, and implementations of commonly-used security algorithms, mechanisms, and protocols.

The Java security APIs span a wide range of areas. Cryptographic and public key infrastructure (PKI) interfaces provide the underlying basis for developing secure applications. Interfaces for performing authentication and access control enable applications to guard against unauthorized access to protected resources.

The APIs allow for multiple interoperable implementations of algorithms and other security services. Services are implemented in *providers*, which are plugged into the JDK through a standard interface that makes it easy for applications to obtain security services without having to know anything about their implementations. This allows developers to focus on how to integrate security into their applications, rather than on how to actually implement complex security mechanisms.

The JDK includes a number of providers that implement a core set of security services. It also allows for additional custom providers to be installed. This enables developers to extend the platform with new security mechanisms.

The JDK is divided into modules. Modules that contain security APIs include the following:

Table 1-1 Modules That Contain Security APIs

Module	Description
`java.base`	Defines the foundational APIs of Java SE; contained packages include `java.security`, `javax.crypto`, `javax.net.ssl`, and `javax.security.auth`
`java.security.jgss`	Defines the Java binding of the IETF Generic Security Services API (GSS-API). This module also contains GSS-API mechanisms including Kerberos v5 and SPNEGO
`java.security.sasl`	Defines Java support for the IETF Simple Authentication and Security Layer (SASL). This module also contains SASL mechanisms including DIGEST-MD5, CRAM-MD5, and NTLM
`java.smartcardio`	Defines the Java Smart Card I/O API
`java.xml.crypto`	Defines the API for XML cryptography
`jdk.jartool`	Defines APIs for signing jar files
`jdk.security.auth`	Provides implementations of the `javax.security.auth.*` interfaces and various authentication modules
`jdk.security.jgss`	Defines Java extensions to the GSS-API and an implementation of the SASL GSS-API mechanism

Java Language Security and Bytecode Verification

The Java language is designed to be type-safe and easy to use. It provides automatic memory management, garbage collection, and range-checking on arrays. This reduces the overall programming burden placed on developers, leading to fewer subtle programming errors and to safer, more robust code.

A compiler translates Java programs into a machine-independent bytecode representation. A bytecode verifier is invoked to ensure that only legitimate bytecodes are executed in the Java runtime. It checks that the bytecodes conform to the Java Language Specification and do not violate Java language rules or namespace restrictions. The verifier also checks for memory management violations, stack

underflows or overflows, and illegal data typecasts. Once bytecodes have been verified, the Java runtime prepares them for execution.

In addition, the Java language defines different access modifiers that can be assigned to Java classes, methods, and fields, enabling developers to restrict access to their class implementations as appropriate. The language defines four distinct access levels:

- `private`: Most restrictive modifier; access is not allowed outside the particular class in which the private member (a method, for example) is defined.
- `protected`: Allows access to any subclass or to other classes within the same package.
- Package-private: If not specified, then this is the default access level; allows access to classes within the same package.
- `public`: No longer guarantees that the element is accessible everywhere; accessibility depends upon whether the package containing that element is exported by its defining module and whether that module is readable by the module containing the code that is attempting to access it.

Basic Security Architecture

The JDK defines a set of APIs spanning major security areas, including cryptography, public key infrastructure, authentication, secure communication, and access control. The APIs allow developers to easily integrate security into their application code.

The APIs are designed around the following principles:

Implementation independence
Applications do not need to implement security themselves. Rather, they can request security services from the JDK. Security services are implemented in providers (see the section Security Providers), which are plugged into the JDK via a standard interface. An application may rely on multiple independent providers for security functionality.

Implementation interoperability
Providers are interoperable across applications. Specifically, an application is not bound to a specific provider if it does not rely on default values from the provider.

Algorithm extensibility
The JDK includes a number of built-in providers that implement a basic set of security services that are widely used today. However, some applications may rely on emerging standards not yet implemented, or on proprietary services. The JDK supports the installation of custom providers that implement such services.

Security Providers

The `java.security.Provider` class encapsulates the notion of a security provider in the Java platform. It specifies the provider's name and lists the security services it implements. Multiple providers may be configured at the same time and are listed in order of preference. When a security service is requested, the highest priority provider that implements that service is selected.

Applications rely on the relevant `getInstance` method to request a security service from an underlying provider.

For example, message digest creation represents one type of service available from providers. To request an implementation of a specific message digest algorithm, call the method `java.security.MessageDigest.getInstance`. The following statement requests a SHA-256 message digest implementation without specifying a provider name:

```
MessageDigest md = MessageDigest.getInstance("SHA-256");
```

The following figure illustrates how this statement obtains a SHA-256 message digest implementation. The providers are searched in preference order, and the implementation from the first provider supplying that particular algorithm, `ProviderB`, is returned.

Figure 1-1 Request SHA-256 Message Digest Implementation Without Specifying Provider

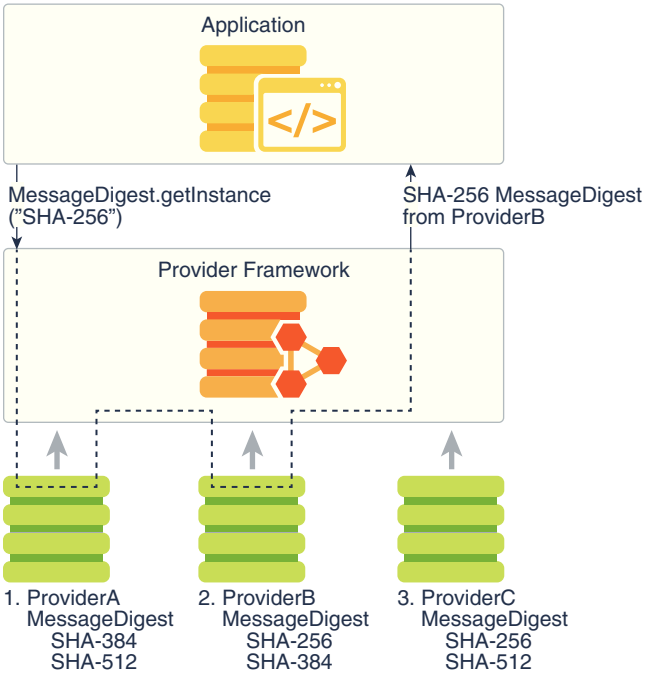

You can optionally request an implementation from a specific provider by specifying the provider's name. The following statement requests a SHA-256 message digest implementation from a specific provider, `ProviderC`:

```
MessageDigest md = MessageDigest.getInstance("SHA-256", "ProviderC");
```

The following figure illustrates how this statement requests a SHA-256 message digest implementation from a specific provider, `ProviderC`. In this case, the implementation from that provider is returned, even though a provider with a higher preference order, `ProviderB`, also supplies a SHA-256 implementation.

Figure 1-2 Request SHA-256 Message Digest Implementation from Specific Provider

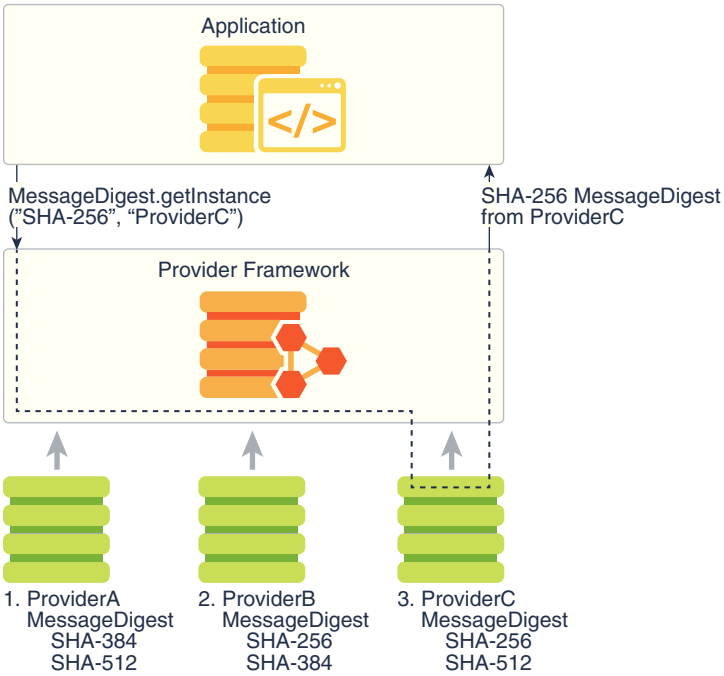

For more information about cryptographic services, such as message digest algorithms, see the section Java Cryptography.

Oracle's implementation of the Java platform includes a number of built-in default providers that implement a basic set of security services that can be used by applications. Note that other vendor implementations of the Java platform may include different sets of providers that encapsulate vendor-specific sets of security services. The term built-in default providers refers to the providers available in Oracle's implementation.

File Locations

The following table lists locations of some security-related files and tools.

Table 1-2 Java security files and tools

File Name or Tool Name	Location	Description
`java.security`	`<java-home>/conf/security`	Certain aspects of Java security, such as configuring the providers, may be customized by setting Security Properties. You may set Security Properties statically in the `java.security` file. Security Properties may also be set dynamically by calling appropriate methods of the `Security` class (in the `java.security` package).
`java.policy`	`<java-home>/conf/security`	This is the default system policy file; see Security Policy.
Cryptographic policy directory	`<java-home>/conf/security/policy`	This directory contains sets of jurisdiction policy files; see Cryptographic Strength Configuration.
`cacerts`	`<java-home>/lib/security`	The `cacerts` file represents a system-wide keystore with Certificate Authority (CA) and other trusted certificates. For information about configuring and managing this file, see keytool in *Java Platform, Standard Edition Tools Reference*.
`keytool`, `jarsigner` Windows only: `kinit`, `klist`, `ktab`	`<java-home>/bin`	For more information about security-related tools, see Security Tools and Commands in *Java Platform, Standard Edition Tools Reference*.

Java Cryptography

The Java cryptography architecture is a framework for accessing and developing cryptographic functionality for the Java platform.

It includes APIs for a large variety of cryptographic services, including the following:

- Message digest algorithms
- Digital signature algorithms
- Symmetric bulk and stream encryption
- Asymmetric encryption
- Password-based encryption (PBE)
- Elliptic Curve Cryptography (ECC)
- Key agreement algorithms

- Key generators
- Message Authentication Codes (MACs)
- Secure Random Number Generators

For historical (export control) reasons, the cryptography APIs are organized into two distinct packages:

- The `java.security` and `java.security.*` packages contains classes that are *not* subject to export controls (like `Signature` and `MessageDigest`)
- The `javax.crypto` package contains classes that are subject to export controls (like `Cipher` and `KeyAgreement`)

The cryptographic interfaces are provider-based, allowing for multiple and interoperable cryptography implementations. Some providers may perform cryptographic operations in software; others may perform the operations on a hardware token (for example, on a smart card device or on a hardware cryptographic accelerator). Providers that implement export-controlled services must be digitally signed by a certificate issued by the Oracle JCE Certificate Authority.

The Java platform includes built-in providers for many of the most commonly used cryptographic algorithms, including the RSA, DSA, and ECDSA signature algorithms, the AES encryption algorithm, the SHA-2 message digest algorithms, and the Diffie-Hellman (DH) and Elliptic Curve Diffie-Hellman (ECDH) key agreement algorithms. Most of the built-in providers implement cryptographic algorithms in Java code.

The Java platform also includes a built-in provider that acts as a bridge to a native PKCS#11 (v2.x) token. This provider, named `SunPKCS11`, allows Java applications to seamlessly access cryptographic services located on PKCS#11-compliant tokens.

On Windows, the Java platform includes a built-in provider that acts as a bridge to the native Microsoft CryptoAPI. This provider, named `SunMSCAPI`, allows Java applications to seamlessly access cryptographic services on Windows through the CryptoAPI.

Public Key Infrastructure

Public Key Infrastructure (PKI) is a term used for a framework that enables secure exchange of information based on public key cryptography. It allows identities (of people, organizations, etc.) to be bound to digital certificates and provides a means of verifying the authenticity of certificates. PKI encompasses keys, certificates, public key encryption, and trusted Certification Authorities (CAs) who generate and digitally sign certificates.

The Java platform includes APIs and provider support for X.509 digital certificates and Certificate Revocation Lists (CRLs), as well as PKIX-compliant certification path building and validation. The classes related to PKI are located in the `java.security` and `java.security.cert` packages.

Key and Certificate Storage

The Java platform provides for long-term persistent storage of cryptographic keys and certificates via key and certificate stores. Specifically, the `java.security.KeyStore` class represents a *key store*, a secure repository of cryptographic keys and/or trusted certificates (to be used, for example, during certification path validation), and the `java.security.cert.CertStore` class represents a *certificate store*, a public and

potentially vast repository of unrelated and typically untrusted certificates. A `CertStore` may also store CRLs.

`KeyStore` and `CertStore` implementations are distinguished by types. The Java platform includes the standard PKCS11 and PKCS12 key store types (whose implementations are compliant with the corresponding PKCS specifications from RSA Security). It also contains a proprietary file-based key store type called JKS (which stands for Java Key Store), and a type called DKS (Domain Key Store) which is a collection of keystores that are presented as a single logical keystore.

The Java platform includes a special built-in key store, `cacerts`, that contains a number of certificates for well-known, trusted CAs. The keytool utility is able to list the certificates included in `cacerts`. See keytool in *Java Platform, Standard Edition Tools Reference*.

The SunPKCS11 provider mentioned in the section Java Cryptography includes a PKCS11 `KeyStore` implementation. This means that keys and certificates residing in secure hardware (such as a smart card) can be accessed and used by Java applications via the `KeyStore` API. Note that smart card keys may not be permitted to leave the device. In such cases, the `java.security.Key` object returned by the `KeyStore` API may simply be a reference to the key (that is, it would not contain the actual key material). Such a `Key` object can only be used to perform cryptographic operations on the device where the actual key resides.

The Java platform also includes an LDAP certificate store type (for accessing certificates stored in an LDAP directory), as well as an in-memory Collection certificate store type (for accessing certificates managed in a `java.util.Collection` object).

Public Key Infrastructure Tools

There are two built-in tools for working with keys, certificates, and key stores:

- `keytool` creates and manages key stores. Use it to perform the following tasks:
  - Create public/private key pairs
  - Display, import, and export X.509 v1, v2, and v3 certificates stored as files
  - Create X.509 certificates
  - Issue certificate (PKCS#10) requests to be sent to CAs
  - Create certificates based on certificate requests
  - Import certificate replies (obtained from the CAs sent certificate requests)
  - Designate public key certificates as trusted
  - Accept a password and store it securely as a secret key
- `jarsigner` signs JAR files and verifies signatures on signed JAR files. The Java ARchive (JAR) file format enables the bundling of multiple files into a single file. Typically, a JAR file contains the class files and auxiliary resources associated with applets and applications.

To digitally sign code, perform the following:

1. Use `keytool` to generate or import appropriate keys and certificates into your key store (if they are not there already).
2. Use the `jar` tool to package the code in a JAR file.

3. Use the `jarsigner` tool (or the `jdk.security.jarsigner` API) to sign the JAR file. The `jarsigner` tool accesses a key store to find any keys and certificates needed to sign a JAR file or to verify the signature of a signed JAR file.

> **Note:**
>
> `jarsigner` can optionally generate signatures that include a timestamp. Systems that verify JAR file signatures can check the timestamp and accept a JAR file that was signed while the signing certificate was valid rather than requiring the certificate to be current. (Certificates typically expire annually, and it is not reasonable to expect JAR file creators to re-sign deployed JAR files annually.)

See keytool and jarsigner in *Java Platform, Standard Edition Tools Reference*.

Authentication

Authentication is the process of determining the identity of a user. In the context of the Java runtime environment, it is the process of identifying the user of an executing Java program. In certain cases, this process may rely on the services described in the section Java Cryptography.

The Java platform provides APIs that enable an application to perform user authentication via pluggable login modules. Applications call into the `LoginContext` class (in the `javax.security.auth.login` package), which in turn references a configuration. The configuration specifies which login module (an implementation of the `javax.security.auth.spi.LoginModule` interface) is to be used to perform the actual authentication.

Since applications solely talk to the standard `LoginContext` API, they can remain independent from the underlying plug-in modules. New or updated modules can be plugged in for an application without having to modify the application itself. The following figure illustrates the independence between applications and underlying login modules:

Figure 1-3 Authentication Login Modules Plugging into the Authentication Framework

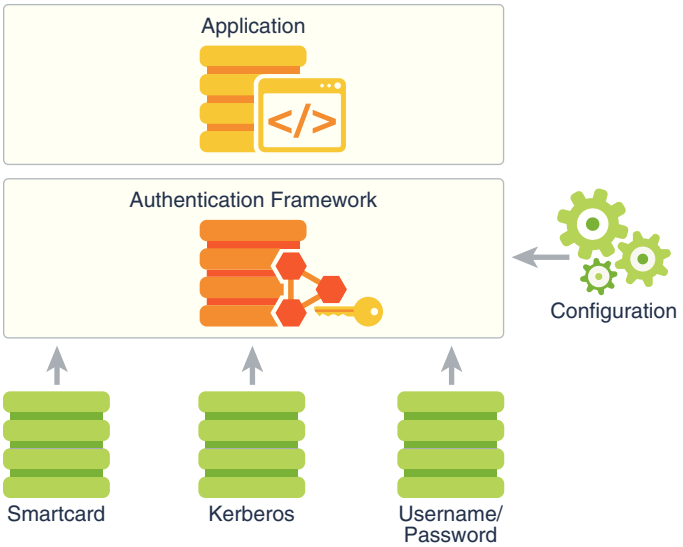

It is important to note that although login modules are pluggable components that can be configured into the Java platform, they are not plugged in via security providers. Therefore, they do not follow the provider searching model as described in the section Security Providers. Instead, as is shown in Figure 1-3, login modules are administered by their own unique configuration.

The Java platform provides the following built-in login modules, all in the `com.sun.security.auth.module` package:

- `Krb5LoginModule` for authentication using Kerberos protocols
- `JndiLoginModule` for username/password authentication using LDAP or NIS databases
- `KeyStoreLoginModule` for logging into any type of key store, including a PKCS#11 token key store

Authentication can also be achieved during the process of establishing a secure communication channel between two peers. The Java platform provides implementations of a number of standard communication protocols, which are discussed in the section Secure Communication.

Secure Communication

The data that travels across a network can be accessed by someone who is not the intended recipient. When the data includes private information, such as passwords and credit card numbers, steps must be taken to make the data unintelligible to unauthorized parties. It is also important to ensure that you are sending the data to the appropriate party, and that the data has not been modified, either intentionally or unintentionally, during transport.

Cryptography forms the basis required for secure communication; see the section Java Cryptography. The Java platform also provides API support and provider implementations for a number of standard secure communication protocols.

TLS and DTLS Protocols

Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols which provide a secure channel between two communication peers. TLS uses a combination of cryptographic processes by providing authentication, confidentiality and integrity properties for communication over a untrusted or potential hostile network. TLS runs over a reliable, stream-oriented transport channel, typically Transport Control Protocol (TCP). TLS is application protocol independent. Higher-level protocols, for example Hypertext Transfer Protocol (HTTP), can layer on top of TLS transparently.

The Datagram Transport Layer Security (DTLS) protocols are based on the stream-oriented TLS protocols and are intended to provider similar security properties for datagram transport, like User Datagram Protocol (UDP), which does not provide reliable or in-order delivery of data.

The JDK provides APIs and an implementation of the SSL, TLS, and DTLS protocols that includes functionality for data encryption, message integrity, and server and client authentication. Applications can use (D)TLS to provide for the secure passage of data between two peers over any application protocol, such as HTTP on top of TCP/IP.

The `javax.net.ssl.SSLSocket` class represents a network socket that encapsulates TLS support on top of a normal stream socket (`java.net.Socket`). Some applications might want to use alternate data transport abstractions (for example, New-I/O); the `javax.net.ssl.SSLEngine` class is available to produce and consume TLS/DTLS packets.

The JDK also includes APIs that support the notion of pluggable (provider-based) key managers and trust managers. *A key manager* is encapsulated by the `javax.net.ssl.KeyManager` class, and manages the keys used to perform authentication. A *trust manager* is encapsulated by the `TrustManager` class (in the same package), and makes decisions about who to trust based on certificates in the key store it manages.

The JDK includes a built-in provider that implements the SSL/TLS/DTLS protocols:

- SSL 3.0
- TLS 1.0
- TLS 1.1
- TLS 1.2
- TLS 1.3
- DTLS 1.0
- DTLS 1.2

Simple Authentication and Security Layer (SASL)

Simple Authentication and Security Layer (SASL) is an Internet standard that specifies a protocol for authentication and optional establishment of a security layer between client and server applications. SASL defines how authentication data is to be exchanged, but does not itself specify the contents of that data. It is a framework into which specific authentication mechanisms that specify the contents and semantics of the authentication data can fit. There are a number of standard SASL mechanisms

defined by the Internet community for various security levels and deployment scenarios.

The Java SASL API, which is in the `java.security.sasl` module, defines classes and interfaces for applications that use SASL mechanisms. It is defined to be mechanism-neutral; an application that uses the API need not be hardwired into using any particular SASL mechanism. Applications can select the mechanism to use based on desired security features. The API supports both client and server applications. The `javax.security.sasl.Sasl` class is used to create `SaslClient` and `SaslServer` objects.

SASL mechanism implementations are supplied in provider packages. Each provider may support one or more SASL mechanisms and is registered and invoked via the standard provider architecture.

The Java platform includes a built-in provider that implements the following SASL mechanisms:

- CRAM-MD5, DIGEST-MD5, EXTERNAL, GSSAPI, NTLM, and PLAIN client mechanisms
- CRAM-MD5, DIGEST-MD5, GSSAPI, and NTLM server mechanisms

Generic Security Service API and Kerberos

The Java platform contains an API with the Java language bindings for the Generic Security Service Application Programming Interface (GSS-API), which is in the java.security.jgss module. GSS-API offers application programmers uniform access to security services atop a variety of underlying security mechanisms. The Java GSS-API currently requires use of a Kerberos v5 mechanism, and the Java platform includes a built-in implementation of this mechanism. At this time, it is not possible to plug in additional mechanisms.

> **Note:**
>
> The `Krb5LoginModule` mentioned in the section Authentication can be used in conjunction with the GSS Kerberos mechanism.

The Java platform also includes a built-in implementation of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) GSS-API mechanism.

Before two applications can use GSS-API to securely exchange messages between them, they must establish a joint security context. The context encapsulates shared state information that might include, for example, cryptographic keys. Both applications create and use an `org.ietf.jgss.GSSContext` object to establish and maintain the shared information that makes up the security context. Once a security context has been established, it can be used to prepare secure messages for exchange.

The Java GSS APIs are in the `org.ietf.jgss` package. The Java platform also defines basic Kerberos classes, like `KerberosPrincipal`, `KerberosTicket`, `KerberosKey`, and `KeyTab`, which are located in the `javax.security.auth.kerberos` package.

Access Control

The access control architecture in the Java platform protects access to sensitive resources (for example, local files) or sensitive application code (for example, methods in a class). All access control decisions are mediated by a security manager, represented by the `java.lang.SecurityManager` class. A `SecurityManager` must be installed into the Java runtime in order to activate the access control checks.

Local applications executed via the `java` command are by default not run with a `SecurityManager` installed. In order to run local applications with a `SecurityManager`, either the application itself must programmatically set one via the `setSecurityManager` method (in the `java.lang.System` class), or `java` must be invoked with a `-Djava.security.manager` argument on the command line.

Permissions

A permission represents access to a system resource. In order for a resource access to be allowed for an applet (or an application running with a security manager), the corresponding permission must be explicitly granted to the code attempting the access.

When Java code is loaded by a class loader into the Java runtime, the class loader automatically associates the following information with that code:

- Where the code was loaded from
- Who signed the code (if anyone)
- Default permissions granted to the code

This information is associated with the code regardless of whether the code is downloaded over an untrusted network (e.g., an applet) or loaded from the filesystem (e.g., a local application). The location from which the code was loaded is represented by a URL, the code signer is represented by the signer's certificate chain, and default permissions are represented by `java.security.Permission` objects.

The default permissions automatically granted to downloaded code include the ability to make network connections back to the host from which it originated. The default permissions automatically granted to code loaded from the local filesystem include the ability to read files from the directory it came from, and also from subdirectories of that directory.

Note that the identity of the user executing the code is not available at class loading time. It is the responsibility of application code to authenticate the end user if necessary (see the section Authentication). Once the user has been authenticated, the application can dynamically associate that user with executing code by invoking the `doAs` method in the `javax.security.auth.Subject` class.

Security Policy

A limited set of default permissions are granted to code by class loaders. Administrators have the ability to flexibly manage additional code permissions via a security policy.

Java SE encapsulates the notion of a security policy in the `java.security.Policy` class. There is only one `Policy` object installed into the Java runtime at any given

time. The basic responsibility of the `Policy` object is to determine whether access to a protected resource is permitted to code (characterized by where it was loaded from, who signed it, and who is executing it). How a `Policy` object makes this determination is implementation-dependent. For example, it may consult a database containing authorization data, or it may contact another service.

Java SE includes a default `Policy` implementation that reads its authorization data from one or more ASCII (UTF-8) files configured in the security properties file. These policy files contain the exact sets of permissions granted to code: specifically, the exact sets of permissions granted to code loaded from particular locations, signed by particular entities, and executing as particular users. The policy entries in each file must conform to a documented proprietary syntax and may be composed via a simple text editor.

Access Control Enforcement

The Java runtime keeps track of the sequence of Java calls that are made as a program executes. When access to a protected resource is requested, the entire call stack, by default, is evaluated to determine whether the requested access is permitted.

As mentioned previously, resources are protected by the `SecurityManager`. Security-sensitive code in the JDK and in applications protects access to resources via code like the following:

```
SecurityManager sm = System.getSecurityManager();
if (sm != null) {
   sm.checkPermission(perm);
}
```

The `Permission` object `perm` corresponds to the requested access. For example, if an attempt is made to read the file `/tmp/abc`, the permission may be constructed as follows:

```
Permission perm = new java.io.FilePermission("/tmp/abc", "read");
```

The default implementation of `SecurityManager` delegates its decision to the `java.security.AccessController` implementation. The `AccessController` traverses the call stack, passing to the installed security `Policy` each code element in the stack, along with the requested permission (for example, the `FilePermission` in the previous example). The `Policy` determines whether the requested access is granted, based on the permissions configured by the administrator. If access is not granted, the `AccessController` throws a `java.lang.SecurityException`.

Figure 1-4 illustrates access control enforcement. In this particular example, there are initially two elements on the call stack, `ClassA` and `ClassB`. `ClassA` invokes a method in `ClassB`, which then attempts to access the file `/tmp/abc` by creating an instance of `java.io.FileInputStream`. The `FileInputStream` constructor creates a `FilePermission`, `perm`, as shown above, and then passes `perm` to the `SecurityManager` class's `checkPermission` method. In this particular case, only the permissions for `ClassA` and `ClassB` need to be checked, because all classes in the `java.base` module, including `FileInputStream`, `SecurityManager`, and `AccessController`, automatically receives all permissions.

In this example, `ClassA` and `ClassB` have different code characteristics – they come from different locations and have different signers. Each may have been granted a different set of permissions. The `AccessController` only grants access to the requested file if the `Policy` indicates that both classes have been granted the required `FilePermission`.

Figure 1-4 Controlling Access to Resources

XML Signature

The Java XML Digital Signature API is a standard Java API for generating and validating XML Signatures.

XML Signatures can be applied to data of any type, XML or binary (see XML Signature Syntax and Processing). The resulting signature is represented in XML. An XML Signature can be used to secure your data and provide data integrity, message authentication, and signer authentication.

The API is designed to support all of the required or recommended features of the W3C Recommendation for XML-Signature Syntax and Processing. The API is extensible and pluggable and is based on the Java Cryptography Service Provider Architecture.

The Java XML Digital Signature API, which is in the `java.xml.crypto` module, consists of six packages:

- `javax.xml.crypto`
- `javax.xml.crypto.dsig`
- `javax.xml.crypto.dsig.keyinfo`
- `javax.xml.crypto.dsig.spec`
- `javax.xml.crypto.dom`
- `javax.xml.crypto.dsig.dom`

Additional Information about Java Security

Find additional Java security documentation at Java SE Security.

> **Note:**
>
> Historically, as new types of security services were added to Java SE (sometimes initially as extensions), various acronyms were used to refer to them. Since these acronyms are still in use in the Java security documentation, here is an explanation of what they represent:
>
> - JSSE (Java Secure Socket Extension) refers to the SSL-related services as described in the section TLS and DTLS Protocols
> - JCE (Java Cryptography Extension) refers to cryptographic services as described in the section Java Cryptography
> - JAAS (Java Authentication and Authorization Service) refers to the authentication and user-based access control services as described in the sections Authentication and Access Control, respectively

Java Security Classes Summary

The following table describes some of the names, packages, and usage of the Java security classes and interfaces..

Table 1-3 Java security packages and classes

Package	Class/Interface Name	Usage	Module
`java.lang`	`SecurityException`	Indicates a security violation	`java.base`
`java.lang`	`SecurityManager`	Mediates all access control decisions	`java.base`
`java.lang`	`System`	Installs the SecurityManager	`java.base`
`java.security`	`AccessController`	Called by default implementation of SecurityManager to make access control decisions	`java.base`
`java.security`	`DomainLoadStoreParameter`	Stores parameters for the Domain keystore (DKS)	`java.base`
`java.security`	`Key`	Represents a cryptographic key	`java.base`
`java.security`	`KeyStore`	Represents a repository of keys and trusted certificates	`java.base`
`java.security`	`MessageDigest`	Represents a message digest	`java.base`
`java.security`	`Permission`	Represents access to a particular resource	`java.base`
`java.security`	`PKCS12Attribute`	Supports attributes in PKCS12 keystores	`java.base`
`java.security`	`Policy`	Encapsulates the security policy	`java.base`
`java.security`	`Provider`	Encapsulates security service implementations	`java.base`
`java.security`	`Security`	Manages security providers and Security Properties	`java.base`
`java.security`	`Signature`	Creates and verifies digital signatures	`java.base`
`java.security.cert`	`Certificate`	Represents a public key certificate	`java.base`
`java.security.cert`	`CertStore`	Represents a repository of unrelated and typically untrusted certificates	`java.base`
`java.security.cert`	`CRL`	Represents a CRL	`java.base`
`javax.crypto`	`Cipher`	Performs encryption and decryption	`java.base`
`javax.crypto`	`KeyAgreement`	Performs a key exchange	`java.base`
`javax.net.ssl`	`KeyManager`	Manages keys used to perform SSL/TLS authentication	`java.base`

Table 1-3 (Cont.) Java security packages and classes

Package	Class/Interface Name	Usage	Module
`javax.net.ssl`	`SSLEngine`	Produces/consumes SSL/TLS packets, allowing the application freedom to choose a transport mechanism	`java.base`
`javax.net.ssl`	`SSLSocket`	Represents a network socket that encapsulates SSL/TLS support on top of a normal stream socket	`java.base`
`javax.net.ssl`	`TrustManager`	Makes decisions about who to trust in SSL/TLS interactions (for example, based on trusted certificates in key stores)	`java.base`
`javax.security.auth`	`Subject`	Represents a user	`java.base`
`javax.security.auth.kerberos`	`KerberosPrincipal`	Represents a Kerberos principal	`java.base`
`javax.security.auth.kerberos`	`KerberosTicket`	Represents a Kerberos ticket	`java.base`
`javax.security.auth.kerberos`	`KerberosKey`	Represents a Kerberos key	`java.base`
`javax.security.auth.kerberos`	`KerberosTab`	Represents a Kerberos keytab file	`java.base`
`javax.security.auth.login`	`LoginContext`	Supports pluggable authentication	`java.base`
`javax.security.auth.spi`	`LoginModule`	Implements a specific authentication mechanism	`java.base`
`javax.security.sasl`	`Sasl`	Creates SaslClient and SaslServer objects	`java.security.sasl`
`javax.security.sasl`	`SaslClient`	Performs SASL authentication as a client	`java.security.sasl`
`javax.security.sasl`	`SaslServer`	Performs SASL authentication as a server	`java.security.sasl`
`jdk.security.jarsigner`	`JarSigner`	Signs a JAR file	`jdk.jartool`
`org.ietf.jgss`	`GSSContext`	Encapsulates a GSS-API security context and provides the security services available via the context	`java.security.jgss`
`com.sun.security.auth.module`	`JndiLoginModule`	Performs username/ password authentication using LDAP or NIS	`jdk.security.auth`
`com.sun.security.auth.module`	`KeyStoreLoginModule`	Performs authentication based on key store login	`jdk.security.auth`
`com.sun.security.auth.module`	`Krb5LoginModule`	Performs authentication using Kerberos protocols	`jdk.security.auth`

Deprecated Security APIs Marked for Removal

The following APIs are deprecated and eligible to be removed in a future release.

You can check the API dependencies using the `jdeprscan` tool. See jdeprscan in *Java Platform, Standard Edition Tools Reference*.

The following classes are deprecated and marked for removal:

- `com.sun.security.auth.PolicyFile`
- `com.sun.security.auth.SolarisNumericGroupPrincipal`
- `com.sun.security.auth.SolarisNumericUserPrincipal`
- `com.sun.security.auth.SolarisPrincipal`
- `com.sun.security.auth.X500Principal`
- `com.sun.security.auth.module.SolarisLoginModule`
- `com.sun.security.auth.module.SolarisSystem`

The following methods are deprecated and marked for removal:

- `java.lang.SecurityManager.getInCheck`
- `java.lang.SecurityManager.checkMemberAccess`
- `java.lang.SecurityManager.classDepth`
- `java.lang.SecurityManager.currentClassLoader`
- `java.lang.SecurityManager.currentLoadedClass`
- `java.lang.SecurityManager.inClass`
- `java.lang.SecurityManager.inClassLoader`
- `java.lang.SecurityManager.checkAwtEventQueueAccess`
- `java.lang.SecurityManager.checkTopLevelWindow`
- `java.lang.SecurityManager.checkSystemClipboardAccess`

The following field is deprecated and marked for removal:

- `java.lang.SecurityManager.incheck`

Security Tools Summary

The following tables describe Java security and Kerberos-related tools.

See Security Tools and Commands in *Java Platform, Standard Edition Tools Reference*.

Table 1-4 Java Security Tools

Tool	Usage
`jar`	Creates Java Archive (JAR) files
`jarsigner`	Signs and verifies signatures on JAR files

Table 1-4 (Cont.) Java Security Tools

Tool	Usage
`keytool`	Creates and manages key stores

There are also three Kerberos-related tools that are shipped with the JDK for Windows. Equivalent functionality is provided in tools of the same name that are automatically part of the Solaris and Linux operating environments.

Table 1-5 Kerberos-related Tools

Tool	Usage
`kinit`	Obtains and caches Kerberos ticket-granting tickets
`klist`	Lists entries in the local Kerberos credentials cache and key table
`ktab`	Manages the names and service keys stored in the local Kerberos key table

Built-In Providers

The Java SE implementation from Oracle includes a number of built-in provider packages. See JDK Providers Documentation.

Java SE Platform Security Architecture

This document gives an overview of the motivation of the major security features implemented for the JDK, describes the classes that are part of the Java security architecture, discusses the impact of this architecture on existing code, and gives thoughts on writing security-sensitive code.

Introduction

Since the inception of Java technology, there has been strong and growing interest around the security of the Java platform as well as new security issues raised by the deployment of Java technology.

From a technology provider's point of view, Java security includes two aspects:

- Provide the Java platform as a secure, ready-built platform on which to run Java-enabled applications in a secure fashion.
- Provide security tools and services implemented in the Java programming language that enable a wider range of security-sensitive applications, for example, in the enterprise world.

This document discusses issues related to the first aspect, where the customers for such technologies include vendors that bundle or embed Java technology in their products (such as browsers and operating systems).

The Original Sandbox Model

The original security model provided by the Java platform is known as the sandbox model, which existed in order to provide a very restricted environment in which to run untrusted code obtained from the open network. The essence of the sandbox model is that local code is trusted to have full access to vital system resources (such as the file system) while downloaded remote code (an applet) is not trusted and can access only the limited resources provided inside the sandbox. This sandbox model is illustrated in the figure below.

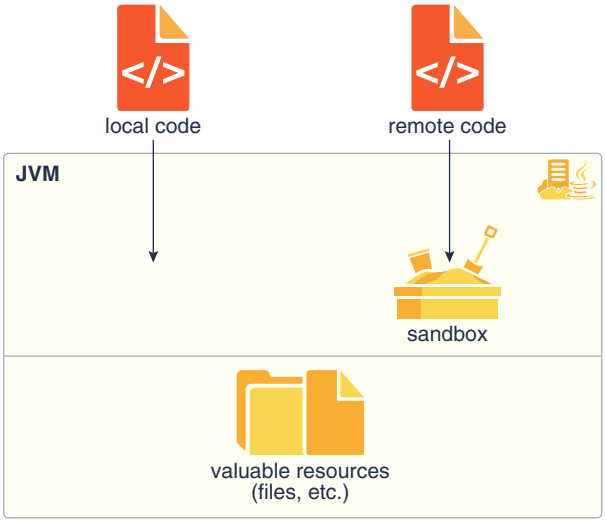

The sandbox model was deployed through the Java Development Kit (JDK), and was generally adopted by applications built with JDK 1.0, including Java-enabled web browsers.

Overall security is enforced through a number of mechanisms. First of all, the language is designed to be type-safe and easy to use. The hope is that the burden on the programmer is such that the likelihood of making subtle mistakes is lessened compared with using other programming languages such as C or C++. Language features such as automatic memory management, garbage collection, and range checking on strings and arrays are examples of how the language helps the programmer to write safe code.

Second, compilers and a bytecode verifier ensure that only legitimate Java bytecodes are executed. The bytecode verifier, together with the Java Virtual Machine, guarantees language safety at run time.

Moreover, a classloader defines a local name space, which can be used to ensure that an untrusted applet cannot interfere with the running of other programs.

Finally, access to crucial system resources is mediated by the Java Virtual Machine and is checked in advance by a `SecurityManager` class that restricts the actions of a piece of untrusted code to the bare minimum.

JDK 1.1 introduced the concept of a "signed applet", as illustrated by the figure below. In that release, a correctly digitally signed applet is treated as if it is trusted local code if the signature key is recognized as trusted by the end system that receives the applet. Signed applets, together with their signatures, are delivered in the JAR (Java Archive) format. In JDK 1.1, unsigned applets still run in the sandbox.

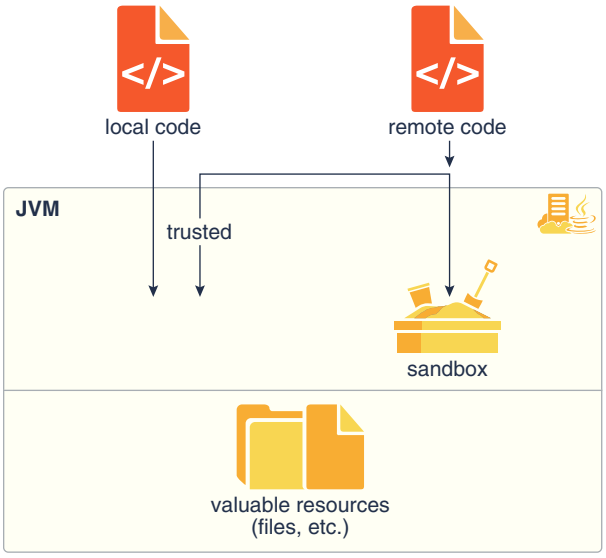

Evolving the Sandbox Model

The new Java SE Platform Security Architecture, illustrated in the figure below, is introduced primarily for the following purposes.

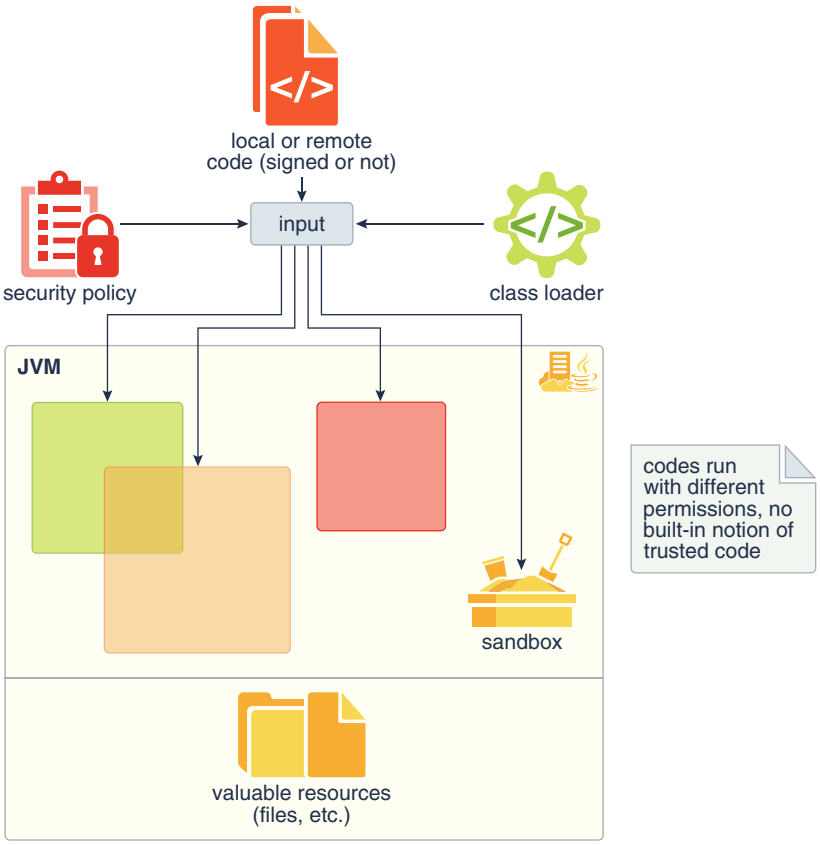

- Fine-grained access control.

  This capability existed in the JDK from the beginning, but to use it, the application writer had to do substantial programming (e.g., by subclassing and customizing the `SecurityManager` and `ClassLoader` classes). The HotJava browser 1.0 is

such an application, as it allows the browser user to choose from a small number of different security levels.

However, such programming is extremely security-sensitive and requires sophisticated skills and in-depth knowledge of computer security. The new architecture will make this exercise simpler and safer.

- Easily configurable security policy.

  Once again, this capability existed previously in the JDK but was not easy to use. Moreover, writing security code is not straightforward, so it is desirable to allow application builders and users to configure security policies without having to program.

- Easily extensible access control structure.

  Up to JDK 1.1, in order to create a new access permission, you had to add a new `check` method to the `SecurityManager` class. The new architecture allows typed permissions (each representing an access to a system resource) and automatic handling of all permissions (including yet-to-be-defined permissions) of the correct type. No new method in the `SecurityManager` class needs to be created in most cases. (In fact, we have so far not encountered a situation where a new method must be created.)

- Extension of security checks to all Java programs, including applications as well as applets.

  There is no longer a built-in concept that all local code is trusted. Instead, local code (e.g., non-system code, application packages installed on the local file system) is subjected to the same security control as applets, although it is possible, if desired, to declare that the policy on local code (or remote code) be the most liberal, thus enabling such code to effectively run as totally trusted. The same principle applies to signed applets and any Java application.

  Finally, an implicit goal is to make internal adjustment to the design of security classes (including the `SecurityManager` and `ClassLoader` classes) to reduce the risks of creating subtle security holes in future programming.

Protection Mechanisms – Overview of Basic Concepts

We now go over, in some detail, the new protection architecture and give a brief explanation of its functionality. We start with an overview of the basic concepts behind the new architecture. We then introduce the major new classes in a natural order, starting with permission specifications, going on to the policy and related features, followed by access control and its usage, and then covering secure class loading and resolution.

A fundamental concept and important building block of system security is the protection domain [Saltzer and Schroeder 75]. A domain can be scoped by the set of objects that are currently directly accessible by a principal, where a principal is an entity in the computer system to which permissions (and as a result, accountability) are granted. The sandbox utilized in JDK 1.0 is one example of a protection domain with a fixed boundary.

The protection domain concept serves as a convenient mechanism for grouping and isolation between units of protection. For example, it is possible (but not yet provided as a built-in feature) to separate protection domains from interacting with each other so that any permitted interaction must be either through trusted system code or

explicitly allowed by the domains concerned. Note that existing object accessibility rules remain valid under the new security architecture.

Protection domains generally fall into two distinct categories: system domain and application domain. It is important that all protected external resources, such as the file system, the networking facility, and the screen and keyboard, be accessible only via system domains. The figure below illustrates the domain composition of a Java application environment.

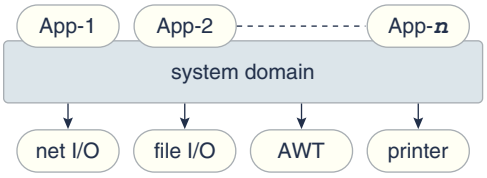

A domain conceptually encloses a set of classes whose instances are granted the same set of permissions. Protection domains are determined by the policy currently in effect. The Java application environment maintains a mapping from code (classes and instances) to their protection domains and then to their permissions, as illustrated by the figure below.

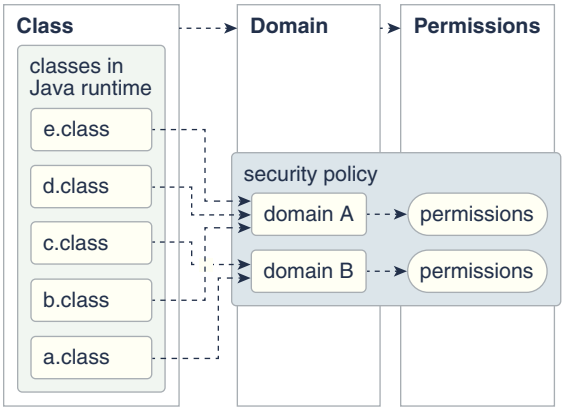

A thread of execution (which is often, but not necessarily tied to, a single Java thread, which in turn is not necessarily tied to the thread concept of the underlying operation system) may occur completely within a single protection domain or may involve an application domain and also the system domain. For example, an application that prints a message out will have to interact with the system domain that is the only access point to an output stream. In this case, it is crucial that at any time the application domain does not gain additional permissions by calling the system domain. Otherwise, there can be serious security implications.

In the reverse situation where a system domain invokes a method from an application domain, such as when the AWT system domain calls an applet's paint method to display the applet, it is again crucial that at any time the effective access rights are the same as current rights enabled in the application domain.

In other words, a less "powerful" domain cannot gain additional permissions as a result of calling or being called by a more powerful domain.

This discussion of one thread involving two protection domains naturally generalizes to a thread that traverses multiple protection domains. A simple and prudent rule of thumb for calculating permissions is the following:

- The permission set of an execution thread is considered to be the intersection of the permissions of all protection domains traversed by the execution thread.

- When a piece of code calls the `doPrivileged` method (see below), the permission set of the execution thread is considered to include a permission if it is allowed by the said code's protection domain and by all protection domains that are called or entered directly or indirectly subsequently.

As you can see, the `doPrivileged` method enables a piece of trusted code to temporarily enable access to more resources than are available directly to the application that called it. This is necessary in some situations. For example, an application may not be allowed direct access to files that contain fonts, but the system utility to display a document must obtain those fonts, on behalf of the user. We provide the `doPrivileged` method for the system domain to deal with this situation, and the method is in fact available to all domains.

During execution, when access to a critical system resource (such as file I/O and network I/O) is requested, the resource-handling code directly or indirectly invokes a special `AccessController` class method that evaluates the request and decides if the request should be granted or denied.

Such an evaluation follows and generalizes the "rule of thumb" given above. The actual way in which the evaluation is conducted can vary between implementations. The basic principle is to examine the call history and the permissions granted to the relevant protection domains, and to return silently if the request is granted or throw a security exception if the request is denied.

Finally, each domain (system or application) may also implement additional protection of its internal resources within its own domain boundary. For example, a banking application may need to support and protect internal concepts such as checking accounts, deposits and withdrawals. Because the semantics of such protection is unlikely to be predictable or enforceable by the JDK, the protection system at this level is best left to the system or application developers. Nevertheless, whenever appropriate, we provide helpful primitives to simplify developers' tasks. One such primitive is the `SignedObject` class, whose detail we will describe later.

Permissions and Security Policy

The Permission Classes

The permission classes represent access to system resources. The `java.security.Permission` class is an abstract class and is subclassed, as appropriate, to represent specific accesses.

As an example of a permission, the following code can be used to produce a permission to read the file named `abc` in the `/tmp` directory:

```
perm = new java.io.FilePermission("/tmp/abc", "read");
```

New permissions are subclassed either from the `Permission` class or one of its subclasses, such as `java.security.BasicPermission`. Subclassed permissions (other than `BasicPermission`) generally belong to their own packages. Thus, `FilePermission` is found in the `java.io` package.

A crucial abstract method that needs to be implemented for each new class of permission is the `implies` method. Basically, "a implies b" means that if one is granted permission "a", one is naturally granted permission "b". This is important when making access control decisions.

Associated with the abstract class `java.security.Permission` are the abstract class named `java.security.PermissionCollection` and the final class `java.security.Permissions`.

Class `java.security.PermissionCollection` represents a collection (i.e., a set that allows duplicates) of `Permission` objects for a single category (such as file permissions), for ease of grouping. In cases where permissions can be added to the `PermissionCollection` object in any order, such as for file permissions, it is crucial that the `PermissionCollection` object ensure that the correct semantics are followed when the `implies` function is called.

Class `java.security.Permissions` represents a collection of collections of `Permission` objects, or in other words, a super collection of heterogeneous permissions.

Applications are free to add new categories of permissions that the system supports. How to add such application-specific permissions is discussed later in this document.

Now we describe the syntax and semantics of all built-in permissions.

java.security.Permission

This abstract class is the ancestor of all permissions. It defines the essential functionalities required for all permissions.

Each permission instance is typically generated by passing one or more string parameters to the constructor. In a common case with two parameters, the first parameter is usually "the name of the target" (such as the name of a file for which the permission is aimed), and the second parameter is the action (such as "read" action on a file). Generally, a set of actions can be specified together as a comma-separated composite string.

java.security.PermissionCollection

This class holds a homogeneous collection of permissions. In other words, each instance of the class holds only permissions of the same type.

java.security.Permissions

This class is designed to hold a heterogeneous collection of permissions. Basically, it is a collection of `java.security.PermissionCollection` objects.

java.security.UnresolvedPermission

Recall that the internal state of a security policy is normally expressed by the permission objects that are associated with each code source. Given the dynamic nature of Java technology, however, it is possible that when the policy is initialized the actual code that implements a particular permission class has not yet been loaded and defined in the Java application environment. For example, a referenced permission class may be in a JAR file that will later be loaded.

The `UnresolvedPermission` class is used to hold such "unresolved" permissions. Similarly, the class `java.security.UnresolvedPermissionCollection` stores a collection of `UnresolvedPermission` permissions.

During access control checking on a permission of a type that was previously unresolved, but whose class has since been loaded, the unresolved permission is

"resolved" and the appropriate access control decision is made. That is, a new object of the appropriate class type is instantiated, if possible, based on the information in the `UnresolvedPermission`. This new object replaces the `UnresolvedPermission`, which is removed. If the permission is still unresolvable at this time, the permission is considered invalid, as if it is never granted in a security policy.

java.io.FilePermission

The targets for this class can be specified in the following ways, where directory and file names are strings that cannot contain white spaces.

```
file
directory (same as directory/)
directory/file
directory/* (all files in this directory)
* (all files in the current directory)
directory/- (all files in the file system under this directory)
- (all files in the file system under the current directory)
"<<ALL FILES>>" (all files in the file system)
```

Note that `<<ALL FILES>>` is a special string denoting all files in the system. On a Solaris, Linux, or macOS system, this includes all files under the root directory. On a Windows system, this includes all files on all drives.

The actions are: **read**, **write**, **delete**, and **execute**. Therefore, the following are valid code samples for creating file permissions:

```
import java.io.FilePermission;

FilePermission p = new FilePermission("myfile", "read,write");
FilePermission p = new FilePermission("/home/gong/", "read");
FilePermission p = new FilePermission("/tmp/mytmp", "read,delete");
FilePermission p = new FilePermission("/bin/*", "execute");
FilePermission p = new FilePermission("*", "read");
FilePermission p = new FilePermission("/-", "read,execute");
FilePermission p = new FilePermission("-", "read,execute");
FilePermission p = new FilePermission("<<ALL FILES>>", "read");
```

The `implies` method in this class correctly interprets the file system. For example, `FilePermission("/-", "read,execute")` implies `FilePermission("/home/gong/public_html/index.html", "read")`, and `FilePermission("bin/*", "execute")` implies `FilePermission("bin/emacs19.31", "execute")`.

Note:

Most of these strings are given in platform-dependent format. For example, to represent read access to the file named `foo` in the `temp` directory on the `C` drive of a Windows system, you would use

```
FilePermission p = new FilePermission("c:\\temp\\foo", "read");
```

The double backslashes are necessary to represent a single backslash because the strings are processed by a tokenizer (`java.io.StreamTokenizer`), which allows `\` to be used as an escape string (e.g., `\n` to indicate a new line) and which thus requires two backslashes to indicate a single backslash. After the tokenizer has processed the above `FilePermission` target string, converting double backslashes to single backslashes, the end result is the actual path:

```
"c:\temp\foo"
```

It is necessary that the strings be given in platform-dependent format until there is a universal file description language. Note also that the use of meta symbols such as `*` and `-` prevents the use of specific file names. We think this is a small limitation that can be tolerated for the moment. Finally, note that `-/` and `<<ALL FILES>>` are the same target on Solaris, Linux, and macOS systems in that they both refer to the entire file system. (They can refer to multiple file systems if they are all available). The two targets are potentially different on other operating systems, such as Windows and macOS.

Also note that a target name that specifies just a directory, with a "read" action, as in

```
FilePermission p = new FilePermission("/home/gong/", "read");
```

means you are only giving permission to list the files in that directory, not read any of them. To allow read access to files, you must specify either an explicit file name, or an `*` or `-`, as in

```
FilePermission p = new FilePermission("/home/gong/myfile", "read");
FilePermission p = new FilePermission("/home/gong/*", "read");
FilePermission p = new FilePermission("/home/gong/-", "read");
```

And finally, note that code always automatically has permission to read files from its same (URL) location, and subdirectories of that location; it does not need explicit permission to do so.

java.net.SocketPermission

This class represents access to a network via sockets. The target for this class can be given as *hostname*:*port_range*, where *hostname* can be given in the following ways:

```
hostname (a single host)
IP address (a single host)
localhost (the local machine)
"" (equivalent to "localhost")
hostname.domain (a single host within the domain)
hostname.subdomain.domain
*.domain (all hosts in the domain)
*.subdomain.domain
* (all hosts)
```

That is, the host is expressed as a DNS name, as a numerical IP address, as `"localhost"` (for the local machine) or as `""` (which is equivalent to specifying `"localhost"`).

The wildcard `*` may be included once in a DNS name host specification. If it is included, it must be in the leftmost position, as in `*.sun.com`.

The *port_range* can be given as follows:

```
N (a single port)
N- (all ports numbered N and above)
-N (all ports numbered N and below)
N1-N2 (all ports between N1 and N2, inclusive)
```

Here `N`, `N1`, and `N2` are non-negative integers ranging from 0 to 65535 (2^{16-1}).

The actions on sockets are **accept**, **connect**, **listen**, and **resolve** (which is basically DNS lookup). Note that implicitly, the action "resolve" is implied by "accept", "connect", and "listen" – i.e., those who can listen or accept incoming connections from or initiate out-going connections to a host should be able to look up the name of the remote host.

Below are some examples of socket permissions.

```
import java.net.SocketPermission;

SocketPermission p = new SocketPermission("java.example.com","accept");
p = new SocketPermission("192.0.2.99","accept");
p = new SocketPermission("*.com","connect");
p = new SocketPermission("*.example.com:80","accept");
p = new SocketPermission("*.example.com:-1023","accept");
p = new SocketPermission("*.example.com:1024-","connect");
p = new SocketPermission("java.example.com:8000-9000",
         "connect,accept");
p = new SocketPermission("localhost:1024-",
          "accept,connect,listen");
```

> **Note:**
>
> `SocketPermission("java.example.com:80,8080","accept")` and `SocketPermission("java.example.com,javasun.example.com","accept")` are not valid socket permissions.
>
> Moreover, because **listen** is an action that applies only to ports on the local host, whereas **accept** is an action that applies to ports on both the local and remote host, both actions are necessary.

java.security.BasicPermission

The `BasicPermission` class extends the `Permission` class. It can be used as the base class for permissions that want to follow the same naming convention as `BasicPermission` (see below).

The name for a `BasicPermission` is the name of the given permission (for example, "exitVM", "setFactory", "queuePrintJob", etc). The naming convention follows the hierarchical property naming convention. An asterisk may appear at the end of the name, following a ".", or by itself, to signify a wildcard match. For example: "java.*" or "*" is valid, "*java" or "a*b" is not valid.

The action string (inherited from `Permission`) is unused. Thus, `BasicPermission` is commonly used as the base class for "named" permissions (ones that contain a name but no actions list; you either have the named permission or you don't.) Subclasses may implement actions on top of `BasicPermission`, if desired.

Some of the `BasicPermission` subclasses are `java.lang.RuntimePermission`, `java.security.SecurityPermission`, `java.util.PropertyPermission`, and `java.net.NetPermission`.

java.util.PropertyPermission

The targets for this class are basically the names of Java properties as set in various property files. Examples are the `java.home` and `os.name` properties. Targets can be specified as "*" (any property), "a.*" (any property whose name has a prefix "a."), "a.b.*", and so on. Note that the wildcard can occur only once and can only be at the rightmost position.

This is one of the `BasicPermission` subclasses that implements actions on top of `BasicPermission`. The actions are read and write. Their meaning is defined as follows: "read" permission allows the `getProperty` method in `java.lang.System` to be called to get the property value, and "write" permission allows the `setProperty` method to be called to set the property value.

java.lang.RuntimePermission

The target for a `RuntimePermission` can be represented by any string, and there is no action associated with the targets. For example, `RuntimePermission("exitVM")` denotes the permission to exit the Java Virtual Machine.

The target names are:

```
createClassLoader
getClassLoader
setContextClassLoader
setSecurityManager
createSecurityManager
exitVM
setFactory
setIO
modifyThread
stopThread
modifyThreadGroup
getProtectionDomain
readFileDescriptor
writeFileDescriptor
loadLibrary.{library name}
accessClassInPackage.{package name}
defineClassInPackage.{package name}
accessDeclaredMembers.{class name}
queuePrintJob
```

java.awt.AWTPermission

This is in the same spirit as the `RuntimePermission`; it's a permission without actions. The targets for this class are:

```
accessClipboard
accessEventQueue
listenToAllAWTEvents
showWindowWithoutWarningBanner
```

java.net.NetPermission

This class contains the following targets and no actions:

```
requestPasswordAuthentication
setDefaultAuthenticator
specifyStreamHandler
```

java.lang.reflect.ReflectPermission

This is the `Permission` class for reflective operations. A `ReflectPermission` is a named permission (like `RuntimePermission`) and has no actions. The only name currently defined is `suppressAccessChecks`, which allows suppressing the standard Java programming language access checks – for public, default (package) access, protected, and private members – performed by reflected objects at their point of use.

java.io.SerializablePermission

This class contains the following targets and no actions:

```
enableSubclassImplementation
enableSubstitution
```

java.security.SecurityPermission

`SecurityPermissions` control access to security-related objects, such as `Security`, `Policy`, `Provider`, `Signer`, and `Identity` objects. This class contains the following targets and no actions:

```
getPolicy
setPolicy
getProperty.{key}
setProperty.{key}
insertProvider.{provider name}
removeProvider.{provider name}
setSystemScope
setIdentityPublicKey
setIdentityInfo
printIdentity
addIdentityCertificate
removeIdentityCertificate
clearProviderProperties.{provider name}
putProviderProperty.{provider name}
removeProviderProperty.{provider name}
getSignerPrivateKey
setSignerKeyPair
```

java.security.AllPermission

This permission implies all permissions. It is introduced to simplify the work of system administrators who might need to perform multiple tasks that require all (or numerous) permissions. It would be inconvenient to require the security policy to iterate through all permissions. Note that `AllPermission` also implies new permissions that are defined in the future.

Clearly, much caution is necessary when considering granting this permission.

javax.security.auth.AuthPermission

`AuthPermission` handles authentication permissions and authentication-related object such as `Subject`, `SubjectDomainCombiner`, `LoginContext`, and `Configuration`. This class contains the following targets and no actions:

```
doAs
doAsPrivileged
getSubject
getSubjectFromDomainCombiner
setReadOnly
```

```
modifyPrincipals
modifyPublicCredentials
modifyPrivateCredentials
refreshCredential
destroyCredential
createLoginContext.{name}
getLoginConfiguration
setLoginConfiguration
refreshLoginConfiguration
```

Discussion of Permission Implications

Recall that permissions are often compared against each other, and to facilitate such comparisons, we require that each permission class defines an `implies` method that represents how the particular permission class relates to other permission classes. For example, `java.io.FilePermission("/tmp/*", "read")` implies `java.io.FilePermission("/tmp/a.txt", "read")` but does not imply any `java.net.NetPermission`.

There is another layer of implication that may not be immediately obvious to some readers. Suppose that one applet has been granted the permission to write to the entire file system. This presumably allows the applet to replace the system binary, including the JVM runtime environment. This effectively means that the applet has been granted all permissions.

Another example is that if an applet is granted the runtime permission to create class loaders, it is effectively granted many more permissions, as a class loader can perform sensitive operations.

Other permissions that are "dangerous" to give out include those that allow the setting of system properties, runtime permissions for defining packages and for loading native code libraries (because the Java security architecture is not designed to and does not prevent malicious behavior at the level of native code), and of course the `AllPermission`.

For more information about permissions, including tables enumerating the risks of assigning specific permissions as well as a table of all the JDK built-in methods that require permissions, see Permissions in the JDK.

How To Create New Types of Permissions

It is essential that no one except Oracle should extend the permissions that are built into the JDK, either by adding new functionality or by introducing additional target keywords into a class such as `java.lang.RuntimePermission`. This maintains consistency.

To create a new permission, the following steps are recommended, as shown by an example. Suppose an application developer from company ABC wants to create a customized permission to "watch TV".

First, create a new class `com.abc.Permission` that extends the abstract class `java.security.Permission` (or one of its subclasses), and another new class `com.abc.TVPermission` that extends the `com.abc.Permission`. Make sure that the `implies` method, among others, is correctly implemented. (Of course,

`com.abc.TVPermission` can directly extend `java.security.Permission`; the intermediate `com.abc.Permission` is not required.)

```
public class com.abc.Permission extends java.security.Permission

public class com.abc.TVPermission extends com.abc.Permission
```

The following figure shows the subclass relationship.

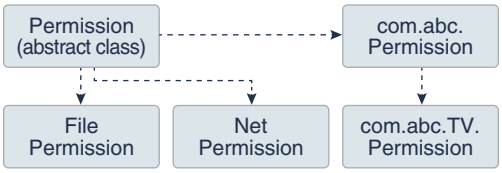

Second, include these new classes with the application package.

Each user that wants to allow this new type of permission for specific code does so by adding an entry in a policy file. (Details of the policy file syntax are given in a later section.) An example of a policy file entry granting code from `http://example.com/` permission to watch channel 5 would be:

```
grant codeBase  "http://example.com/" {
    permission com.abc.TVPermission "channel-5", "watch";
}
```

In the application's resource management code, when checking to see if a permission should be granted, call `AccessController`'s `checkPermission` method using a `com.abc.TVPermission` object as the parameter.

```
com.abc.TVPermission tvperm = new
     com.abc.TVPermission("channel-5", "watch");
AccessController.checkPermission(tvperm);
```

Note that, when adding a new permission, one should create a new (permission) class and not add a new method to the security manager. (In the past, in order to enable checking of a new type of access, you had to add a new method to the `SecurityManager` class.)

If more elaborate `TVPermissions` such as "channel-1:13" or "channel-*" are allowed, then it may be necessary to implement a `TVPermissionCollection` object that knows how to deal with the semantics of these pseudo names.

New code should always invoke a permission check by calling the `checkPermission` method of the `AccessController` class in order to exercise the built-in access control algorithm. There is no essential need to examine whether there is a `ClassLoader` or a `SecurityManager`. On the other hand, if the algorithm should be left to the installed security manager class, then the method `SecurityManager.checkPermission` should be invoked instead.

java.security.CodeSource

This class extends the concept of a codebase within HTML to encapsulate not only the code location (URL) but also the certificate(s) containing public keys that should be used to verify signed code originating from that location. Note that this is not the equivalent of the `CodeBase` tag in HTML files. Each certificate is represented as a `java.security.cert.Certificate`, and each URL as a `java.net.URL`.

java.security.Policy

The system security policy for a Java application environment, specifying which permissions are available for code from various sources, is represented by a `Policy` object. More specifically, it is represented by a `Policy` subclass providing an implementation of the abstract methods in the `Policy` class.

In order for an applet (or an application running under a `SecurityManager`) to be allowed to perform secured actions, such as reading or writing a file, the applet (or application) must be granted permission for that particular action. The only exception is that code always automatically has permission to read files from its same `CodeSource`, and subdirectories of that `CodeSource`; it does not need explicit permission to do so.

There could be multiple instances of the `Policy` object, although only one is "in effect" at any time. The currently-installed `Policy` object can be obtained by calling the `getPolicy` method, and it can be changed by a call to the `setPolicy` method (by code with permission to reset the Policy).

The source location for the policy information utilized by the `Policy` object is up to the Policy implementation. The policy configuration may be stored, for example, as a flat ASCII file, as a serialized binary file of the `Policy` class, or as a database. There is a `Policy` reference implementation that obtains its information from static policy configuration files.

Policy File Format

In the Policy reference implementation, the policy can be specified within one or more policy configuration files. The configuration files indicate what permissions are allowed for code from specified code sources. Each configuration file must be encoded in UTF-8.

A policy configuration file essentially contains a list of entries. It may contain a keystore entry, and contains zero or more grant entries.

A keystore is a database of private keys and their associated digital certificates such as X.509 certificate chains authenticating the corresponding public keys. The `keytool` utility is used to create and administer keystores. The keystore specified in a policy configuration file is used to look up the public keys of the signers specified in the grant entries of the file. A keystore entry must appear in a policy configuration file if any grant entries specify signer aliases, or if any grant entries specify a principal alias (see below).

At this time, there can be only one keystore entry in the policy file (others after the first one are ignored), and it can appear anywhere outside the file's grant entries . It has the following syntax:

```
keystore "some_keystore_url", "keystore_type";
```

Here, *some_keystore_url* specifies the URL location of the keystore, and *keystore_type* specifies the keystore type. The latter is optional. If not specified, the type is assumed to be that specified by the *keystore.type* property in the security properties file.

The URL is relative to the policy file location. Thus if the policy file is specified in the security properties file as:

```
policy.url.1=http://foo.bar.example.com/blah/some.policy
```

and that policy file has an entry:

```
keystore ".keystore";
```

then the keystore will be loaded from:

```
http://foo.bar.example.com/blah/.keystore
```

The URL can also be absolute.

A keystore type defines the storage and data format of the keystore information, and the algorithms used to protect private keys in the keystore and the integrity of the keystore itself. The Oracle JDK's default keystore type is PKCS12.

Each grant entry in a policy file essentially consists of a `CodeSource` and its permissions. Actually, a `CodeSource` consists of a URL and a set of certificates, while a policy file entry includes a URL and a list of signer names. The system creates the corresponding `CodeSource` after consulting the keystore to determine the certificate(s) of the specified signers.

Each grant entry in the policy file is of the following format, where the leading `grant` is a reserved word that signifies the beginning of a new entry and optional items appear in brackets. Within each entry, a leading `permission` is another reserved word that marks the beginning of a new permission in the entry. Each grant entry grants a set of permissions to a specified code source and principals.

```
grant [SignedBy "signer_names"] [, CodeBase "URL"]
      [, Principal [principal_class_name] "principal_name"]
      [, Principal [principal_class_name] "principal_name"] ... {
    permission permission_class_name [ "target_name" ]
               [, "action"] [, SignedBy "signer_names"];
    permission ...
};
```

White spaces are allowed immediately before or after any comma. The name of the permission class must be a fully qualified class name, such as

`java.io.FilePermission`, and cannot be abbreviated (for example, to `FilePermission`).

Note that the action field is optional in that it can be omitted if the permission class does not require it. If it is present, then it must come immediately after the target field.

The exact meaning of a CodeBase URL value depends on the characters at the end. A CodeBase with a trailing "/" matches all class files (not JAR files) in the specified directory. A CodeBase with a trailing "/*" matches all files (both class and JAR files) contained in that directory. A CodeBase with a trailing "/-" matches all files (both class and JAR files) in the directory and recursively all files in subdirectories contained in that directory.

The CodeBase field (URL) is optional in that, if it is omitted, it signifies "any code base".

The first signer name field is a string alias that is mapped, via a separate mechanism, to a set of public keys (within certificates in the keystore) that are associated with the signers. These keys are used to verify that certain signed classes are really signed by these signers.

This signer field can be a comma-separated string containing names of multiple signers, an example of which is `Adam,Eve,Charles`, which means signed by Adam and Eve and Charles (i.e., the relationship is AND, not OR).

This field is optional in that, if it is omitted, it signifies "any signer", or in other words, "It doesn't matter whether the code is signed or not".

The second signer field, inside a permission entry, represents the alias to the keystore entry containing the public key corresponding to the private key used to sign the bytecodes that implemented the said permission class. This permission entry is effective (i.e., access control permission will be granted based on this entry) only if the bytecode implementation is verified to be correctly signed by the said alias.

A principal value specifies a class_name/principal_name pair which must be present within the executing threads principal set. The principal set is associated with the executing code by way of a Subject. The principal field is optional in that, if it is omitted, it signifies "any principals".

> **Note:**
>
> Regarding keystore alias replacement: If the principal class_name/principal_name pair is specified as a single quoted string, it is treated as a keystore alias. The keystore is consulted and queried (via the alias) for an X509 Certificate. If one is found, the principal_class is automatically treated as `javax.security.auth.x500.X500Principal`, and the principal_name is automatically treated as the subject distinguished name from the certificate. If an X509 Certificate mapping is not found, the entire grant entry is ignored.

The order between the CodeBase, SignedBy, and Principal fields does not matter.

An informal BNF grammar for the policy file format is given below, where non-capitalized terms are terminals:

```
PolicyFile -> PolicyEntry | PolicyEntry; PolicyFile
PolicyEntry -> grant {PermissionEntry}; |
           grant SignerEntry {PermissionEntry} |
           grant CodebaseEntry {PermissionEntry} |
           grant PrincipalEntry {PermissionEntry} |
           grant SignerEntry, CodebaseEntry {PermissionEntry} |
           grant CodebaseEntry, SignerEntry {PermissionEntry} |
           grant SignerEntry, PrincipalEntry {PermissionEntry} |
           grant PrincipalEntry, SignerEntry {PermissionEntry} |
           grant CodebaseEntry, PrincipalEntry {PermissionEntry} |
           grant PrincipalEntry, CodebaseEntry {PermissionEntry} |
           grant SignerEntry, CodebaseEntry, PrincipalEntry
{PermissionEntry} |
           grant CodebaseEntry, SignerEntry, PrincipalEntry
{PermissionEntry} |
           grant SignerEntry, PrincipalEntry, CodebaseEntry
{PermissionEntry} |
           grant CodebaseEntry, PrincipalEntry, SignerEntry
{PermissionEntry} |
           grant PrincipalEntry, CodebaseEntry, SignerEntry
{PermissionEntry} |
           grant PrincipalEntry, SignerEntry, CodebaseEntry
{PermissionEntry} |
           keystore "url"
SignerEntry -> signedby (a comma-separated list of strings)
CodebaseEntry -> codebase (a string representation of a URL)
PrincipalEntry -> OnePrincipal | OnePrincipal, PrincipalEntry
OnePrincipal -> principal [ principal_class_name ] "principal_name" (a
principal)
PermissionEntry -> OnePermission | OnePermission PermissionEntry
OnePermission -> permission permission_class_name
                 [ "target_name" ] [, "action_list"]
                 [, SignerEntry];
```

Now we give some examples. The following policy grants permission `a.b.Foo` to code signed by `Roland`:

```
grant signedBy "Roland" {
    permission a.b.Foo;
};
```

The following grants a `FilePermission` to all code (regardless of the signer and/or CodeBase):

```
grant {
   permission java.io.FilePermission ".tmp", "read";
};
```

The following grants two permissions to code that is signed by both `Li` and `Roland`:

```
grant signedBy "Roland,Li" {
  permission java.io.FilePermission "/tmp/*", "read";
  permission java.util.PropertyPermission "user.*";
};
```

The following grants two permissions to code that is signed by `Li` and that comes from `http://example.com`:

```
grant codeBase "http://example.com/*", signedBy "Li" {
    permission java.io.FilePermission "/tmp/*", "read";
    permission java.io.SocketPermission "*", "connect";
};
```

The following grants two permissions to code that is signed by both `Li` and `Roland`, and only if the bytecodes implementing `com.abc.TVPermission` are genuinely signed by `Li`.

```
grant signedBy "Roland,Li" {
  permission java.io.FilePermission "/tmp/*", "read";
  permission com.abc.TVPermission "channel-5", "watch",
     signedBy "Li";
};
```

The reason for including the second signer field is to prevent spoofing when a permission class does not reside with the Java runtime installation. For example, a copy of the `com.abc.TVPermission` class can be downloaded as part of a remote JAR archive, and the user policy might include an entry that refers to it. Because the archive is not long-lived, the second time the `com.abc.TVPermission` class is downloaded, possibly from a different web site, it is crucial that the second copy is authentic, as the presence of the permission entry in the user policy might reflect the user's confidence or belief in the first copy of the class bytecode.

The reason we chose to use digital signatures to ensure authenticity, rather than storing (a hash value of) the first copy of the bytecodes and using it to compare with the second copy, is because the author of the permission class can legitimately update the class file to reflect a new design or implementation.

Note:

The strings for a file path must be specified in a platform-dependent format; this is necessary until there is a universal file description language. The above examples have shown strings appropriate on Solaris. On Windows, when you directly specify a file path in a string, you need to include two backslashes for each actual single backslash in the path, as in

```
grant signedBy "Roland" {
  permission java.io.FilePermission "C:\\users\\Cathy\\*", "read";
};
```

This is because the strings are processed by a tokenizer (`java.io.StreamTokenizer`), which allows "\" to be used as an escape string (e.g., "\n" to indicate a new line) and which thus requires two backslashes to indicate a single backslash. After the tokenizer has processed the above `FilePermission` target string, converting double backslashes to single backslashes, the end result is the actual path:

```
"C:\users\Cathy\*"
```

Finally, here are some principal-based grant entries:

```
grant principal javax.security.auth.x500.X500Principal "cn=Alice" {
  permission java.io.FilePermission "/home/Alice", "read, write";
};
```

This permits any code executing as the `X500Principal`, `cn=Alice`, permission to read and write to `/home/Alice`.

The following example shows a grant statement with both codesource and principal information.

```
grant codebase "http://www.games.example.com",
      signedBy "Duke",
      principal javax.security.auth.x500.X500Principal "cn=Alice" {
  permission java.io.FilePermission "/tmp/games", "read, write";
};
```

This allows code downloaded from `www.games.example.com`, signed by `Duke`, and executed by `cn=Alice`, permission to read and write into the `/tmp/games` directory.

he following example shows a grant statement with KeyStore alias replacement:

```
keystore "http://foo.bar.example.com/blah/.keystore";

grant principal "alice" {
  permission java.io.FilePermission "/tmp/games", "read, write";
};
```

`alice` will be replaced by `javax.security.auth.x500.X500Principal` `cn=Alice` assuming the X.509 certificate associated with the keystore alias, *alice*, has a subject distinguished name of `cn=Alice`. This allows code executed by the `X500Principal` `cn=Alice` permission to read and write into the `/tmp/games` directory.

Property Expansion in Policy Files

Property expansion is possible in policy files and in the security properties file. Property expansion is similar to expanding variables in a shell. That is, when a string like `${some.property}` appears in a policy file, or in the security properties file, it will be expanded to the value of the specified system property. For example,

```
permission java.io.FilePermission "${user.home}", "read";
```

will expand `${user.home}` to use the value of the `user.home` system property. If that property's value is `/home/cathy`, then the above is equivalent to

```
permission java.io.FilePermission "/home/cathy", "read";
```

In order to assist in platform-independent policy files, you can also use the special notation of `${/}`, which is a shortcut for `${file.separator}`. This allows permission designations such as

```
permission java.io.FilePermission "${user.home}${/}*", "read";
```

If `user.home` is `/home/cathy`, and you are on Solaris, the above gets converted to:

```
permission java.io.FilePermission "/home/cathy/*", "read";
```

If on the other hand `user.home` is `C:\users\cathy` and you are on a Windows system, the above gets converted to:

```
permission java.io.FilePermission "C:\users\cathy\*", "read";
```

Also, as a special case, if you expand a property in a codebase, such as

```
grant codeBase "file:/${java.home}/lib/ext/"
```

then any `file.separator` characters will be automatically converted to slashes (`/`), which is desirable since codebases are URLs. Thus on a Windows system, even if `java.home` is set to `C:\j2sdk1.2`, the above would get converted to

```
grant codeBase "file:/C:/j2sdk1.2/lib/ext/"
```

Thus you don't need to use `${/}` in codebase strings (and you shouldn't).

Property expansion takes place anywhere a double quoted string is allowed in the policy file. This includes the `signedby`, `codebase`, target names, and action fields.

Whether or not property expansion is allowed is controlled by the value of the `policy.expandProperties` property in the Security Properties file. If the value of this Security Property is true (the default), expansion is allowed.

Please note: You can't use nested properties; they will not work. For example,

```
"${user.${foo}}"
```

doesn't work, even if the `foo` property is set to `home`. The reason is the property parser doesn't recognize nested properties; it simply looks for the first `${`, and then keeps looking until it finds the first `}` and tries to interpret the result `${user.$foo}` as a property, but fails if there is no such property.

Also note: If a property can't be expanded in a grant entry, permission entry, or keystore entry, that entry is ignored. For example, if the system property `foo` is not defined and you have:

```
grant codeBase "${foo}" {
  permission ...;
  permission ...;
};
```

then all the permissions in this grant entry are ignored. If you have

```
grant {
  permission Foo "${foo}";
  permission Bar;
};
```

then only the `permission Foo "${foo}";` entry is ignored. And finally, if you have

```
keystore "${foo}";
```

then the keystore entry is ignored.

One final note: On Windows systems, when you directly specify a file path in a string, you need to include two backslashes for each actual single backslash in the path, as in

```
"C:\\users\\cathy\\foo.bat"
```

This is because the strings are processed by a tokenizer (`java.io.StreamTokenizer`), which allows the backslash (`\`) to be used as an escape string (e.g., `\n` to indicate a new line) and which thus requires two backslashes to indicate a single backslash. After the tokenizer has processed the above string, converting double backslashes to single backslashes, the end result is

```
"C:\users\cathy\foo.bat"
```

Expansion of a property in a string takes place **after** the tokenizer has processed the string. Thus if you have the string

```
"${user.home}\\foo.bat"
```

then first the tokenizer processes the string, converting the double backslashes to a single backslash, and the result is

```
"${user.home}\foo.bat"
```

Then the `${user.home}` property is expanded and the end result is

```
"C:\users\cathy\foo.bat"
```

assuming the `user.home` value is `C:\users\cathy`. Of course, for platform independence, it would be better if the string was initially specified without any explicit slashes, i.e., using the `${/}` property instead, as in

```
"${user.home}${/}foo.bat"
```

General Expansion in Policy Files

Generalized forms of expansion are also supported in policy files. For example, permission names may contain a string of the form: `${{protocol:protocol_data}}` If such a string occurs in a permission name, then the value in *`protocol`* determines the exact type of expansion that should occur, and *`protocol_data`* is used to help perform the expansion. *`protocol_data`* may be empty, in which case the above string should simply take the form:

`${{protocol}}`

There are two protocols supported in the default policy file implementation:

1. `${{self}}`

   The protocol, `self`, denotes a replacement of the entire string, `${{self}}`, with one or more principal class/name pairs. The exact replacement performed depends upon the contents of the grant clause to which the permission belongs.

   If the grant clause does not contain any principal information, the permission will be ignored (permissions containing `${{self}}` in their target names are only valid in the context of a principal-based grant clause). For example, `BarPermission` will always be ignored in the following grant clause:

   ```
   grant codebase "www.foo.example.com", signedby "duke" {
       permission BarPermission "... ${{self}} ...";
   };
   ```

   If the grant clause contains principal information, `${{self}}` will be replaced with that same principal information. For example, `${{self}}` in `BarPermission` will be

replaced by `javax.security.auth.x500.X500Principal "cn=Duke"` in the following grant clause:

```
grant principal javax.security.auth.x500.X500Principal "cn=Duke" {
    permission BarPermission "... ${{self}} ...";
};
```

If there is a comma-separated list of principals in the grant clause, then `${{self}}` will be replaced by the same comma-separated list or principals. In the case where both the principal class and name are wildcarded in the grant clause, `${{self}}` is replaced with all the principals associated with the `Subject` in the current `AccessControlContext`.
The following example describes a scenario involving both `self` and KeyStore alias replacement together:

```
keystore "http://foo.bar.example.com/blah/.keystore";

grant principal "duke" {
    permission BarPermission "... ${{self}} ...";
};
```

In the above example, "duke" will first be expanded into `javax.security.auth.x500.X500Principal "cn=Duke"` assuming the X.509 certificate associated with the KeyStore alias, `"duke"`, has a subject distinguished name of `"cn=Duke"`. Next, `${{self}}` will be replaced with the same principal information that just got expanded in the grant clause: `javax.security.auth.x500.X500Principal "cn=Duke"`.

2. `${{alias:`*`alias_name`*`}}`

   The protocol, `alias`, denotes a `java.security.KeyStore` alias substitution. The `KeyStore` used is the one specified in the KeyStore entry; see Policy File Format. *`alias_name`* represents an alias into the `KeyStore`. `${{alias:`*`alias_name`*`}}` is replaced with `javax.security.auth.x500.X500Principal "`*`DN`*`"`, where *`DN`* represents the subject distinguished name of the certificate belonging to *`alias_name`*. For example:

```
keystore "http://foo.bar.example.com/blah/.keystore";

grant codebase "www.foo.example.com" {
    permission BarPermission "... ${{alias:duke}} ...";
};
```

   In the above example the X.509 certificate associated with the alias, `duke`, is retrieved from the `KeyStore`, `foo.bar.example.com/blah/.keystore`. Assuming duke's certificate specifies `"o=dukeOrg, cn=duke"` as the subject distinguished name, then `${{alias:duke}}` is replaced with `javax.security.auth.x500.X500Principal "o=dukeOrg, cn=duke"`.
   The permission entry is ignored under the following error conditions:

   - The keystore entry is unspecified
   - The *`alias_name`* is not provided

- The certificate for *alias_name* cannot be retrieved
- The certificate retrieved is not an X.509 certificate

Assigning Permissions

When a principal executes a class that originated from a particular `CodeSource`, the security mechanism consults the policy object to determine what permissions to grant. This is done by invoking the `getPermissions` or `implies` method on the `Policy` object that is installed in the VM.

Clearly, a given code source in a `ProtectionDomain` can match the code source given in multiple entries in the policy, for example because the wildcard (`*`) is allowed.

The following algorithm is used to locate the appropriate set of permissions in the policy.

1. Match the public keys, if code is signed.
2. If a key is not recognized in the policy, then ignore the key.

   If every key is ignored, then treat the code as unsigned.
3. If the keys are matched or no signer was specified, then try to match all URLs in the policy for the keys.
4. If the keys are matched (or no signer was specified) and the URLs are matched (or no codebase was specified), then try to match all principals in the policy with the principals associated with the current executing thread.
5. If either key, URL, or principals are not matched, then use the built-in default permission, which is the original sandbox permission.

The exact meaning of a policy entry codeBase URL value depends on the characters at the end. A codeBase with a trailing `"/"` matches all class files (not JAR files) in the specified directory. A codeBase with a trailing `"/*"` matches all files (both class and JAR files) contained in that directory. A codeBase with a trailing `"/-"` matches all files (both class and JAR files) in the directory and recursively all files in subdirectories contained in that directory.

As an example, given `"http://example.com/-"` in the policy, then any code base that is on this web site matches the policy entry. Matching code bases include `"http://example.com/j2se/sdk/"` and `"http://example.com/people/gong/appl.jar"`.

If multiple entries are matched, then all the permissions given in those entries are granted. In other words, permission assignment is additive. For example, if code signed with key A gets permission X and code signed by key B gets permission Y and no particular codebase is specified, then code signed by both A and B gets permissions X and Y. Similarly, if code with codeBase `"http://example.com/-"` is given permission X, and `"http://example.com/people/*"` is given permission Y, and no particular signers are specified, then an applet from "http://example.com/people/applet.jar" gets both X and Y.

Note that URL matching here is purely syntactic. For example, a policy can give an entry that specifies a URL `"ftp://ftp.example.com"`. Such an entry is useful only when one can obtain Java code directly from ftp for execution.

To specify URLs for the local file system, a file URL can be used. For example, to specify files in the `/home/cathy/temp` directory in a Solaris system, you'd use

```
"file:/home/cathy/temp/*"
```

To specify files in the `temp` directory on the C drive in a Windows system, use

```
"file:/c:/temp/*"
```

Note: codeBase URLs always use slashes (no backlashes), regardless of the platform they apply to.

You can also use an absolute path name such as

```
"/home/gong/bin/MyWonderfulJava"
```

Default System and User Policy Files

In the Policy reference implementation, the policy can be specified within one or more policy configuration files. The configuration files specify what permissions are allowed for code from specified code sources. A policy file can be composed via a simple text editor. There is by default a single system-wide policy file, and a single user policy file. The system policy file is by default located at

- `{java.home}/conf/security/java.policy` (UNIX-based systems)
- `{java.home}\conf\security\java.policy` (Windows)

Here, `java.home` is a system property specifying the directory into which the JDK was installed. The user policy file is by default located at

- `{user.home}/.java.policy` (UNIX-based systems)
- `{user.home}\.java.policy` (Windows)

Here, `user.home` is a system property specifying the user's home directory.

When the Policy is initialized, the system policy is loaded in first, and then the user policy is added to it. If neither policy is present, a built-in policy is used. This built-in policy is the same as the original sandbox policy. Policy file locations are specified in the security properties file, which is located at

- `{java.home}/conf/security/java.security` (UNIX-based systems)
- `{java.home}\conf\security\java.security` (Windows)

The policy file locations are specified as the values of properties whose names are of the form

```
policy.url.n
```

Here, *n* is a number. You specify each such property value in a line of the following form:

```
policy.url.n=URL
```

Here, *URL* is a URL specification. For example, the default system and user policy files are defined in the security properties file as

```
policy.url.1=file:${java.home}/conf/security/java.policy
policy.url.2=file:${user.home}/.java.policy
```

You can actually specify a number of URLs, including ones of the form "http://", and all the designated policy files will get loaded. You can also comment out or change the second one to disable reading the default user policy file.

The algorithm starts at `policy.url.1`, and keeps incrementing until it does not find a URL. Thus if you have `policy.url.1` and `policy.url.3`, `policy.url.3` will never be read.

It is also possible to specify an additional or a different policy file when invoking execution of an application. This can be done via the `-Djava.security.policy` command-line argument, which sets the value of the `java.security.policy` property. For example, consider the following example:

```
java -Djava.security.manager -Djava.security.policy=pURL SomeApp
```

Here, *pURL* is a URL specifying the location of a policy file, then the specified policy file will be loaded in addition to all the policy files that are specified in the security properties file. (The `-Djava.security.manager` argument ensures that the default security manager is installed, and thus the application is subject to policy checks, as described in Managing Applets and Applications. It is not required if the application `SomeApp` installs a security manager.)

If you use the following, with a double equals sign (==), then just the specified policy file will be used; all others will be ignored.

```
java -Djava.security.manager -Djava.security.policy==pURL SomeApp
```

> **Note:**
>
> Use the double equals sign (==) with care as it overrides the built-in JDK policy file, which grants a set of default permissions that are designed to provide a secure, out-of-the-box configuration for the JDK. Overriding this policy may result in unexpected behavior (JDK code may not be granted the right permissions) and should only be done by experienced users.

> **Note:**
>
> The `-Djava.security.policy` policy file value will be ignored (for both `java` and `appletviewer` commands) if the `policy.allowSystemProperty` property in the security properties file is set to false. The default is true.

Customizing Policy Evaluation

The current design of the `Policy` class is not as comprehensive as it could be. We have given the issues much thought and are progressing cautiously, partly to ensure that we define method calls that are appropriate for the most common cases. For the meantime, an alternative policy class can be given to replace the default policy class, as long as the former is a subclass of the abstract `Policy` class and implements the `getPermissions` method (and other methods as necessary).

The `Policy` reference implementation can be changed by resetting the value of the `policy.provider` Security Property (in the Security Properties file, *<java-home>*/`conf/security/java.security`) to the fully qualified name of the desired `Policy` implementation class.

The Security Property `policy.provider` specifies the name of the policy class, and the default is the following:

```
policy.provider=sun.security.provider.PolicyFile
```

To customize, you can change the property value to specify another class, as in

```
policy.provider=com.mycom.MyPolicy
```

Note that the `MyPolicy` class must be a subclass of `java.security.Policy`. It is perhaps worth emphasizing that such an override of the policy class is a temporary solution and a more comprehensive policy API will probably make this unnecessary.

java.security.GeneralSecurityException

This is an exception class that is a subclass of `java.lang.Exception`. The intention is that there should be two types of exceptions associated with security and the security packages.

- `java.lang.SecurityException` and its subclasses should be runtime exceptions (unchecked, not declared) that are likely to cause the execution of a program to stop.

  Such an exception is thrown only when some sort of security violation is detected. For example, such an exception is thrown when some code attempts to access a file, but it does not have permission for the access. Application developers may catch these exceptions, if they want.

- `java.security.GeneralSecurityException`, which is a subclass of `java.lang.Exception` (must be declared or caught) that is thrown in all other cases from within the security packages.

  Such an exception is security related but non-vital. For example, passing in an invalid key is probably not a security violation and should be caught and dealt with by a developer.

There are currently still two exceptions within the `java.security` package that are subclasses from `RuntimeException`. We at this moment cannot change these due to backward compatibility requirements. We will revisit this issue in the future.

Access Control Mechanisms and Algorithms

java.security.ProtectionDomain

The `ProtectionDomain` class encapsulates the characteristics of a domain. Such a domain encloses a set of classes whose instances are granted a set of permissions when being executed on behalf of a given set of `Principals`.

A `ProtectionDomain` is constructed with a `CodeSource`, a `ClassLoader`, an array of Principals, and a collection of `Permissions`. The `CodeSource` encapsulates the codebase (`java.net.URL`) for all classes in this domain, as well as a set of certificates (of type `java.security.cert.Certificate`) for public keys that correspond to the private keys that signed all code in this domain. The `Principals` represent the user on whose behalf the code is running.

The permissions passed in at `ProtectionDomain` construction time represent a static set of permissions bound to the domain regardless of the `Policy` in force. The `ProtectionDomain` subsequently consults the current policy during each security check to retrieve dynamic permissions granted to the domain.

Classes from different `CodeSources`, or that are being executed on behalf of different principals, belong to different domains.

Today all code shipped as part of the JDK is considered system code and run inside the unique system domain. Each applet or application runs in its appropriate domain, determined by policy.

It is possible to ensure that objects in any non-system domain cannot automatically discover objects in another non-system domain. This partition can be achieved by careful class resolution and loading, for example, using different classloaders for different domains. However, `SecureClassLoader` (or its subclasses) can, at its choice, load classes from different domains, thus allowing these classes to co-exist within the same name space (as partitioned by a classloader).

java.security.AccessController

The `AccessController` class is used for three purposes, each of which is described in further detail in sections below:

- to decide whether an access to a critical system resource is to be allowed or denied, based on the security policy currently in effect,
- to mark code as being "privileged", thus affecting subsequent access determinations, and
- to obtain a "snapshot" of the current calling context so access-control decisions from a different context can be made with respect to the saved context.

Any code that controls access to system resources should invoke `AccessController` methods if it wishes to use the specific security model and access control algorithm utilized by these methods. If, on the other hand, the application wishes to defer the security model to that of the `SecurityManager` installed at runtime, then it should instead invoke corresponding methods in the `SecurityManager` class.

For example, the typical way to invoke access control has been the following code (taken from an earlier version of JDK):

```
ClassLoader loader = this.getClass().getClassLoader();
if (loader != null) {
    SecurityManager security = System.getSecurityManager();
    if (security != null) {
        security.checkRead("path/file");
    }
}
```

Under the current architecture, the check typically should be invoked whether or not there is a classloader associated with a calling class. It could be simply, for example:

```
FilePermission perm = new FilePermission("path/file", "read");
AccessController.checkPermission(perm);
```

The `AccessController` `checkPermission` method examines the current execution context and makes the right decision as to whether or not the requested access is allowed. If it is, this check returns quietly. Otherwise, an `AccessControlException` (a subclass of `java.lang.SecurityException`) is thrown.

Note that there are (legacy) cases, for example, in some browsers, where whether there is a `SecurityManager` installed signifies one or the other security state that may result in different actions being taken. For backward compatibility, the `checkPermission` method on `SecurityManager` can be used.

```
SecurityManager security = System.getSecurityManager();
if (security != null) {
    FilePermission perm = new FilePermission("path/file", "read");
    security.checkPermission(perm);
}
```

We currently do not change this aspect of the `SecurityManager` usage, but would encourage application developers to use new techniques introduced in the JDK in their future programming when the built-in access control algorithm is appropriate.

The default behavior of the `SecurityManager` `checkPermission` method is actually to call the `AccessController` `checkPermission` method. A different `SecurityManager` implementation may implement its own security management approach, possibly including the addition of further constraints used in determining whether or not an access is permitted.

Algorithm for Checking Permissions

Suppose access control checking occurs in a thread of computation that has a chain of multiple callers (think of this as multiple method calls that cross the protection domain boundaries), as illustrated in the next figure.

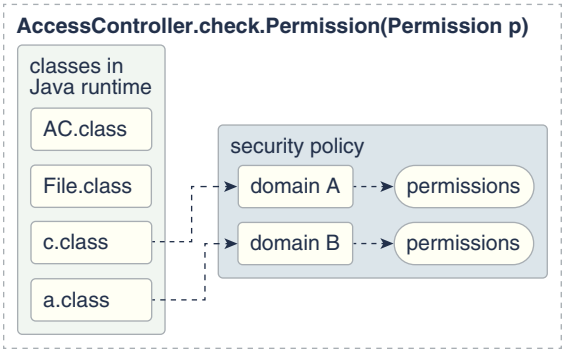

When the `checkPermission` method of the `AccessController` is invoked by the most recent caller (e.g., a method in the File class), the basic algorithm for deciding whether to allow or deny the requested access is as follows.

If any caller in the call chain does not have the requested permission, `AccessControlException` is thrown, unless the following is true – a caller whose domain is granted the said permission has been marked as "privileged" (see the next section) and all parties subsequently called by this caller (directly or indirectly) all have the said permission.

There are obviously two implementation strategies:

- In an "eager evaluation" implementation, whenever a thread enters a new protection domain or exits from one, the set of effective permissions is updated dynamically.

  The benefit is that checking whether a permission is allowed is simplified and can be faster in many cases. The disadvantage is that, because permission checking occurs much less frequently than cross-domain calls, a large percentage of permission updates are likely to be useless effort.

- In a "lazy evaluation" implementation, whenever permission checking is requested, the thread state (as reflected by the current state, including the current thread's call stack or its equivalent) is examined and a decision is reached to either deny or grant the particular access requested.

  One potential downside of this approach is performance penalty at permission checking time, although this penalty would have been incurred anyway in the "eager evaluation" approach (albeit at earlier times and spread out among each cross-domain call). Our implementation so far has yielded acceptable performance, so we feel that lazy evaluation is the most economical approach overall.

Therefore, the algorithm for checking permissions is currently implemented as "lazy evaluation". Suppose the current thread traversed m callers, in the order of caller 1 to caller 2 to caller `m`. Then caller m invoked the `checkPermission` method. The basic algorithm `checkPermission` uses to determine whether access is granted or denied is the following (see subsequent sections for refinements):

```
for (int i = m; i > 0; i--) {

    if (caller i's domain does not have the permission)
        throw AccessControlException

    else if (caller i is marked as privileged) {
        if (a context was specified in the call to doPrivileged)
```

```
                context.checkPermission(permission)
            if (limited permissions were specified in the call to
doPrivileged) {
                for (each limited permission) {
                    if (the limited permission implies the requested
permission)
                        return;
                }
            } else
                return;
        }
    }

    // Next, check the context inherited when the thread was created.
    // Whenever a new thread is created, the AccessControlContext at
    // that time is stored and associated with the new thread, as the
    // "inherited" context.

    inheritedContext.checkPermission(permission);
```

Handling Privileges

A static method in the `AccessController` class allows code in a class instance to inform the `AccessController` that a body of its code is "privileged" in that it is solely responsible for requesting access to its available resources, no matter what code caused it to do so.

That is, a caller can be marked as being "privileged" when it calls the `doPrivileged` method. When making access control decisions, the `checkPermission` method stops checking if it reaches a caller that was marked as "privileged" via a `doPrivileged` call without a context argument (see a subsequent section for information about a context argument). If that caller's domain has the specified permission, no further checking is done and `checkPermission` returns quietly, indicating that the requested access is allowed. If that domain does not have the specified permission, an exception is thrown, as usual.

The normal use of the "privileged" feature is as follows:

If you don't need to return a value from within the "privileged" block, do the following:

```
somemethod() {
     ...normal code here...
     AccessController.doPrivileged(new PrivilegedAction() {
         public Object run() {
             // privileged code goes here, for example:
             System.loadLibrary("awt");
             return null; // nothing to return
         }
     });
     ...normal code here...
}
```

`PrivilegedAction` is an interface with a single method, named `run`, that returns an Object. The above example shows creation of an anonymous inner class implementing that interface; a concrete implementation of the `run` method is supplied.

When the call to `doPrivileged` is made, an instance of the `PrivilegedAction` implementation is passed to it. The `doPrivileged` method calls the `run` method from the `PrivilegedAction` implementation after enabling privileges, and returns the `run` method's return value as the `doPrivileged` return value, which is ignored in this example. (For more information about inner classes, see Nested Classes in the Java Tutorials.

If you need to return a value, you can do something like the following:

```
somemethod() {
     ...normal code here...
     String user = (String) AccessController.doPrivileged(
       new PrivilegedAction() {
         public Object run() {
             return System.getProperty("user.name");
         }
       }
     );
     ...normal code here...
}
```

If the action performed in your `run` method could throw a "checked" exception (one listed in the `throws` clause of a method), then you need to use the `PrivilegedExceptionAction` interface instead of the `PrivilegedAction` interface:

```
somemethod() throws FileNotFoundException {
     ...normal code here...
   try {
     FileInputStream fis = (FileInputStream)
      AccessController.doPrivileged(
       new PrivilegedExceptionAction() {
         public Object run() throws FileNotFoundException {
             return new FileInputStream("someFile");
         }
       }
     );
   } catch (PrivilegedActionException e) {
     // e.getException() should be an instance of
     // FileNotFoundException,
     // as only "checked" exceptions will be "wrapped" in a
     // <code>PrivilegedActionException</code>.
     throw (FileNotFoundException) e.getException();
   }
     ...normal code here...
}
```

Some important points about being privileged: Firstly, this concept only exists within a single thread. As soon as the privileged code completes, the privilege is guaranteed to be erased or revoked.

Secondly, in this example, the body of code in the `run` method is privileged. However, if it calls less trustworthy code that is less privileged, that code will not gain any privileges as a result; a permission is only granted if the privileged code has the

permission *and* so do all the subsequent callers in the call chain up to the `checkPermission` call.

A variant of `AccessController.doPrivileged` enables code to assert a subset of its privileges without preventing the full traversal of the stack to check for other permissions. See Asserting a Subset of Privileges.

For more information about marking code as "privileged," see Appendix A: API for Privileged Blocks.

Inheritance of Access Control Context

When a thread creates a new thread, a new stack is created. If the current security context was not retained when this new thread was created, then when `AccessController.checkPermission` was called inside the new thread, a security decision would be made based solely upon the new thread's context, not taking into consideration that of the parent thread.

This clean stack issue would not be a security problem per se, but it would make the writing of secure code, and especially system code, more prone to subtle errors. For example, a non-expert developer might assume, quite reasonably, that a child thread (e.g., one that does not involve untrusted code) inherits the same security context from the parent thread (e.g., one that involves untrusted code). This would cause unintended security holes when accessing controlled resources from inside the new thread (and then passing the resources along to less trusted code), if the parent context was not in fact saved.

Thus, when a new thread is created, we actually ensure (via thread creation and other code) that it automatically inherits the parent thread's security context at the time of creation of the child thread, in such a way that subsequent `checkPermission` calls in the child thread will take into consideration the inherited parent context.

In other words, the logical thread context is expanded to include both the parent context (in the form of an `AccessControlContext`, described in the next section) and the current context, and the algorithm for checking permissions is expanded to the following. (Recall there are m callers up to the call to `checkPermission`, and see the next section for information about the `AccessControlContext` `checkPermission` method.)

```
    for (int i = m; i > 0; i--) {

        if (caller i's domain does not have the permission)
            throw AccessControlException

        else if (caller i is marked as privileged) {
            if (a context was specified in the call to doPrivileged)
                context.checkPermission(permission)
            if (limited permissions were specified in the call to
doPrivileged) {
                for (each limited permission) {
                    if (the limited permission implies the requested
permission)
                        return;
                }
            } else
                return;
```

```
    }
}

// Next, check the context inherited when the thread was created.
// Whenever a new thread is created, the AccessControlContext at
// that time is stored and associated with the new thread, as the
// "inherited" context.

inheritedContext.checkPermission(permission);
```

Note that this inheritance is transitive so that, for example, a grandchild inherits both from the parent and the grandparent. Also note that the inherited context snapshot is taken when the new child is created, and not when the child is first run. There is no public API change for the inheritance feature.

java.security.AccessControlContext

Recall that the `AccessController` `checkPermission` method performs security checks within the context of the current execution thread (including the inherited context). A difficulty arises when such a security check can only be done in a different context. That is, sometimes a security check that should be made within a given context will actually need to be done from within a different context. For example, when one thread posts an event to another thread, the second thread serving the requesting event would not have the proper context to complete access control, if the service requests access to controller resources.

To address this issue, we provide the `AccessController` `getContext` method and `AccessControlContext` class. The `getContext` method takes a "snapshot" of the current calling context, and places it in an `AccessControlContext` object, which it returns. A sample call is the following:

```
AccessControlContext acc = AccessController.getContext();
```

This context captures relevant information so that an access control decision can be made by checking, from within a different context, against this context information. For example, one thread can post a request event to a second thread, while also supplying this context information. `AccessControlContext` itself has a `checkPermission` method that makes access decisions based on the context it encapsulates, rather than that of the current execution thread. Thus, the second thread can perform an appropriate security check if necessary by invoking the following:

```
acc.checkPermission(permission);
```

The above method call is equivalent to performing the same security check in the context of the first thread, even though it is done in the second thread.

There are also times where one or more permissions must be checked against an access control context, but it is unclear a priori which permissions are to be checked. In these cases you can use the `doPrivileged` method that takes a context:

```
somemethod() {
    AccessController.doPrivileged(new PrivilegedAction() {
        public Object run() {
```

```
                // Code goes here. Any permission checks from
                // this point forward require both the current
                // context and the snapshot's context to have
                // the desired permission.
            }
        });
        ...normal code here...
```

Now the complete algorithm utilized by the `AccessController` `checkPermission` method can be given. Suppose the current thread traversed `m` callers, in the order of caller 1 to caller 2 to caller `m`. Then caller `m` invoked the `checkPermission` method. The algorithm `checkPermission` uses to determine whether access is granted or denied is the following

```
    for (int i = m; i > 0; i--) {

        if (caller i's domain does not have the permission)
            throw AccessControlException

        else if (caller i is marked as privileged) {
            if (a context was specified in the call to doPrivileged)
                context.checkPermission(permission)
            if (limited permissions were specified in the call to
doPrivileged) {
                for (each limited permission) {
                    if (the limited permission implies the requested
permission)
                        return;
                }
            } else
                return;
        }
    }

    // Next, check the context inherited when the thread was created.
    // Whenever a new thread is created, the AccessControlContext at
    // that time is stored and associated with the new thread, as the
    // "inherited" context.

    inheritedContext.checkPermission(permission);
```

Secure Class Loading

Dynamic class loading is an important feature of the Java Virtual Machine because it provides the Java platform with the ability to install software components at run-time. It has a number of unique characteristics. First of all, lazy loading means that classes are loaded on demand and at the last moment possible. Second, dynamic class loading maintains the type safety of the Java Virtual Machine by adding link-time checks, which replace certain run-time checks and are performed only once. Moreover, programmers can define their own class loaders that, for example, specify the remote location from which certain classes are loaded, or assign appropriate security attributes to them. Finally, class loaders can be used to provide separate name spaces for various software components. For example, a browser can load

applets from different web pages using separate class loaders, thus maintaining a degree of isolation between those applet classes. In fact, these applets can contain classes of the same name – these classes are treated as distinct types by the Java Virtual Machine.

The class loading mechanism is not only central to the dynamic nature of the Java programming language. It also plays a critical role in providing security because the class loader is responsible for locating and fetching the class file, consulting the security policy, and defining the class object with the appropriate permissions.

Class Loader Class Hierarchies

When loading a class, because there can be multiple instances of class loader objects in one Java Virtual Machine, an important question is how do we determine which class loader to use. The JDK has introduced multiple class loader classes are introduced that have distinct properties, so another important question is what type of class loader we should use.

The root of the class loader class hierarchy is an abstract class called `java.lang.ClassLoader`. Class `java.security.SecureClassLoader` is a subclass and a concrete implementation of the abstract `ClassLoader` class. Class `java.net.URLClassLoader` is a subclass of `SecureClassLoader`.

When creating a custom class loader class, one can subclass from any of the above class loader classes, depending on the particular needs of the custom class loader.

The Primordial Class Loader

Because each class is loaded by its class loader, and each class loader itself is a class and must be loaded by another class loader, we seem to have the obvious chicken-and-egg problem, i.e., where does the first class loader come from? There is a "primordial" class loader that bootstraps the class loading process. The primordial class loader is generally written in a native language, such as C, and does not manifest itself within the Java context. The primordial class loader often loads classes from the local file system in a platform-dependent manner.

Some classes, such as those defined in the `java.*` package, are essential for the correct functioning of the Java Virtual Machine and runtime system. They are often referred to as base classes. Due to historical reasons, all such classes have a class loader that is a null. This null class loader is perhaps the only sign of the existence of a primordial class loader. In fact, it is easier to simply view the null class loader as the primordial class loader.

Given all classes in one Java application environment, we can easily form a class loading tree to reflect the class loading relationship. Each class that is not a class loader is a leaf node. Each class's parent node is its class loader, with the null class loader being the root class. Such a structure is a tree because there cannot be cycles – a class loader cannot have loaded its own ancestor class loader.

Class Loader Delegation

When one class loader is asked to load a class, this class loader either loads the class itself or it can ask another class loader to do so. In other words, the first class loader can delegate to the second class loader. The delegation relationship is virtual in the sense that it has nothing to do with which class loader loads which other class loader. Instead, the delegation relationship is formed when class loader objects are created,

and in the form of a parent-child relationship. Nevertheless, the system class loader is the delegation root ancestor of all class loaders. Care must be taken to ensure that the delegation relationship does not contain cycles. Otherwise, the delegation process may enter into an infinite loop.

Class Resolution Algorithm

The default implementation of the JDK `ClassLoader` method for loading a class searches for classes in the following order:

- Check if the class has already been loaded.
- If the current class loader has a specified delegation parent, delegate to the parent to try to load this class. If there is no parent, delegate to the primordial class loader.
- Call a customizable method to find the class elsewhere.

Here, the first step looks into the class loader's local cache (or its functional equivalent, such as a global cache) to see if a loaded class matches the target class. The last step provides a way to customize the mechanism for looking for classes; thus a custom class loader can override this method to specify how a class should be looked up. For example, an applet class loader can override this method to go back to the applet host and try to locate the class file and load it over the network.

If at any step a class is located, it is returned. If the class is not found using the above steps, a `ClassNotFound` exception is thrown.

Observe that it is critical for type safety that the same class not be loaded more than once by the same class loader. If the class is not among those already loaded, the current class loader attempts to delegate the task to the parent class loader. This can occur recursively. This ensures that the appropriate class loader is used. For example, when locating a system class, the delegation process continues until the system class loader is reached.

We have seen the delegation algorithm earlier. But, given the name of any class, which class loader do we start with in trying to load the class? The rules for determining the class loader are the following:

- When loading the first class of an application, a new instance of the `URLClassLoader` is used.
- When loading the first class of an applet, a new instance of the `AppletClassLoader` is used.
- When `java.lang.Class.ForName` is directly called, the primordial class loader is used.
- If the request to load a class is triggered by a reference to it from an existing class, the class loader for the existing class is asked to load the class.

Note that rules about the use of `URLClassLoader` and `AppletClassLoader` instances have exceptions and can vary depending on the particular system environment. For example, a web browser may choose to reuse an existing `AppletClassLoader` to load applet classes from the same web page.

Due to the power of class loaders, we severely restrict who can create class loader instances. On the other hand, it is desirable to provide a convenient mechanism for applications or applets to specify URL locations and load classes from them. We

provide static methods to allow any program to create instances of the `URLClassLoader` class, although not other types of class loaders.

Security Management

Managing Applets and Applications

Currently, all JDK system code invokes `SecurityManager` methods to check the policy currently in effect and perform access control checks. There is typically a security manager (`SecurityManager` implementation) installed whenever an applet is running; the `appletviewer` and most browsers install a security manager.

A security manager is not automatically installed when an application is running. To apply the same security policy to an application found on the local file system as to downloaded applets, either the user running the application must invoke the Java Virtual Machine with the `-Djava.security.manager` command-line argument (which sets the value of the `java.security.manager` property), as in

```
java -Djava.security.manager SomeApp
```

or the application itself must call the `setSecurityManager` method in the `java.lang.System` class to install a security manager.

It is possible to specify on the command line a particular security manager to be utilized, by following `-Djava.security.manager` with an equals and the name of the class to be used as the security manager, as in

```
java -Djava.security.manager=COM.abc.MySecMgr SomeApp
```

If no security manager is specified, the built-in default security manager is utilized (unless the application installs a different security manager). All of the following are equivalent and result in usage of the default security manager:

```
java -Djava.security.manager SomeApp
java -Djava.security.manager="" SomeApp
java -Djava.security.manager=default SomeApp
```

The JDK includes a property named `java.class.path`. Classes that are stored on the local file system but should not be treated as base classes (e.g., classes built into the SDK) should be on this path. Classes on this path are loaded with a secure class loader and are thus subjected to the security policy being enforced.

There is also a `-Djava.security.policy` command-line argument whose usage determines what policy files are utilized. This command-line argument is described in detail in Default Policy Implementation and Policy File Syntax. Basically, if you don't include `-Djava.security.policy` on the command line, then the policy files specified in the security properties file will be used.

You can use a `-Djava.security.policy` command-line argument to specify an additional or a different policy file when invoking execution of an application. For example, if you type the following, where `pURL` is a URL specifying the location of a

policy file, then the specified policy file will be loaded in addition to all the policy files specified in the security properties file:

```
java -Djava.security.manager -Djava.security.policy=pURL SomeApp
```

If you instead type the following command, using a double equals, then just the specified policy file will be used; all others will be ignored:

```
java -Djava.security.manager -Djava.security.policy==pURL SomeApp
```

SecurityManager versus AccessController

The new access control mechanism is fully backward compatible. For example, all `check` methods in `SecurityManager` are still supported, although most of their implementations are changed to call the new `SecurityManager` `checkPermission` method, whose default implementation calls the `AccessController` `checkPermission` method. Note that certain internal security checks may stay in the `SecurityManager` class, unless it can be parameterized.

We have not at this time revised any system code to call `AccessController` instead of calling `SecurityManager` (and checking for the existence of a classloader), because of the potential of existing third-party applications that subclass the `SecurityManager` and customize the `check` methods. In fact, we added a new method `SecurityManager.checkPermission` that by default simply invokes `AccessController.checkPermission`.

To understand the relationship between `SecurityManager` and `AccessController`, it is sufficient to note that `SecurityManager` represents the concept of a central point of access control, while `AccessController` implements a particular access control algorithm, with special features such as the `doPrivileged` method. By keeping `SecurityManager` up to date, we maintain backward compatibility (e.g., for those applications that have written their own security manager classes based on earlier versions of the JDK) and flexibility (e.g., for someone wanting to customize the security model to implement mandatory access control or multilevel security). By providing `AccessController`, we build in the algorithm that we believe is the most restrictive and that relieves the typical programmer from the burden of having to write extensive security code in most scenarios.

We encourage the use of `AccessController` in application code, while customization of a security manager (via subclassing) should be the last resort and should be done with extreme care. Moreover, a customized security manager, such as one that always checks the time of the day before invoking standard security checks, could and should utilize the algorithm provided by `AccessController` whenever appropriate.

One thing to remember is that, when you implement your own `SecurityManager`, you should install it as trusted software and grant it `java.security.AllPermission`. You can do this by adjusting the policy file to grant `AllPermission` to your `SecurityManager`. For more information, see Default Policy Implementation and Policy File Syntax.

Auxiliary Tools

This section briefly describes the usage of two tools that assist in the deployment of security features.

The Key and Certificate Management Tool

`keytool` is a key and certificate management utility. It enables users to administer their own public/private key pairs and associated certificates for use in self-authentication (where the user authenticates himself/herself to other users/services) or data integrity and authentication services, using digital signatures. The authentication information includes both a sequence (chain) of X.509 certificates, and an associated private key, which can be referenced by a so-called "alias". This tool also manages certificates (that are "trusted" by the user), which are stored in the same database as the authentication information, and can be referenced by an "alias".

`keytool` stores the keys and certificates in a so-called *keystore*. The default keystore implementation implements the keystore as a file. It protects private keys with a password.

The chains of X.509 certificates are provided by organizations called Certification Authorities, or CAs. Identities (including CAs) use their private keys to authenticate their association with objects (such as with channels which are secured using SSL), with archives of code they signed, or (for CAs) with X.509 certificates they have issued. As a bootstrapping tool, certificates generated using the `-gencert` option may be used until a Certification Authority returns a certificate chain.

The private keys in this database are always stored in encrypted form, to make it difficult to disclose these private keys inappropriately. A password is required to access or modify the database. These private keys are encrypted using the "password", which should be several words long. If the password is lost, those authentication keys cannot be recovered.

In fact, each private key in the keystore can be protected using its own individual password, which may or may not be the same as the password that protects the keystore's overall integrity.

This tool is (currently) intended to be used from the command line, where one simply types `keytool` as a shell prompt. `keytool` is a script that executes the appropriate Java classes and is built together with the SDK.

The command line options for each command may be provided in any order. Typing an incorrect option or typing `keytool -help` will cause the tool's usage to be summarized on the output device (such as a shell window).

The JAR Signing and Verification Tool

The `jarsigner` tool can be used to digitally sign Java archives (JAR files), and to verify such signatures. This tool depends on the keystore that is managed by `keytool`.

> **Note:**
>
> You can also use the `jdk.security.jarsigner` API to sign JAR files.

GuardedObject and SignedObject

java.security.GuardedObject and java.security.Guard

Recall that the class `AccessControlContext` is useful when an access control decision has to be made in a different context. There is another such scenario, where the supplier of a resource is not in the same thread as the consumer of that resource, and the consumer thread cannot provide the supplier thread the access control context information (because the context is security-sensitive, or the context is too large to pass, or for other reasons). For this case, we provide a class called `GuardedObject` to protect access to the resource, illustrated in the figure below.

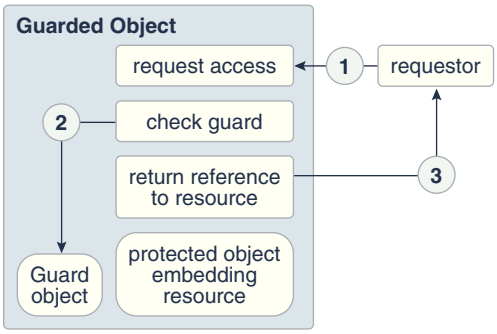

The basic idea is that the supplier of the resource can create an object representing the resource, create a `GuardedObject` that embeds the resource object inside, and then provide the `GuardedObject` to the consumer. In creating the `GuardedObject`, the supplier also specifies a `Guard` object such that anyone (including the consumer) can only obtain the resource object if certain (security) checks inside the `Guard` are satisfied.

`Guard` is an interface, so any object can choose to become a `Guard`. The only method in this interface is called `checkGuard`. It takes an `Object` argument and it performs certain (security) checks. The `Permission` class in `java.security` implements the `Guard` interface.

For example, suppose a system thread is asked to open a file `/a/b/c.txt` for read access, but the system thread does not know who the requestor is or under what circumstances the request is made. Therefore, the correct access control decision cannot be made at the server side. The system thread can use `GuardedObject` to delay the access control checking, as follows.

```
FileInputStream f = new FileInputStream("/a/b/c.txt");
FilePermission p = new FilePermission("/a/b/c.txt", "read");
GuardedObject g = new GuardedObject(f, p);
```

Now the system thread can pass g to the consumer thread. For that thread to obtain the file input stream, it has to call

```
FileInputStream fis = (FileInputStream) g.getObject();
```

This method in turn invokes the `checkGuard` method on the `Guard` object `p`, and because `p` is a `Permission`, its `checkGuard` method is in fact:

```
SecurityManager sm = System.getSecurityManager();
if (sm != null) sm.checkPermission(this);
```

This ensures that a proper access control check takes place within the consumer context. In fact, one can replace often-used hash tables and access control lists in many cases and simply store a hash table of `GuardedObjects`.

This basic pattern of `GuardedObject` and `Guard` is very general, and we expect that by extending the basic `Guard` and `GuardedObject` classes, developers can easily obtain quite powerful access control tools. For example, per-method invocation can be achieved with an appropriate `Guard` for each method, and a `Guard` can check the time of the day, the signer or other identification of the caller, or any other relevant information.

Note that certain typing information is lost because `GuardedObject` returns an `Object`. `GuardedObject` is intended to be used between cooperating parties so that the receiving party should know what type of object to expect (and to cast for). In fact, we envision that most usage of `GuardedObject` involves `subclassing` it (say to form a `GuardedFileInputStream` class), thus encapsulating typing information, and casting can happen suitably in the subclass.

java.security.SignedObject

This class is an essential building block for other security primitives. `SignedObject` contains another `Serializable` object, the (to-be-)signed object and its signature. If the signature is not null, it contains a valid digital signature of the signed object. This is illustrated in the figure below.

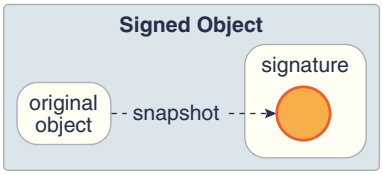

The underlying signing algorithm is set through a Signature object as a parameter to the `sign` method call, and the algorithm can be, among others, the NIST standard DSA, using DSA and SHA-256. The algorithm is specified using the same convention for signatures, such as "SHA/DSA".

The signed object is a "deep copy" (in serialized form) of an original object. Once the copy is made, further manipulation of the original object has no side effect on the copy. A signed object is immutable.

A typical example of creating a signed object is the following:

```
    Signature signingEngine = Signature.getInstance(algorithm,provider);
    SignedObject so = new SignedObject(myobject, signingKey,
signingEngine);
```

A typical example of verification is the following (having received `SignedObject so`), where the first line is not needed if the name of the algorithm is known:

```
    String algorithm = so.getAlgorithm();
    Signature verificationEngine = Signature.getInstance(algorithm,
provider);
    so.verify(verificationEngine);
```

Potential applications of `SignedObject` include:

- It can be used internally to any Java application environment as an unforgeable authorization token – one that can be passed around without the fear that the token can be maliciously modified without being detected.
- It can be used to sign and serialize data/object for storage outside the Java runtime (e.g., storing critical access control data on disk).
- Nested `SignedObjects` can be used to construct a logical sequence of signatures, resembling a chain of authorization and delegation.

It is intended that this class can be subclassed in the future to allow multiple signatures on the same signed object. In that case, existing method calls in this base class will be fully compatible in semantics. In particular, any `get` method will return the unique value if there is only one signature, and will return an arbitrary one from the set of signatures if there is more than one signature.

Discussion and Future Directions

Resource Consumption Management

Resource consumption management is relatively easy to implement in some cases (e.g., to limit the number of windows any application can pop up at any one time), while it can be quite hard to implement efficiently in other cases (e.g., to limit memory or file system usage). We plan to coherently address such issues in the future.

Arbitrary Grouping of Permissions

Sometimes it is convenient to group a number of permissions together and use a short-hand name to refer to them. For example, if we would like to have a permission called `SuperPermission` to include (and imply) both `FilePermission("-", "read,write")` and `SocketPermission("*", "connect,accept")`, technically we can use the class `Permissions` or a similar class to implement this super permission by using the `add` methods to add the required permissions. And such grouping can be arbitrarily complicated.

The more difficult issues are the following. First, to understand what actual permissions one is granting when giving out such a super permission, either a fixed and named permission class is created to denote a statically specified group of permissions, or the member permissions need to be spelled out in the policy file. Second, processing the policy (file) can become more complicated because the grouped permissions may need to be expanded. Moreover, nesting of grouped permission increases complexity even more.

Object-Level Protection

Given the object-oriented nature of the Java programming language, it is conceivable that developers will benefit from a set of appropriate object-level protection mechanisms that (1) goes beyond the natural protection provided by the Java programming language and that (2) supplements the thread-based access control mechanism.

One such mechanism is `SignedObject`. Another is the `SealedObject` class, which uses encryption to hide the content of an object.

`GuardedObject` is a general way to enforce access control at a per class/object per method level. This method, however, should be used only selectively, partly because this type of control can be difficult to administer at a high level.

Subdividing Protection Domains

A potentially useful concept not currently implemented is that of "subdomains." A subdomain is one that is enclosed in another. A subdomain would not have more permissions or privileges than the domain of which it is a subpart. A domain could be created, for example, to selectively further limit what a program can do.

Often a domain is thought of as supporting inheritance: a subdomain would automatically inherit the parent domain's security attributes, except in certain cases where the parent further restricts the subdomain explicitly. Relaxing a subdomain by right amplification is a possibility with the notion of trusted code.

For convenience, we can think of the system domain as a single, big collection of all system code. For better protection, though, system code should be run in multiple system domains, where each domain protects a particular type of resource and is given a special set of rights. For example, if file system code and network system code run in separate domains, where the former has no rights to the networking resources and the latter has no rights to the file system resources, the risks and consequence of an error or security flaw in one system domain is more likely to be confined within its boundary.

Running Applets with Signed Content

The JAR and Manifest specifications on code signing allow a very flexible format. Classes within the same archive can be signed with different keys, and a class can be unsigned, signed with one key, or signed with multiple keys. Other resources within the archive, such as audio clips and graphic images, can also be signed or unsigned, just like classes can.

This flexibility brings about the issue of interpretation. The following questions need to be answered, especially when keys are treated differently:

1. Should images and audios be required to be signed with the same key if any class in the archive is signed?
2. If images and audios are signed with different keys, can they be placed in the same `appletviewer` (or browser page), or should they be sent to different viewers for processing?

These questions are not easy to answer, and require consistency across platforms and products to be the most effective. Our intermediate approach is to provide a

simple answer – all images and audio clips are forwarded to be processed within the same applet classloader, whether they are signed or not. This temporary solution will be improved once a consensus is reached.

Moreover, if a digital signature cannot be verified because the bytecode content of a class file does not match the signed hash value in the JAR, a security exception is thrown, as the original intention of the JAR author is clearly altered. Previously, there was a suggestion to run such code as untrusted. This idea is undesirable because the applet classloader allows the loading of code signed by multiple parties. This means that accepting a partially modified JAR file would allow an untrusted piece of code to run together with and access other code through the same classloader.

Appendix A: API for Privileged Blocks

This section explains what privileged code is and what it is used for. It also shows you how to use the doPrivileged API.

- Using the doPrivileged API
- What It Means to Have Privileged Code
- Reflection

Using the doPrivileged API

This section describes the doPrivileged API and the use of the privileged feature.

- No Return Value, No Exception Thrown
- Accessing Local Variables
- Handling Exceptions
- Asserting a Subset of Privileges
- Least Privilege
- More Privilege

No Return Value, No Exception Thrown

If you do not need to return a value from within the *privileged* block, your call to `doPrivileged` can look like Example 1-1.

Note that the invocation of `doPrivileged` with a lambda expression explicitly casts the lambda expression as of type `PrivilegedAction<Void>`. Another version of the method `doPrivileged` exists that takes an object of type `PrivilegedExceptionAction`; see Handling Exceptions.

`PrivilegedAction` is a functional interface with a single abstract method, named `run`, that returns a value of type specified by its type parameter.

Note that this example ignores the return value of the `run` method. Also, depending on what *privileged code* actually consists of, you might have to make some changes due to the way inner classes work. For example, if *privileged code* throws an exception or attempts to access local variables, then you will have to make some changes, which is described later.

Be *very* careful in your use of the *privileged* construct, and always remember to make the privileged code section as small as possible. That is, try to limit the code within the

`run` method to only what needs to be run with privileges, and do more general things outside the `run` method. Also note that the call to `doPrivileged` should be made in the code that wants to enable its privileges. Do not be tempted to write a utility class that itself calls `doPrivileged` as that could lead to security holes. You can write utility classes for `PrivilegedAction` classes though, as shown in the preceding example. See Guideline 9-3: Safely invoke `java.security.AccessController.doPrivileged` in Secure Coding Guidelines for the Java Programming Language.

Example 1-1 Sample Code for Privileged Block

- In a class that implements the interface `PrivilegedAction`.
- In an anonymous class.
- In a lambda expression.

```
import java.security.*;

public class NoReturnNoException {

    class MyAction implements PrivilegedAction<Void> {
        public Void run() {
            // Privileged code goes here, for example:
            System.loadLibrary("awt");
            return null; // nothing to return
        }
    }

    public void somemethod() {

        MyAction mya = new MyAction();

        // Become privileged:
        AccessController.doPrivileged(mya);

        // Anonymous class
        AccessController.doPrivileged(new PrivilegedAction<Void>() {
            public Void run() {
                // Privileged code goes here, for example:
                System.loadLibrary("awt");
                return null; // nothing to return
            }
        });

        // Lambda expression
        AccessController.doPrivileged((PrivilegedAction<Void>)
            () -> {
                // Privileged code goes here, for example:
                System.loadLibrary("awt");
                return null; // nothing to return
            }
        );
    }

    public static void main(String... args) {
        NoReturnNoException myApplication = new NoReturnNoException();
```

```
        myApplication.somemethod();
    }
}
```

Returning Values

If you need to return a value, then you can do something like the following:

```
System.out.println(
    AccessController.doPrivileged((PrivilegedAction<String>)
        () -> System.getProperty("user.name")
    )
);
```

Accessing Local Variables

If you are using a lambda expression or anonymous inner class, then any local variables you access must be `final` or effectively final.

For example:

```
String lib = "awt";
AccessController.doPrivileged((PrivilegedAction<Void>)
    () -> {
        System.loadLibrary(lib);
        return null; // nothing to return
    }
);

AccessController.doPrivileged(new PrivilegedAction<Void>() {
    public Object run() {
        System.loadLibrary(lib);
        return null; // nothing to return
    }
});
```

The variable `lib` is effectively final because its value has not been modified. For example, suppose you add the following assignment statement after the declaration of the variable `lib`:

```
lib = "swing";
```

The compiler generates the following errors when it encounters the invocation `System.loadLibrary` both in the lambda expression and the anonymous class:

- `error: local variables referenced from a lambda expression must be final or effectively final`
- `error: local variables referenced from an inner class must be final or effectively final`

See Accessing Members of an Enclosing Class in Local Classes for more information.

If there are cases where you cannot make an existing variable effectively final (because it gets set multiple times), then you can create a new `final` variable right before invoking the `doPrivileged` method, and set that variable equal to the other variable. For example:

```
String lib;

// The lib variable gets set multiple times so you can't make it
// effectively final.

// Create a final String that you can use inside of the run method
final String fLib = lib;

AccessController.doPrivileged((PrivilegedAction<Void>)
    () -> {
        System.loadLibrary(fLib);
        return null; // nothing to return
    }
);
```

Handling Exceptions

If the action performed in your `run` method could throw a *checked* exception (one that must be listed in the `throws` clause of a method), then you need to use the `PrivilegedExceptionAction` interface instead of the `PrivilegedAction` interface.

Example 1-2 Sample for Handling Exceptions

If a checked exception is thrown during execution of the `run` method, then it is placed in a `PrivilegedActionException` *wrapper* exception that is then thrown and should be caught by your code, as illustrated in the following example:

```
public void processSomefile() throws IOException {

    try {
        Path path = FileSystems.getDefault().getPath("somefile");
        BufferedReader br = AccessController.doPrivileged(
            (PrivilegedExceptionAction<BufferedReader>)
                () -> Files.newBufferedReader(path)
        );
        // ... read from file and do something
    } catch (PrivilegedActionException e) {

        // e.getException() should be an instance of IOException
        // as only checked exceptions will be wrapped in a
        // PrivilegedActionException.
        throw (IOException) e.getException();
    }
}
```

Asserting a Subset of Privileges

As of JDK 8, a variant of doPrivileged is available that enables code to assert a subset of its privileges, without preventing the full traversal of the stack to check for other permissions. This variant of the doPrivileged variant has three parameters, one of which you use to specify this subset of privileges. For example, the following excerpt asserts a privilege to retrieve system properties:

```
// Returns the value of the specified property. All code
// is allowed to read the app.version and app.vendor
// properties.

public String getProperty(final String prop) {
    return AccessController.doPrivileged(
        (PrivilegedAction<String>) () -> System.getProperty(prop),
        null,
        new java.util.PropertyPermission("app.version", "read"),
        new java.util.PropertyPermission("app.vendor", "read")
    );
    }
```

The first parameter of this version of doPrivileged is of type java.security.PrivilegedAction. In this example, the first parameter is a lambda expression that implements the functional interface PrivilegedAction whose run method returns the value of the system property specified by the parameter prop.

The second parameter of this version of doPrivileged is of type AccessControlContext. Sometimes you need to perform an additional security check within a different context, such as a worker thread. You can obtain an AccessControlContext instance from a particular calling context with the method AccessControlContext.getContext. If you specify null for this parameter (as in this example), then the invocation of doPrivileged does not perform any additional security checks.

The third parameter of this version of doPrivileged is of type Permission..., which is a varargs parameter. This means that you can specify one or more Permission parameters or an array of Permission objects, as in Permission[]. In this example, the invocation of doPrivileged can retrieve the properties app.version and app.vendor.

You can use this three parameter variant of doPrivileged in a mode of least privilege or a mode of more privilege.

Least Privilege

The typical use case of the doPrivileged method is to enable the method that invokes it to perform one or more actions that require permission checks without requiring the callers of the current method to have all the necessary permissions.

For example, the current method might need to open a file or make a network request for its own internal implementation purposes.

Before JDK 8, calls to doPrivileged methods had only two parameters. They worked by granting temporary privileges to the calling method and stopping the normal full traversal of the stack for access checking when it reached that class, rather than continuing up the call stack where it might reach a method whose defining class does not have the required permission. Typically, the class that is calling doPrivileged

might have additional permissions that are not required in that code path and which might also be missing from some caller classes.

Normally, these extra permissions are not exercised at runtime. Not elevating them through use of `doPrivileged` helps to block exploitation of any incorrect code that could perform unintended actions. This is especially true when the `PrivilegedAction` is more complex than usual, or when it calls code outside the class or package boundary that might evolve independently over time.

The three-parameter variant of `doPrivileged` is generally safer to use because it avoids unnecessarily elevating permissions that are not intended to be required. However, it executes less efficiently so simple or performance-critical code paths might choose not to use it.

More Privilege

When coding the current method, you want to temporarily extend the permission of the calling method to perform an action.

For example, a framework I/O API might have a general purpose method for opening files of a particular data format. This API would take a normal file path parameter and use it to open an underlying `FileInputStream` using the calling code's permissions. However, this might also allow any caller to open the data files in a special directory that contains some standard demonstration samples.

The callers of this API could be directly granted a `FilePermission` for *read* access. However, it might not be convenient or possible for the security policy of the calling code to be updated. For example, the calling code could be a sandboxed applet.

One way to implement this is for the code to check the incoming path and determine if it refers to a file in the special directory. If it does, then it would call `doPrivileged`, enabling all permissions, then open the file inside the `PrivilegedAction`. If the file was not in the special directory, the code would open the file without using `doPrivileged`.

This technique requires the implementation to carefully handle the requested file path to determine if it refers to the special shared directory. The file path must be canonicalized before calling `doPrivileged` so that any relative path will be processed (and permission to read the `user.dir` system property will be checked) prior to determining if the path refers to a file in the special directory. It must also prevent malicious "../" path elements meant to escape out of the special directory.

A simpler and better implementation would use the variant of `doPrivileged` with the third parameter. It would pass a `FilePermission` with *read* access to the special directory as the third parameter. Then any manipulation of the file would be inside the `PrivilegedAction`. This implementation is simpler and much less prone to contain a security flaw.

What It Means to Have Privileged Code

Marking code as *privileged* enables a piece of trusted code to temporarily enable access to more resources than are available directly to the code that called it.

The policy for a JDK installation specifies what permissions which types of system resource accesses – are allowed for code from specified code sources. A *code source* (of type `CodeSource`) essentially consists of the code location (URL) and a reference to the certificates containing the public keys corresponding to the private keys used to sign the code (if it was signed).

The policy is represented by a `Policy` object. More specifically, it is represented by a `Policy` subclass providing an implementation of the abstract methods in the `Policy` class (which is in the `java.security` package).

The source location for the policy information used by the `Policy` object depends on the `Policy` implementation. The `Policy` reference implementation obtains its information from policy configuration files. See Default Policy Implementation and Policy File Syntax for information about the `Policy` reference implementation and the syntax that must be used in policy files it reads.

A *protection domain* encompasses a `CodeSource` instance and the permissions granted to code from that `CodeSource`, as determined by the security policy currently in effect. Thus, classes signed by the same keys and from the same URL are typically placed in the same domain, and a class belongs to one and only one protection domain. (However, classes signed by the same keys and from the same URL but loaded by separate class loader instances are typically placed in separate domains.) Classes that have the same permissions but are from different code sources belong to different domains.

Classes shipped with the JDK run-time image and loaded by the bootstrap class loader are granted `AllPermission`. However, classes shipped with the JDK run-time image and loaded by the platform class loader are granted permissions as specified by the default policy of the JDK. Each module's classes are assigned a unique protection domain using the `jrt` URL scheme and may only be granted the permissions necessary for them to function correctly, and not necessarily `AllPermission`.

Each applet or application runs in its appropriate domain, determined by its code source. For an applet (or an application running under a security manager) to be allowed to perform a secured action (such as reading or writing a file), the applet or application must be granted permission for that particular action.

More specifically, whenever a resource access is attempted, *all* code traversed by the execution thread up to that point must have permission for that resource access, *unless some code on the thread has been marked as* ***privileged***. That is, suppose that access control checking occurs in a thread of execution that has a chain of multiple callers. (Think of this as multiple method calls that potentially cross the protection domain boundaries.) When the `AccessController.checkPermission` method is invoked by the most recent caller, the basic algorithm for deciding whether to allow or deny the requested access is as follows: If the code for any caller in the call chain does not have the requested permission, then an `AccessControlException` is thrown, *unless* the following is true: a caller whose code is granted the said permission has been marked as *privileged*, and all parties subsequently called by this caller (directly or indirectly) have the said permission.

> **Note:**
>
> The method `AccessController.checkPermission` is normally invoked indirectly through invocations of specific `SecurityManager` methods that begin with the word `check` such as `checkConnect` or through the method `SecurityManager.checkPermission`. Normally, these checks only occur if a `SecurityManager` has been installed; code checked by the `AccessController.checkPermission` method first checks if the method `System.getSecurityManager` returns null.

Marking code as *privileged* enables a piece of trusted code to temporarily enable access to more resources than are available directly to the code that called it. This is necessary in some situations. For example, an application might not be allowed direct access to files that contain fonts, but the system utility to display a document must obtain those fonts, on behalf of the user. The system utility must become privileged in order to obtain the fonts.

Reflection

The `doPrivileged` method can be invoked reflectively using the `java.lang.reflect.Method.invoke` method.

One subtlety that must be considered is the interaction of this API with reflection. The `doPrivileged` method can be invoked reflectively using the `java.lang.reflect.Method.invoke` method. In this case, the privileges granted in privileged mode are not those of `Method.invoke` but of the non-reflective code that invoked it. Otherwise, system privileges could erroneously (or maliciously) be conferred on user code. Note that similar requirements exist when using reflection in the existing API.

Appendix B: Acknowledgments

The design and implementation of new security features in Java 2 SDK is the work of primarily members of the JavaSoft security group. Other (past and present) members of the JavaSoft community provided invaluable insight, detailed reviews, and much needed technical assistance. Significant contributors, in alphabetical order, include but are not limited to: Gigi Ankeny, Josh Bloch, Satya Dodda, Charlie Lai, Sheng Liang, Jan Luehe, Marianne Mueller, Jeff Nisewanger, Hemma Prafullchandra, Roger Riggs, Nakul Saraiya, Bill Shannon, Roland Schemers, and Vijay Srinivasan.

This work is not possible without strong support from JavaSoft management (our thanks go to Dick Neiss, Jon Kannegaard, and Alan Baratz), and the testing and documentation groups (especially Mary Dageforde). We are grateful for technical guidance from James Gosling, Graham Hamilton, and Jim Mitchell.

We received numerous suggestions from our corporate partners and licensees, whom we could not fully list here.

Appendix C: References

M. Gasser. Building a Secure Computer System. Van Nostrand Reinhold Co., New York, 1988.

L. Gong, "Java Security: Present and Near Future". IEEE Micro, 17(3):14--19, May/June 1997.

L. Gong, T.M.A. Lomas, R.M. Needham, and J.H. Saltzer, "Protecting Poorly Chosen Secrets from Guessing Attacks". IEEE Journal on Selected Areas in Communications, 11(5):648--656, June, 1993.

J. Gosling, Bill Joy, and Guy Steele. The Java Language Specification. Addison-Wesley, Menlo Park, California, August 1996.

A.K. Jones. Protection in Programmed Systems. Ph.D. dissertation, Carnegie-Mellon University, Pittsburgh, PA 15213, June 1973.

B.W. Lampson. Protection. In Proceedings of the 5th Princeton Symposium on Information Sciences and Systems, Princeton University, March 1971. Reprinted in ACM Operating Systems Review, 8(1):18--24, January, 1974.

T. Lindholm and F. Yellin. The Java Virtual Machine Specification. Addison-Wesley, Menlo Park, California, 1997.

P.G. Neumann. Computer-Related Risks. Addison-Wesley, Menlo Park, California, 1995.

U.S. General Accounting Office. Information Security: Computer Attacks at Department of Defense Pose Increasing Risks. Technical Report GAO/AIMD-96-84, Washington, D.C. 20548, May 1996.

J.H. Saltzer. Protection and the Control of Information Sharing in Multics. Communications of the ACM, 17(7):388--402, July 1974.

J.H. Saltzer and M.D. Schroeder. The Protection of Information in Computer Systems}. Proceedings of the IEEE, 63(9):1278--1308, September 1975.

M.D. Schroeder. Cooperation of Mutually Suspicious Subsystems in a Computer Utility. Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA 02139, September 1972.

W.A. Wulf, R. Levin, and S.P. Harbison. HYDRA/C.mmp -- An Experimental Computer System. McGraw-Hill, 1981.

Standard Algorithm Names

See Java Security Standard Algorithm Names Specification for information about the set of standard names for algorithms, certificate and keystore types that Java SE requires and uses.

Permissions in the JDK

A permission represents access to a system resource. In order for a resource access to be allowed for an applet (or an application running with a security manager), the corresponding permission must be explicitly granted to the code attempting the access.

A permission typically has a name (often referred to as a "target name") and, in some cases, a comma-separated list of one or more actions. For example, the following code creates a `FilePermission` object representing read access to the file named `abc` in the `/tmp` directory:

```
perm = new java.io.FilePermission("/tmp/abc", "read");
```

Here, the target name is "`/tmp/abc`" and the action string is "`read`".

> **Important:**
>
> The above statement creates a permission object. A permission object represents, but does not grant access to, a system resource. Permission objects are constructed and assigned ("granted") to code based on the policy in effect. When a permission object is assigned to some code, that code is granted the permission to access the system resource specified in the permission object, in the specified manner. A permission object may also be constructed by the current security manager when making access decisions. In this case, the (target) permission object is created based on the requested access, and checked against the permission objects granted to and held by the code making the request.

The policy for a Java application environment is represented by a Policy object. In the `"JavaPolicy"` Policy implementation, the policy can be specified within one or more policy configuration files. The policy file(s) specify what permissions are allowed for code from specified code sources. The following is a sample policy file entry that grants code from the `/home/sysadmin` directory read access to the file `/tmp/abc`:

```
grant codeBase "file:/home/sysadmin/" {
    permission java.io.FilePermission "/tmp/abc", "read";
};
```

To know more about policy file locations and granting permissions in policy files, see Default Policy Implementation and Policy File Syntax.

Technically, whenever a resource access is attempted, *all* code traversed by the execution thread up to that point must have permission for that resource access, unless some code on the thread has been marked as "privileged." See Appendix A: API for Privileged Blocks.

Permission Descriptions and Risks

The following is a list of all built-in JDK permission types. The class summary for each permission type discusses the risks of granting each permission.

- `java.awt.AWTPermission`
- `java.io.FilePermission`
- `java.io.SerializablePermission`
- `java.lang.RuntimePermission`
- `java.lang.management.ManagementPermission`
- `java.lang.reflect.ReflectPermission`
- `java.net.NetPermission`
- `java.net.URLPermission`
- `java.net.SocketPermission`
- `java.nio.file.LinkPermission`
- `java.security.AllPermission`

- `java.security.SecurityPermission`
- `java.security.UnresolvedPermission`
- `java.sql.SQLPermission`
- `java.util.logging.LoggingPermission`
- `java.util.PropertyPermission`
- `javax.management.MBeanPermission`
- `javax.management.MBeanServerPermission`
- `javax.management.MBeanTrustPermission`
- `javax.management.remote.SubjectDelegationPermission`
- `javax.net.ssl.SSLPermission`
- `javax.security.auth.AuthPermission`
- `javax.security.auth.PrivateCredentialPermission`
- `javax.security.auth.kerberos.DelegationPermission`
- `javax.security.auth.kerberos.ServicePermission`
- `javax.sound.sampled.AudioPermission`

> **Note:**
>
> See Appendix A: FilePermission Path Name Canonicalization Disabled By Default for important information about a change in how `FilePermission` path names are canonicalized.

Methods and the Permissions They Require

The following table is a list of methods that require permissions, which `SecurityManager` method they call, and which permission is checked by the default implementation of that `SecurityManager` method.

> **Note:**
>
> This list is not complete; other methods exist that require permissions. See the Java SE and JDK API Specification for additional information on methods that throw `SecurityException` and the permissions that are required.

In the default `SecurityManager` method implementations, a call to a method in the **Method** column can only be successful if the permission specified in the corresponding entry in the **SecurityManager Method** column is allowed by the policy currently in effect. For example, consider the following table row:

Method	SecurityManager Method Called	Permission

`java.awt.Toolkit`
`public final EventQueue`
`getSystemEventQueue()`
	`checkPermission`	`java.awt.AWTPermission "accessEventQueue";`

This table row specifies that a call to the `getSystemEventQueue` method in the `java.awt.Toolkit` class results in a call to the `checkPermission` `SecurityManager` method, which can only be successful if the following permission is granted to code on the call stack:

```
java.awt.AWTPermission "accessEventQueue";
```

The table rows have the following format, where the runtime value of `foo` replaces the string `{foo}` in the permission name.

Method	SecurityManager Method Called	Permission

`some.package.class`
`public static void`
`someMethod(String foo);`
	`checkXXX`	`SomePermission "{foo}";`

As an example, here is one table entry:

Method	SecurityManager Method Called	Permission

`java.io.FileInputStream`
 	
`FileInputStream(String name)`
	`checkRead(String)`	`java.io.FilePermission "{name}", "read";`

If the `FileInputStream` method (in this case, a constructor) is called with `"/test/MyTestFile"` as the `name` argument, as in

```
FileInputStream("/test/MyTestFile");
```

then in order for the call to succeed, the following permission must be set in the current policy, allowing read access to the file `"/test/MyTestFile"`:

```
java.io.FilePermission "/test/MyTestFile", "read";
```

More specifically, the permission must either be explicitly set, as above, or implied by another permission, such as the following:

```
java.io.FilePermission "/test/*", "read";
```

which allows read access to any files in the `"/test"` directory.

In some cases, a term in braces is not exactly the same as the name of a specific method argument but is meant to represent the relevant value. Here is an example:

Method	SecurityManager Method Called	Permission
`java.net.DatagramSocket public synchronized void receive(DatagramPacket p);`	`checkAccept({host}, {port})`	`java.net.SocketPermission "{host}:{port}", "accept";`

Here, the appropriate host and port values are calculated by the `receive` method and passed to `checkAccept`.

In most cases, just the name of the SecurityManager method called is listed. Where the method is one of multiple methods of the same name, the argument types are also listed, for example for `checkRead(String)` and `checkRead(FileDescriptor)`. In other cases where arguments may be relevant, they are also listed.

The following table is ordered by package name; the methods in classes in the `java.awt` package are listed first, followed by methods in classes in the `java.beans` package, and so on:

Table 1-6 Methods and the Permissions

Method	SecurityManager Method	Permission
`java.awt.Graphics2d public abstract void setComposite(Composite comp)`	`checkPermission`	`java.awt.AWTPermission "readDisplayPixels"` if this `Graphics2D` context is drawing to a `Component` on the display screen and the `Composite` is a custom object rather than an instance of the `AlphaComposite` class. Note: The `setComposite` method is actually abstract and thus can't invoke security checks. Each actual implementation of the method should call the `java.lang.SecurityManager checkPermission` method with a `java.awt.AWTPermission("readDisplayPixels")` permission under the conditions noted.

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.awt.Robot`
`public Robot()`
`public Robot(GraphicsDevice screen)`
	`checkPermission`	`java.awt.AWTPermission "createRobot"`

`java.awt.Toolkit`
`public void addAWTEventListener( AWTEventListener listener, long eventMask)`
`public void removeAWTEventListener( AWTEventListener listener)`
	`checkPermission`	`java.awt.AWTPermission "listenToAllAWTEvents"`

`java.awt.Toolkit` `public abstract PrintJob getPrintJob( Frame frame, String jobtitle, Properties props)`	`checkPrintJobAccess`	`java.lang.RuntimePermission "queuePrintJob"` Note: The `getPrintJob` method is actually abstract and thus can't invoke security checks. Each actual implementation of the method should call the `java.lang.SecurityManager checkPrintJobAccess` method, which is successful only if the `java.lang.RuntimePermission "queuePrintJob"` permission is currently allowed.
`java.awt.Toolkit` `public abstract Clipboard getSystemClipboard()`	`checkPermission`	`java.awt.AWTPermission "accessClipboard"` Note: The `getSystemClipboard` method is actually abstract and thus can't invoke security checks. Each actual implementation of the method should call the `checkPermission` method, which is successful only if the `java.awt.AWTPermission "accessClipboard"` permission is currently allowed.

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
`java.awt.Toolkit` `public final EventQueue getSystemEventQueue()`	`checkPermission`	`java.awt.AWTPermission "accessEventQueue"`
`java.awt.Window Window()`	`checkPermission`	If `java.awt.AWTPermission "showWindowWithoutWarningBanner"` is set, the window will be displayed without a banner warning that the window was created by an applet. It it's not set, such a banner will be displayed.

`java.beans.Beans`
`public static void setDesignTime( boolean isDesignTime)`
`public static void setGuiAvailable( boolean isGuiAvailable)`
 	
`java.beans.Introspector`
`public static synchronized void setBeanInfoSearchPath(String path[])`
 	
`java.beans.PropertyEditorManager`
`public static void registerEditor( Class targetType, Class editorClass)`
`public static synchronized void setEditorSearchPath(String path[])`
	`checkPropertiesAccess`	`java.util.PropertyPermission "*", "read,write"`

`java.io.File`
`public boolean delete()`
`public void deleteOnExit()`
	`checkDelete(String)`	`java.io.FilePermission "{name}", "delete"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.io.FileInputStream`
 	
`FileInputStream(FileDescriptor fdObj)`
	`checkRead(FileDescriptor)`	`java.lang.RuntimePermission "readFileDescriptor"`

`java.io.FileInputStream`
`  FileInputStream(String name)`
`  FileInputStream(File file)`
 	
`java.io.File`
`  public boolean exists()`
`  public boolean canRead()`
`  public boolean isFile()`
`  public boolean isDirectory()`
`  public boolean isHidden()`
`  public long lastModified()`
`  public long length()`
`  public String[] list()`
`  public String[] list(`
`  FilenameFilter filter)`
`  public File[] listFiles()`
`  public File[] listFiles(`
`  FilenameFilter filter)`
`  public File[] listFiles(`
`  FileFilter filter)`
 	
`java.io.RandomAccessFile`
`  RandomAccessFile(String name, String mode)`
`  RandomAccessFile(File file, String mode)`
 	
(where mode is "r" in both `RandomAccessFile` constructurs)
	`checkRead(String)`	`java.io.FilePermission "{name}", "read"`

`java.io.FileOutputStream`
 	
`FileOutputStream(FileDescriptor fdObj)`
	`checkWrite(FileDescriptor)`	`java.lang.RuntimePermission "writeFileDescriptor"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
java.io.FileOutputStream FileOutputStream(File file) FileOutputStream(String name) FileOutputStream(String name, boolean append) java.io.File public boolean canWrite() public boolean createNewFile() public static File createTempFile(String prefix, String suffix) public static File createTempFile(String prefix, String suffix, File directory) public boolean mkdir() public boolean mkdirs() public boolean renameTo(File dest) public boolean setLastModified(long time) public boolean setReadOnly()	checkWrite(String)	java.io.FilePermission "{name}", "write"
java.io.ObjectInputStream protected final boolean enableResolveObject(boolean enable); java.io.ObjectOutputStream protected final boolean enableReplaceObject(boolean enable)	checkPermission	java.io.SerializablePermission "enableSubstitution"

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.io.ObjectInputStream protected ObjectInputStream()`
 	
`java.io.ObjectOutputStream protected ObjectOutputStream()`
	`checkPermission`	`java.io.SerializablePermission "enableSubclassImplementation"`

`java.io.RandomAccessFile RandomAccessFile(String name, String mode)` (where mode is "rw")	`checkRead(String)` and `checkWrite(String)`	`java.io.FilePermission "{name}", "read,write"`
`java.lang.Class public static Class forName( String name, boolean initialize, ClassLoader loader)`	`checkPermission`	If `loader` is null, and the caller's class loader is not null, then `java.lang.RuntimePermission("getClassLoader")`
`java.lang.Class public ClassLoader getClassLoader()`	`checkPermission`	If the caller's class loader is null, or is the same as or an ancestor of the class loader for the class whose class loader is being requested, no permission is needed. Otherwise, `java.lang.RuntimePermission "getClassLoader"` is required.

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.lang.Class`
`public Class[] getDeclaredClasses()`
`public Field[] getDeclaredFields()`
`public Method[] getDeclaredMethods()`
`public Constructor[] getDeclaredConstructors()`
`public Field getDeclaredField( String name)`
`public Method getDeclaredMethod(...)`
`public Constructor getDeclaredConstructor(...)`
	`checkMemberAccess(this, Member.DECLARED)` and, if this class is in a package, `checkPackageAccess({pkgName})`	Default `checkMemberAccess` does not require any permissions if "this" class's class loader is the same as that of the caller. Otherwise, it requires `java.lang.RuntimePermission "accessDeclaredMembers"`. If this class is in a package, `java.lang.RuntimePermission "accessClassInPackage.{pkgName}"` is also required.

`java.lang.Class`
`public Class[] getClasses()`
`public Field[] getFields()`
`public Method[] getMethods()`
`public Constructor[] getConstructors()`
`public Field getField(String name)`
`public Method getMethod(...)`
`public Constructor getConstructor(...)`
	`checkMemberAccess(this, Member.PUBLIC)` and, if class is in a package, `checkPackageAccess({pkgName})`	Default `checkMemberAccess` does not require any permissions when the access type is `Member.PUBLIC`. If this class is in a package, `java.lang.RuntimePermission "accessClassInPackage.{pkgName}"` is required.

`java.lang.Class` `public ProtectionDomain getProtectionDomain()`	`checkPermission`	`java.lang.RuntimePermission "getProtectionDomain"`

`java.lang.ClassLoader`
`ClassLoader()`
`ClassLoader(ClassLoader parent)`
	`checkCreateClassLoader`	`java.lang.RuntimePermission "createClassLoader"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.lang.ClassLoader`
`public static ClassLoader getSystemClassLoader()`
`public ClassLoader getParent()`
	`checkPermission`	If the caller's class loader is null, or is the same as or an ancestor of the class loader for the class whose class loader is being requested, no permission is needed. Otherwise, `java.lang.RuntimePermission "getClassLoader"` is required.

`java.lang.Runtime`
`public Process exec(String command)`
`public Process exec( String command, String envp[])`
`public Process exec(String cmdarray[])`
`public Process exec( String cmdarray[], String envp[])`
	`checkExec`	`java.io.FilePermission "{command}", "execute"`

`java.lang.Runtime`
`public void exit(int status)`
`public static void runFinalizersOnExit(boolean value)`
 	
`java.lang.System`
`public static void exit(int status)`
`public static void runFinalizersOnExit(boolean value)`
	`checkExit(status)` where status is 0 for `runFinalizersOnExit`	`java.lang.RuntimePermission "exitVM.{status}"`

`java.lang.Runtime`
`public void addShutdownHook(Thread hook)`
`public boolean removeShutdownHook(Thread hook)`
	`checkPermission`	`java.lang.RuntimePermission "shutdownHooks"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.lang.Runtime`
`public void load(String lib)`
`public void loadLibrary(String lib)`
 	
`java.lang.System`
`public static void load(String filename)`
`public static void loadLibrary(String libname)`
	`checkLink({libName})` where `{libName}` is the lib, filename or libname argument	`java.lang.RuntimePermission "loadLibrary.{libName}"`

`java.lang.SecurityManager` methods	`checkPermission`	See Table 1-7.

`java.lang.System`
`public static Properties getProperties()`
`public static void setProperties(Properties props)`
	`checkPropertiesAccess`	`java.util.PropertyPermission "*", "read,write"`

`java.lang.System`
`public static String getProperty(String key)`
`public static String getProperty(String key, String def)`
	`checkPropertyAccess`	`java.util.PropertyPermission "{key}", "read"`

`java.lang.System`
`public static void setIn(InputStream in)`
`public static void setOut(PrintStream out)`
`public static void setErr(PrintStream err)`
	`checkPermission`	`java.lang.RuntimePermission "setIO"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.lang.System`
`public static String`
`setProperty(String key, String value)`
	`checkPermission`	`java.util.PropertyPermission "{key}", "write"`

`java.lang.System`
`public static synchronized void`
 	
`setSecurityManager(SecurityManager s)`
	`checkPermission`	`java.lang.RuntimePermission "setSecurityManager"`

`java.lang.Thread` `public ClassLoader getContextClassLoader()`	`checkPermission`	If the caller's class loader is null, or is the same as or an ancestor of the context class loader for the thread whose context class loader is being requested, no permission is needed. Otherwise, `java.lang.RuntimePermission "getClassLoader"` is required.

`java.lang.Thread`
`public void setContextClassLoader(`
`ClassLoader cl)`
	`checkPermission`	`java.lang.RuntimePermission "setContextClassLoader"`

`java.lang.Thread`
`public final void checkAccess()`
`public void interrupt()`
`public final void suspend()`
`public final void resume()`
`public final void setPriority(`
`int newPriority)`
`public final void setName(`
`String name)`
`public final void setDaemon(`
`boolean on)`
	`checkAccess(this)`	`java.lang.RuntimePermission "modifyThread"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.lang.Thread`
`public static int`
`enumerate(Thread tarray[])`
	`checkAccess({threadGroup})`	`java.lang.RuntimePermission "modifyThreadGroup"`

`java.lang.Thread` `public final void stop()`	`checkAccess(this)`. Also `checkPermission` if the current thread is trying to stop a thread other than itself.	`java.lang.RuntimePermission "modifyThread".` . Also `java.lang.RuntimePermission "stopThread"` if the current thread is trying to stop a thread other than itself.

`java.lang.Thread`
`public final synchronized void`
`stop(Throwable obj)`
	`checkAccess(this)`. Also `checkPermission` if the current thread is trying to stop a thread other than itself or `obj` is not an instance of `ThreadDeath`.	`java.lang.RuntimePermission "modifyThread".` Also `java.lang.RuntimePermission "stopThread"` if the current thread is trying to stop a thread other than itself or `obj` is not an instance of `ThreadDeath`.

`java.lang.Thread`
`Thread()`
`Thread(Runnable target)`
`Thread(String name)`
`Thread(Runnable target, String name)`
 	
`java.lang.ThreadGroup`
`ThreadGroup(String name)`
`ThreadGroup(`
`ThreadGroup parent,`
`String name)`
	`checkAccess({parentThreadGroup})`	`java.lang.RuntimePermission "modifyThreadGroup"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.lang.Thread`
`Thread(ThreadGroup group, ...)`
 	
`java.lang.ThreadGroup`
`public final void checkAccess()`
`public int enumerate(Thread list[])`
`public int enumerate(Thread list[], boolean recurse)`
`public int enumerate(ThreadGroup list[])`
`public int enumerate(ThreadGroup list[], boolean recurse)`
`public final ThreadGroup getParent()`
`public final void setDaemon(boolean daemon)`
`public final void setMaxPriority(int pri)`
`public final void suspend()`
`public final void resume()`
`public final void destroy()`
	`checkAccess(this)` for `ThreadGroup` methods, or `checkAccess(group)` for `Thread` methods	`java.lang.RuntimePermission "modifyThreadGroup"`

`java.lang.ThreadGroup` `public final void interrupt()`	`checkAccess(this)`	Requires `java.lang.RuntimePermission "modifyThreadGroup"`. Also requires `java.lang.RuntimePermission "modifyThread"`, since the `java.lang.Thread interrupt()` method is called for each thread in the thread group and in all of its subgroups. See the `Thread interrupt()` method.

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
`java.lang.ThreadGroup` `public final void stop()`	`checkAccess(this)`	Requires `java.lang.RuntimePermission "modifyThreadGroup"`. Also requires `java.lang.RuntimePermission "modifyThread"` and possibly `java.lang.RuntimePermission "stopThread"`, since the `java.lang.Thread stop()` method is called for each thread in the thread group and in all of its subgroups. See the `Thread stop()` method.

`java.lang.reflect.AccessibleObject`
`public static void setAccessible(...)`
`public void setAccessible(...)`
	`checkPermission`	`java.lang.reflect.ReflectPermission "suppressAccessChecks"`

`java.net.Authenticator`
`public static PasswordAuthentication requestPasswordAuthentication(`
`InetAddress addr,`
`int port,`
`String protocol,`
`String prompt,`
`String scheme)`
	`checkPermission`	`java.net.NetPermission "requestPasswordAuthentication"`

`java.net.Authenticator` `public static void setDefault(Authenticator a)`	`checkPermission`	`java.net.NetPermission "setDefaultAuthenticator"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.net.MulticastSocket`
`public void joinGroup(InetAddress mcastaddr)`
`public void leaveGroup(InetAddress mcastaddr)`
	`checkMulticast(InetAddress)`	`java.net.SocketPermission( mcastaddr.getHostAddress(), "accept,connect")`

`java.net.DatagramSocket`
`public void send(DatagramPacket p)`
	`checkMulticast(p.getAddress())` or `checkConnect( p.getAddress().getHostAddress(), p.getPort())`	`if (p.getAddress().isMulticastAddress()) {`
		`java.net.SocketPermission(`
		
		`(p.getAddress()).getHostAddress(),`
		`"accept,connect")`
		`} else {`
		`port = p.getPort();`
		`host = p.getAddress().getHostAddress();`
		`if (port == -1)`
		
		`java.net.SocketPermission "{host}","resolve";`
		`else`
		
		`java.net.SocketPermission "{host}:{port}","connect";`
		`}`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.net.MulticastSocket`
`public synchronized void send(DatagramPacket p, byte ttl)`
	`checkMulticast(p.getAddress(), ttl)` or `checkConnect( p.getAddress().getHostAddress(), p.getPort())`	`if (p.getAddress().isMulticastAddress()) {`
		`java.net.SocketPermission(`
		`(p.getAddress()).getHostAddress(), "accept,connect")`
		`} else {`
		`port = p.getPort();`
		`host = p.getAddress().getHostAddress();`
		`if (port == -1)`
		`java.net.SocketPermission "{host}","resolve";`
		`else`
		`java.net.SocketPermission "{host}:{port}","connect"`
		`}`

`java.net.InetAddress`
`public String getHostName()`
`public static InetAddress[] getAllByName(String host)`
`public static InetAddress getLocalHost()`
 	
`java.net.DatagramSocket`
`public InetAddress getLocalAddress()`
	`checkConnect({host}, -1)`	`java.net.SocketPermission "{host}", "resolve"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.net.ServerSocket`
`ServerSocket(...)`
 	
`java.net.DatagramSocket`
`DatagramSocket(...)`
 	
`java.net.MulticastSocket`
`MulticastSocket(...)`
	`checkListen({port})`	`java.net.SocketPermission`
		`"localhost:{port}","listen";`

`java.net.ServerSocket`
`public Socket accept()`
`protected final void`
`implAccept(Socket s)`
	`checkAccept({host},`
	`{port})`
		`java.net.SocketPermission`
		`"{host}:{port}", "accept"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
java.net.ServerSocket public static synchronized void setSocketFactory(...) java.net.Socket public static synchronized void setSocketImplFactory(...) java.net.URL public static synchronized void setURLStreamHandlerFactory(...) java.net.URLConnection public static synchronized void setContentHandlerFactory(...) public static void setFileNameMap(FileNameMap map) java.net.HttpURLConnection public static void setFollowRedirects(boolean set) java.rmi.activation.ActivationGroup public static synchronized ActivationGroup createGroup(...) public static synchronized void setSystem(ActivationSystem system) java.rmi.server.RMISocketFactory	checkSetFactory	java.lang.RuntimePermission "setFactory"

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
public synchronized static void setSocketFactory(...)		
java.net.Socket Socket(...)	checkConnect({host}, {port})	java.net.SocketPermission "{host}:{port}", "connect"
java.net.DatagramSocket public synchronized void receive(DatagramPacket p)	checkAccept({host}, {port})	java.net.SocketPermission "{host}:{port}", "accept"
java.net.URL URL(...)	checkPermission	java.net.NetPermission "specifyStreamHandler"
java.net.URLClassLoader URLClassLoader(...)	checkCreateClassLoader	java.lang.RuntimePermission "createClassLoader"
java.security.AccessControlContext public AccessControlContext(AccessControlContext acc, DomainCombiner combiner) public DomainCombiner getDomainCombiner()	checkPermission	java.security.SecurityPermission "createAccessControlContext"
java.security.Identity public void addCertificate(...)	checkSecurityAccess("addIdentityCertificate")	java.security.SecurityPermission "addIdentityCertificate"
java.security.Identity public void removeCertificate(...)	checkSecurityAccess("removeIdentityCertificate")	java.security.SecurityPermission "removeIdentityCertificate"

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
`java.security.Identity public void setInfo(String info)`	`checkSecurityAccess("setIdentityInfo")`	`java.security.SecurityPermission "setIdentityInfo"`
`java.security.Identity public void setPublicKey(PublicKey key)`	`checkSecurityAccess("setIdentityPublicKey")`	`java.security.SecurityPermission "setIdentityPublicKey"`
`java.security.Identity public String toString(...)`	`checkSecurityAccess("printIdentity")`	`java.security.SecurityPermission "printIdentity"`
`java.security.IdentityScope protected static void setSystemScope()`	`checkSecurityAccess("setSystemScope")`	`java.security.SecurityPermission "setSystemScope"`
`java.security.Permission public void checkGuard(Object object)`	`checkPermission(this)`	This `Permission` object is the permission checked.
`java.security.Policy public static Policy getPolicy()`	`checkPermission`	`java.security.SecurityPermission "getPolicy"`
`java.security.Policy public static void setPolicy(Policy policy)`	`checkPermission`	`java.security.SecurityPermission "setPolicy"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission

`java.security.Policy`
`public static Policy`
`getInstance(`
`String type,`
`SpiParameter params)`
`getInstance(`
`String type,`
`SpiParameter params,`
`String provider)`
`getInstance(`
`String type,`
`SpiParameter params,`
`Provider provider)`
	`checkPermission`	`java.security.SecurityPermission "createPolicy.{type}"`

`java.security.Provider` `public synchronized void clear()`	`checkSecurityAccess("clearProviderProperties."+{name})`	`java.security.SecurityPermission "clearProviderProperties.{name}"` where `name` is the provider name.
`java.security.Provider` `public synchronized Object put(Object key, Object value)`	`checkSecurityAccess("putProviderProperty."+{name})`	`java.security.SecurityPermission "putProviderProperty.{name}"` where `name` is the provider name.
`java.security.Provider` `public synchronized Object remove(Object key)`	`checkSecurityAccess("removeProviderProperty."+{name})`	`java.security.SecurityPermission "removeProviderProperty.{name}"` where `name` is the provider name.
`java.security.SecureClassLoader` `SecureClassLoader(...)`	`checkCreateClassLoader`	`java.lang.RuntimePermission "createClassLoader"`
`java.security.Security` `public static void getProperty(String key)`	`checkPermission`	`java.security.SecurityPermission "getProperty.{key}"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
`java.security.Security public static int addProvider(Provider provider) public static int insertProviderAt( Provider provider, int position);`	`checkSecurityAccess("insertProvider."+provider.getName())`	`java.security.SecurityPermission "insertProvider.{name}"`
`java.security.Security public static void removeProvider(String name)`	`checkSecurityAccess("removeProvider."+name)`	`java.security.SecurityPermission "removeProvider.{name}"`
`java.security.Security public static void setProperty(String key, String datum)`	`checkSecurityAccess("setProperty."+key)`	`java.security.SecurityPermission "setProperty.{key}"`
`java.security.Signer public PrivateKey getPrivateKey()`	`checkSecurityAccess("getSignerPrivateKey")`	`java.security.SecurityPermission "getSignerPrivateKey"`
`java.security.Signer public final void setKeyPair(KeyPair pair)`	`checkSecurityAccess("setSignerKeypair")`	`java.security.SecurityPermission "setSignerKeypair"`
`java.sql.DriverManager public static synchronized void setLogWriter(PrintWriter out)`	`checkPermission`	`java.sql.SQLPermission "setLog"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
`java.sql.DriverManager public static synchronized void setLogStream(PrintWriter out)`	`checkPermission`	`java.sql.SQLPermission "setLog"`
`java.util.Locale public static synchronized void setDefault(Locale newLocale)`	`checkPermission`	`java.util.PropertyPermission "user.language","write"`
`java.util.zip.ZipFile ZipFile(String name)`	`checkRead`	`java.io.FilePermission "{name}","read"`
`javax.security.auth.Subject public static Subject getSubject( final AccessControlContext acc)`	`checkPermission`	`javax.security.auth.AuthPermission "getSubject"`
`javax.security.auth.Subject public void setReadOnly()`	`checkPermission`	`javax.security.auth.AuthPermission "setReadOnly"`
`javax.security.auth.Subject public static Object doAs( final Subject subject, final PrivilegedAction action)`	`checkPermission`	`javax.security.auth.AuthPermission "doAs"`

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
javax.security.auth.Subject public static Object doAs(final Subject subject, final PrivilegedExceptionAction action) throws java.security.PrivilegedActionException	checkPermission	javax.security.auth.AuthPermission "doAs"
javax.security.auth.Subject public static Object doAsPrivileged(final Subject subject, final PrivilegedAction action, final AccessControlContext acc)	checkPermission	javax.security.auth.AuthPermission "doAsPrivileged"
javax.security.auth.Subject public static Object doAsPrivileged(final Subject subject, final PrivilegedExceptionAction action, final AccessControlContext acc) throws java.security.PrivilegedActionException	checkPermission	javax.security.auth.AuthPermission "doAsPrivileged"
javax.security.auth.SubjectDomainCombiner public Subject getSubject()	checkPermission	javax.security.auth.AuthPermission "getSubjectFromDomainCombiner"

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
javax.security.auth.SubjectDomainCombiner public Subject getSubject()	checkPermission	javax.security.auth.AuthPermission "getSubjectFromDomainCombiner"
javax.security.auth.login.LoginContext public LoginContext(String name) throws LoginException	checkPermission	javax.security.auth.AuthPermission "createLoginContext.{name}"
javax.security.auth.login.LoginContext public LoginContext(String name, Subject subject) throws LoginException	checkPermission	javax.security.auth.AuthPermission "createLoginContext.{name}"
javax.security.auth.login.LoginContext public LoginContext(String name, CallbackHandler callbackHandler) throws LoginException	checkPermission	javax.security.auth.AuthPermission "createLoginContext.{name}"
javax.security.auth.login.LoginContext public LoginContext(String name, Subject subject, CallbackHandler callbackHandler) throws LoginException	checkPermission	javax.security.auth.AuthPermission "createLoginContext.{name}"

Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
`javax.security.auth.login.Configuration public static Configuration getConfiguration()`	`checkPermission`	`javax.security.auth.AuthPermission "getLoginConfiguration"`
`javax.security.auth.login.Configuration public static void setConfiguration( Configuration configuration)`	`checkPermission`	`javax.security.auth.AuthPermission "setLoginConfiguration"`
`javax.security.auth.login.Configuration public static void refresh()`	`checkPermission`	`javax.security.auth.AuthPermission "refreshLoginConfiguration"`
`javax.security.auth.login.Configuration public static Configuration getInstance( String type, SpiParameter params) getInstance( String type, SpiParameter params, String provider) getInstance(String type, SpiParameter params, Provider provider)`	`checkPermission`	`javax.security.auth.AuthPermission "createLoginConfiguration.{type}"`

java.lang.SecurityManager Method Permission Checks

The following table shows which permissions are checked by the default implementations of the `java.lang.SecurityManager` methods.

Each of the specified `check` methods calls the `SecurityManager` `checkPermission` method with the specified permission, except for the `checkConnect` and `checkRead` methods that take a context argument. Those methods expect the context to be an

`AccessControlContext` and they call the context's `checkPermission` method with the specified permission.

Table 1-7 java.lang.SecurityManager Methods and Permissions

Method	Permission
`public void checkAccept(String host, int port);`	`java.net.SocketPermission "{host}:{port}", "accept";`
`public void checkAccess(Thread t);`	`java.lang.RuntimePermission "modifyThread";`
`public void checkAccess(ThreadGroup g);`	`java.lang.RuntimePermission "modifyThreadGroup";`
`public void checkAwtEventQueueAccess();` **Note:** This method is deprecated; use instead `public void checkPermission(Permission perm);`	`java.awt.AWTPermission "accessEventQueue";`

`public void checkConnect(String host, int port);`	`if (port == -1)`
	`  java.net.SocketPermission "{host}","resolve";`
	`else`
	`  java.net.SocketPermission "{host}:{port}","connect";`

`public void checkConnect(`
`  String host,`
`  int port,`
`  Object context);`
	`if (port == -1)`
	`  java.net.SocketPermission "{host}","resolve";`
	`else`
	`  java.net.SocketPermission "{host}:{port}","connect";`

Table 1-7 (Cont.) java.lang.SecurityManager Methods and Permissions

Method	Permission
`public void checkCreateClassLoader();`	`java.lang.RuntimePermission "createClassLoader";`
`public void checkDelete(String file);`	`java.io.FilePermission "{file}", "delete";`

`public void checkExec(String cmd);`	if cmd is an absolute path:
	`java.io.FilePermission "{cmd}", "execute";`
	else
	`java.io.FilePermission "<<ALL_FILES>>", "execute";`

`public void checkExit(int status);`	`java.lang.RuntimePermission "exitVM.{status}";`
`public void checkLink(String lib);`	`java.lang.RuntimePermission "loadLibrary.{lib}";`
`public void checkListen(int port);`	`java.net.SocketPermission "localhost:{port}","listen";`

`public void checkMemberAccess(Class clazz, int which);`
Note:
This method is deprecated; use instead `public void checkPermission(Permission perm);`
	`if (which != Member.PUBLIC) {`
	`  if (currentClassLoader() != clazz.getClassLoader()) {`
	`  checkPermission(`
	`  new java.lang.RuntimePermission(`
	`  "accessDeclaredMembers"));`
	`  }`
	`}`

Table 1-7 (Cont.) java.lang.SecurityManager Methods and Permissions

Method	Permission
`public void checkMulticast(InetAddress maddr);`	`java.net.SocketPermission( maddr.getHostAddress(),"accept,connect");`
`public void checkMulticast(InetAddress maddr, byte ttl);` **Note:** This method is deprecated; use instead `public void checkPermission(Permission perm);`	`java.net.SocketPermission( maddr.getHostAddress(),"accept,connect");`
`public void checkPackageAccess(String pkg);`	`java.lang.RuntimePermission "accessClassInPackage.{pkg}";`
`public void checkPackageDefinition(String pkg);`	`java.lang.RuntimePermission "defineClassInPackage.{pkg}";`
`public void checkPrintJobAccess();`	`java.lang.RuntimePermission "queuePrintJob";`
`public void checkPropertiesAccess();`	`java.util.PropertyPermission "*", "read,write";`
`public void checkPropertyAccess(String key);`	`java.util.PropertyPermission "{key}", "read,write";`
`public void checkRead(FileDescriptor fd);`	`java.lang.RuntimePermission "readFileDescriptor";`

Table 1-7 (Cont.) java.lang.SecurityManager Methods and Permissions

Method	Permission
`public void checkRead(String file);`	`java.io.FilePermission "{file}", "read";`
`public void checkRead(String file, Object context);`	`java.io.FilePermission "{file}", "read";`
`public void checkSecurityAccess(String target);`	`java.security.SecurityPermission "{target}";`
`public void checkSetFactory();`	`java.lang.RuntimePermission "setFactory";`
`public void checkSystemClipboardAccess();` **Note:** This method is deprecated; use instead `public void checkPermission(Permission perm);`	`java.awt.AWTPermission "accessClipboard";`
`public boolean checkTopLevelWindow(Object window);` **Note:** This method is deprecated; use instead `public void checkPermission(Permission perm);`	`java.awt.AWTPermission "showWindowWithoutWarningBanner";`
`public void checkWrite(FileDescriptor fd);`	`java.lang.RuntimePermission "writeFileDescriptor";`

Table 1-7 (Cont.) java.lang.SecurityManager Methods and Permissions

Method	Permission
`public void checkWrite(String file);`	`java.io.FilePermission "{file}", "write";`
`public SecurityManager();`	`java.lang.RuntimePermission "createSecurityManager";`

JDK Supported Permissions

The following permissions are not standard but the JDK supports them; you may need to grant them in policy files.

- `jdk.net.NetworkPermission "setOption.SO_FLOW_SLA";`
- `com.sun.tools.attach.AttachPermission "attachVirtualMachine";`
- `com.sun.jdi.JDIPermission "virtualMachineManager";`
- `com.sun.security.jgss.InquireSecContextPermission "*";`
- `jdk.jfr.FlightRecorderPermission "accessFlightRecorder", "registerEvent";`
- `javax.smartcardio.CardPermission "*", "connect,reset,exclusive,transmitControl,getBasicChannel,openLogicalChannel";`

Default Policy Implementation and Policy File Syntax

The policy for a Java programming language application environment (specifying which permissions are available for code from various sources, and executing as various principals) is represented by a Policy object. More specifically, it is represented by a `Policy` subclass providing an implementation of the abstract methods in the `Policy` class (which is in the `java.security` package).

The source location for the policy information utilized by the Policy object is up to the Policy implementation. The Policy reference implementation obtains its information from static policy configuration files.

The rest of this document pertains to the Policy reference implementation and the syntax that must be used in policy files it reads:

- Default Policy Implementation
- Default Policy File Locations
- Modifying the Policy Implementation
- Policy File Syntax
- Policy File Examples

- Property Expansion in Policy Files
- Windows Systems, File Paths, and Property Expansion
- General Expansion in Policy Files

Default Policy Implementation

In the Policy reference implementation, the policy can be specified within one or more policy configuration files. The configuration file(s) specify what permissions are allowed for code from a specified code source, and executed by a specified principal. Each configuration file must be encoded in UTF-8.

There is by default a single system-wide policy file, and a single (optional) user policy file. By default, permissions required by JDK modules that are loaded by the platform class loader or its ancestors are always granted.

The Policy reference implementation is initialized the first time its `getPermissions` method is called, or whenever its `refresh` method is called. Initialization involves parsing the policy configuration file(s) (see Policy File Syntax), and then populating the `Policy` object.

Default Policy File Locations

There is by default a single system-wide policy file, and a single (optional) user policy file. When the Policy is initialized, the system policy is loaded in first, and then the user policy is added to it. If neither policy is present, a built-in policy is used. This built-in policy is the same as the `java.policy` file installed with the JRE.

System Policy File Locations

By default, the system policy file is `<`*`java-home`*`>/conf/security/java.policy`.

The system policy file is meant to grant system-wide code permissions. The `java.policy` file installed with the JDK allows anyone to listen on dynamic ports, and allows any code to read certain "standard" properties that are not security-sensitive, such as the `os.name` and `file.separator` properties.

User Policy File Location

By default, the user policy file is `<`*`user-home`*`>/.java.policy`.

Policy File Location and Format

Policy file locations are specified in the security properties file `<`*`java-home`*`>/conf/security/java.security`.

The policy file locations are specified as the values of properties whose names are of the following form:

```
policy.url.n
```

Here, *n* is a number. You specify each such property value in a line of the following form:

```
policy.url.n=URL
```

Here, *URL* is a URL specification. For example, the default system and user policy files are defined in the security properties file as:

```
policy.url.1=file:${java.home}/conf/security/java.policy
policy.url.2=file:${user.home}/.java.policy
```

(See Property Expansion in Policy Files for information about specifying property values via a special syntax, such as specifying the `java.home` property value via `${java.home}`.)

You can actually specify a number of URLs (including ones of the form "`http://`"), and all the designated policy files will get loaded. You can also comment out or change the second one to disable reading the default user policy file.

The algorithm starts at `policy.url.1`, and keeps incrementing until it does not find a URL. Thus if you have `policy.url.1` and `policy.url.3`, and `policy.url.3` will never be read.

Specifying an Additional Policy File at Runtime

It is also possible to specify an additional or a different policy file when invoking execution of an application. This can be done via the `-Djava.security.policy` command line argument, which sets the value of the `java.security.policy` property. For example, if you use following command, where *someURL* is a URL specifying the location of a policy file, then the specified policy file will be loaded in addition to all the policy files that are specified in the security properties file.

```
java -Djava.security.manager -Djava.security.policy=someURL SomeApp
```

The URL can be any regular URL or simply the name of a policy file in the current directory, as in:

```
java -Djava.security.manager -Djava.security.policy=mypolicy SomeApp
```

The `-Djava.security.manager` option ensures that the default security manager is installed, and thus the application is subject to policy checks. It is not required if the application `SomeApp` installs a security manager.

If you use the following command (note the double equals) then *just* the specified policy file will be used; all the ones indicated in the security properties file will be ignored.

```
java -Djava.security.manager -Djava.security.policy==someURL SomeApp
```

> **Note:**
>
> The policy file value of the `-Djava.security.policy` option is ignored if the `policy.allowSystemProperty` property in the security properties file is set to false. The default is true.

Modifying the Policy Implementation

The Policy reference implementation can be modified by editing the security properties file, which is the `java.security` file in the `conf/security` directory of the JDK.

An alternative policy class can be given to replace the Policy reference implementation class, as long as the former is a subclass of the abstract Policy class and implements the `getPermissions` method (and other methods as necessary).

One of the types of properties you can set in `java.security` is of the following form:

```
policy.provider=PolicyClassName
```

PolicyClassName must specify the fully qualified name of the desired `Policy` implementation class.

The default security properties file entry for this property is the following:

```
policy.provider=sun.security.provider.PolicyFile
```

To customize, you can change the property value to specify another class, as in

```
policy.provider=com.mycom.MyPolicy
```

Policy File Syntax

The policy configuration file(s) for a JDK installation specifies what permissions (which types of system resource accesses) are granted to code from a specified code source, and executed as a specified principal.

For an applet (or an application running under a security manager) to be allowed to perform secured actions (such as reading or writing a file), the applet (or application) must be granted permission for that particular action. In the Policy reference implementation, that permission must be granted by a grant entry in a policy configuration file. See below and the Java Security Architecture Specification for more information. (The only exception is that code always automatically has permission to read files from its same (URL) location, and subdirectories of that location; it does not need explicit permission to do so.)

A policy configuration file essentially contains a list of entries. It may contain a "keystore" entry, and contains zero or more "grant" entries.

Keystore Entry

A *keystore* is a database of private keys and their associated digital certificates such as X.509 certificate chains authenticating the corresponding public keys. The keytool utility is used to create and administer keystores. The keystore specified in a policy configuration file is used to look up the public keys of the signers specified in the grant

entries of the file. A keystore entry must appear in a policy configuration file if any grant entries specify signer aliases, or if any grant entries specify principal aliases.

At this time, there can be only one `keystore`/`keystorePasswordURL` entry in the policy file (other entries following the first one are ignored). This entry can appear anywhere outside the file's grant entries. It has the following syntax:

```
keystore "some_keystore_url", "keystore_type", "keystore_provider";
keystorePasswordURL "some_password_url";
```

Here,

some_keystore_url
Specifies the URL location of the keystore.

some_password_url
Specifies the URL location of the keystore password.

keystore_type
Specifies the keystore type.

keystore_provider
Specifies the keystore provider.

> **Note:**
>
> - The input stream from `some_keystore_url` is passed to the `KeyStore.load` method.
> - If NONE is specified as the URL, then a null stream is passed to the `KeyStore.load` method. NONE should be specified in the URL if the KeyStore is not file-based. For example, if it resides on a hardware token device.
> - The URL is relative to the policy file location. If the policy file is specified in the security properties file as:
>
>  ```
>  policy.url.1=http://foo.example.com/fum/some.policy
>  ```
>
> and that policy file has an entry:
>
>  ```
>  keystore ".keystore";
>  ```
>
> then the keystore will be loaded from:
>
>  ```
>  http://foo.example.com/fum/.keystore
>  ```
> - The URL can also be absolute.

A **keystore type** defines the storage and data format of the keystore information, and the algorithms used to protect private keys in the keystore and the integrity of the

keystore itself. The default type is "PKCS12". Thus, if the keystore type is "PKCS12", it does not need to be specified in the keystore entry.

Grant Entries

Code being executed is always considered to come from a particular "code source" (represented by an object of type `CodeSource`). The code source includes not only the location (URL) where the code originated from, but also a reference to the certificate(s) containing the public key(s) corresponding to the private key(s) used to sign the code. Certificates in a code source are referenced by symbolic alias names from the user's keystore. Code is also considered to be executed as a particular principal (represented by an object of type `Principal`), or group of principals.

Each **grant entry** includes one or more "permission entries" preceded by optional `codeBase`, `signedBy`, and principal name/value pairs that specify which code you want to grant the permissions. The basic format of a grant entry is the following:

```
grant signedBy "signer_names", codeBase "URL",
      principal principal_class_name "principal_name",
      principal principal_class_name "principal_name",
      ... {

    permission permission_class_name "target_name", "action",
        signedBy "signer_names";
    permission permission_class_name "target_name", "action",
        signedBy "signer_names";
    ...
};
```

All non-italicized items above must appear as is (although case doesn't matter and some are optional, as noted below). Italicized items represent variable values.

A grant entry must begin with the word `grant`.

The SignedBy, Principal, and CodeBase Fields

The `signedBy`, `codeBase`, and `principal` values are optional, and the order of these fields does not matter.

signedBy Value

A `signedBy` value indicates the alias for a certificate stored in the keystore. The public key within that certificate is used to verify the digital signature on the code; you grant the permission(s) to code signed by the private key corresponding to the public key in the keystore entry specified by the alias.

The `signedBy` value can be a comma-separated list of multiple aliases. An example is "Adam,Eve,Charles", which means "signed by Adam and Eve and Charles"; the relationship is AND, not OR. To be more exact, a statement like "Code signed by Adam" means "Code in a class file contained in a JAR which is signed using the private key corresponding to the public key certificate in the keystore whose entry is aliased by Adam".

The `signedBy` field is optional in that, if it is omitted, it signifies "any signer". It doesn't matter whether the code is signed or not or by whom.

principal Value

A principal value specifies a `class_name`/`principal_name` pair which must be present within the executing thread's principal set. The principal set is associated with the executing code by way of a Subject.

The `principal_class_name` may be set to the wildcard value, *, which allows it to match any `Principal` class. In addition, the `principal_name` may also be set to the wildcard value, *, allowing it to match any `Principal` name. When setting the `principal_class_name` or `principal_name` to *, do not surround the * with quotes. Also, if you specify a wildcard principal class, you must also specify a wildcard principal name.

The principal field is optional in that, if it is omitted, it signifies "any principals".

Keystore Alias Replacement

If the principal `class_name`/`principal_name` pair is specified as a single quoted string, then it is treated as a keystore alias. The keystore is consulted and queried (via the alias) for an X509 Certificate. If one is found, the principal `class_name` is automatically treated as `javax.security.auth.x500.X500Principal`, and the `principal_name` is automatically treated as the subject distinguished name from the certificate. If an X509 Certificate mapping is not found, the entire grant entry is ignored.

codeBase Value

A `codeBase` value indicates the code source location; you grant the permission(s) to code from that location. An empty `codeBase` entry signifies "any code"; it doesn't matter where the code originates from.

> **Note:**
>
> A `codeBase` value is a URL and thus should always utilize slashes (never backslashes) as the directory separator, even when the code source is actually on a Windows system. Thus, if the source location for code on a Windows system is actually `C:\somepath\api\`, then the policy `codeBase` entry should look like:
>
> ```
> grant codeBase "file:/C:/somepath/api/" {
>  ...
> };
> ```

The exact meaning of a `codeBase` value depends on the characters at the end. A `codeBase` with a trailing `"/"` matches all class files (not JAR files) in the specified directory. A `codeBase` with a trailing `"/*"` matches all files (both class and JAR files) contained in that directory. A `codeBase` with a trailing `"/-"` matches all files (both class and JAR files) in the directory and recursively all files in subdirectories contained in that directory. The following table illustrates the different cases:

Table 1-8 How Codebase URLs in Downloaded Code Match Those in Policy Files

Codebase URL of Downloaded Code	Codebase URL in Policy File	Match?
www.example.com/usr/ann/	www.example.com/usr/ann	Yes
www.example.com/usr/ann/	www.example.com/usr/ann/	Yes
www.example.com/usr/ann/	www.example.com/usr/ann/*	Yes
www.example.com/usr/ann/	www.example.com/usr/ann/-	Yes
www.example.com/usr/ann/appl.jar	www.example.com/usr/ann/	No
www.example.com/usr/ann/appl.jar	www.example.com/usr/ann/-	Yes
www.example.com/usr/ann/appl.jar	www.example.com/usr/ann/*	Yes
www.example.com/usr/ann/appl.jar	www.example.com/usr/-	Yes
www.example.com/usr/ann/appl.jar	www.example.com/usr/*	No
www.example.com/usr/ann/	www.example.com/usr/-	Yes
www.example.com/usr/ann/	www.example.com/usr/*	No

If you are using a modular runtime image (see the `jlink` tool), you can grant permissions to the application and library modules in the image by specifying a `jrt` URL as the `codeBase` value in a policy file. See JEP 220: Modular Run-Time Images for more information about `jrt` URLs.

The following example grants permission to read the `foo` property to the module `com.greetings`:

```
grant codeBase "jrt:/com.greetings" {
    permission java.util.PropertyPermission "foo", "read";
};
```

The Permission Entries

A **permission entry** must begin with the word `permission`. The word *`permission_class_name`* in the template above would actually be a specific permission type, such as `java.io.FilePermission` or `java.lang.RuntimePermission`.

The "*action*" is required for many permission types, such as `java.io.FilePermission` (where it specifies what type of file access is permitted). It is not required for categories such as `java.lang.RuntimePermission` where it is not necessary, you either have the permission specified by the "*`target_name`*" value following the *permission_class_name* or you don't.

The `signedBy` name/value pair for a permission entry is optional. If present, it indicates a signed permission. That is, the permission class itself must be signed by the given alias(es) in order for the permission to be granted. For example, suppose you have the following grant entry:

```
grant {
    permission Foo "foobar", signedBy "FooSoft";
};
```

Then this permission of type *Foo* is granted if the `Foo.class` permission was placed in a JAR file and the JAR file was signed by the private key corresponding to the public key in the certificate specified by the "FooSoft" alias, or if `Foo.class` is a system class, since system classes are not subject to policy restrictions.

Items that appear in a permission entry must appear in the specified order (`permission`, *permission_class_name*, "*target_name*", "*action*", and `signedBy` "*signer_names*"). An entry is terminated with a semicolon.

Case is unimportant for the identifiers (`permission`, `signedBy`, `codeBase`, etc.) but is significant for the *permission_class_name* or for any string that is passed in as a value.

> **Note:**
>
> See Appendix A: FilePermission Path Name Canonicalization Disabled By Default for important information about a change in how `FilePermission` path names are canonicalized.

File Path Specifications on Windows Systems

When you are specifying a `java.io.FilePermission`, the "*target_name*" is a file path. On Windows systems, whenever you directly specify a file path in a string (but not in a codebase URL), you need to include two backslashes for each actual single backslash in the path, as in

```
    grant {
        permission java.io.FilePermission "C:\\users\\cathy\\foo.bat",
"read";
    };
```

The reason this is necessary is because the strings are processed by a tokenizer (`java.io.StreamTokenizer`), which allows "\" to be used as an escape string (for example, "\n" to indicate a new line) and which thus requires two backslashes to indicate a single backslash. After the tokenizer has processed the above file path string, converting double backslashes to single backslashes, the end result is

```
    "C:\users\cathy\foo.bat"
```

Policy File Examples

The following policy configuration file contains two entries:

```
// If the code is signed by "Duke", grant it read/write access to all
// files in /tmp:
grant signedBy "Duke" {
    permission java.io.FilePermission "/tmp/*", "read,write";
};

// Grant everyone the following permission:
```

```
grant {
    permission java.util.PropertyPermission "java.vendor", "read";
};
```

The following policy configuration file specifies that *only* code that satisfies the following conditions can call methods in the `Security` class to add or remove providers or to set Security Properties:

- The code was loaded from a signed JAR file that is in the "`/home/sysadmin/`" directory on the local file system.
- The signature can be verified using the public key referenced by the alias name "`sysadmin`" in the keystore.

```
grant signedBy "sysadmin", codeBase "file:/home/sysadmin/*" {
    permission java.security.SecurityPermission
"Security.insertProvider.*";
    permission java.security.SecurityPermission
"Security.removeProvider.*";
    permission java.security.SecurityPermission "Security.setProperty.*";
};
```

Either component of the policy entry (or both) may be missing.

The following is a policy configuration file where `codeBase` is missing:

```
grant signedBy "sysadmin" {
    permission java.security.SecurityPermission
"Security.insertProvider.*";
    permission java.security.SecurityPermission
"Security.removeProvider.*";
};
```

If this policy is in effect, then code that comes in a JAR file signed by "`sysadmin`" can add/remove providers, regardless of where the JAR file originated from.

The following is a policy configuration file without a signer:

```
grant codeBase "file:/home/sysadmin/-" {
    permission java.security.SecurityPermission
"Security.insertProvider.*";
    permission java.security.SecurityPermission
"Security.removeProvider.*";
};
```

In this case, code that comes from anywhere in the "`home/sysadmin/`" directory on the local file system can add/remove providers. The code does not need to be signed.

The following is a policy configuration file where neither `codeBase` nor `signedBy` is included:

```
grant {
    permission java.security.SecurityPermission
"Security.insertProvider.*";
```

```
      permission java.security.SecurityPermission
"Security.removeProvider.*";
  };
```

Here, with both code source components missing, any code (regardless of where it originated from, or whether or not it is signed, or who signed it) can add/remove providers.

The following represents a principal-based entry:

```
grant principal javax.security.auth.x500.X500Principal "cn=Alice" {
    permission java.io.FilePermission "/home/Alice", "read, write";
};
```

This permits any code executing as the X500Principal, "`cn=Alice`", permission to read and write to "`/home/Alice`".

The following represents a principal-based entry with a wildcard value:

```
grant principal javax.security.auth.x500.X500Principal * {
    permission java.io.FilePermission "/tmp", "read, write";
};
```

This permits any code executing as an X500Principal (regardless of the distinguished name), permission to read and write to "`/tmp`".

The following example shows a grant statement with both codesource and principal information:

```
grant codebase "http://www.games.example.com",
      signedBy "Duke",
      principal javax.security.auth.x500.X500Principal "cn=Alice" {
    permission java.io.FilePermission "/tmp/games", "read, write";
};
```

This allows code downloaded from "`www.games.example.com`", signed by "`Duke`", and executed by "`cn=Alice`", permission to read and write into the "`/tmp/games`" directory.

The following example shows a grant statement with KeyStore alias replacement:

```
keystore "http://foo.example.com/blah/.keystore";

grant principal "alice" {
    permission java.io.FilePermission "/tmp/games", "read, write";
};
```

"`alice`" will be replaced by the following:

```
  javax.security.auth.x500.X500Principal "cn=Alice"
```

This assumes that X.509 certificate associated with the keystore alias, `alice`, has a subject distinguished name of "`cn=Alice`". This allows code executed by the X500Principal "`cn=Alice`" permission to read and write into the "`/tmp/games`" directory.

Property Expansion in Policy Files

Property expansion is possible in policy files and in the security properties file.

Property expansion is similar to expanding variables in a shell. That is, when a string like

```
${some.property}
```

appears in a policy file, or in the security properties file, it will be expanded to the value of the system property. For example,

```
permission java.io.FilePermission "${user.home}", "read";
```

will expand "`${user.home}`" to use the value of the "user.home" system property. If that property's value is "`/home/cathy`", then the above is equivalent to

```
permission java.io.FilePermission "/home/cathy", "read";
```

In order to assist in platform-independent policy files, you can also use the special notation of "`${/}`", which is a shortcut for `${file.separator}`". This allows things like

```
permission java.io.FilePermission "${user.home}${/}*", "read";
```

If the value of the "`user.home` " property is `/home/cathy`, and you are on Solaris, Linux, or macOS, the above gets converted to:

```
permission java.io.FilePermission "/home/cathy/*", "read";
```

If on the other hand the "`user.home`" value is `C:\users\cathy` and you are on a Windows system, the above gets converted to:

```
permission java.io.FilePermission "C:\users\cathy\*", "read";
```

Also, as a special case, if you expand a property in a codebase, such as

```
grant codeBase "file:${my.libraries}/api/"
```

then any file separator characters will be automatically converted to `/` characters. For example, suppose the value of `my.libraries` is `C:\Users\me\lib`. Thus on a Windows system, the above would get converted to

```
grant codeBase "file:C:/Users/me/lib/api/"
```

Thus you don't need to use `${/}` in codebase strings (and you shouldn't). Property expansion takes place anywhere a double quoted string is allowed in the policy file. This includes the *"signer_names"*, *"URL"*, *"target_name"*, and *"action"* fields. Whether or not property expansion is allowed is controlled by the value of the "`policy.expandProperties`" property in the security properties file. If the value of this property is true (the default), expansion is allowed.

> **Note:**
>
> You can't use nested properties; they will not work. For example,
>
> ```
> "${user.${foo}}"
> ```
>
> doesn't work, even if the "`foo`" property is set to "`home`". The reason is the property parser doesn't recognize nested properties; it simply looks for the first "`${`", and then keeps looking until it finds the first "`}`" and tries to interpret the result (in this case, "`${user.$foo}`") as a property, but fails if there is no such property.

> **Note:**
>
> If a property can't be expanded in a grant entry, permission entry, or keystore entry, that entry is ignored. For example, if the system property "`foo`" is not defined and you have:
>
> ```
> grant codeBase "${foo}" {
>  permission ...;
>  permission ...;
> };
> ```
>
> then all the permissions in this grant entry are ignored. If you have
>
> ```
> grant {
>  permission Foo "${foo}";
>  permission Bar "barTarget";
> };
> ```
>
> then only the "`permission Foo...`" entry is ignored. And finally, if you have
>
> ```
> keystore "${foo}";
> ```
>
> then the keystore entry is ignored.

Windows Systems, File Paths, and Property Expansion

The file path specifications on Windows systems should include two backslashes for each actual single backslash.

As mentioned in File Path Specifications on Windows Systems, on Windows systems, when you directly specify a file path in a string (but not in a codebase URL), you need to include two backslashes for each actual single backslash in the path, as in

```
    grant {
        permission java.io.FilePermission "C:\\users\\cathy\\foo.bat",
"read";
    };
```

This is because the strings are processed by a tokenizer (`java.io.StreamTokenizer`), which allows "\" to be used as an escape string (e.g., "\n" to indicate a new line) and which thus requires two backslashes to indicate a single backslash. After the tokenizer has processed the above file path string, converting double backslashes to single backslashes, the end result is

```
"C:\users\cathy\foo.bat"
```

Expansion of a property in a string takes place after the tokenizer has processed the string. Thus if you have the string

```
"${user.home}\\foo.bat"
```

then first the tokenizer processes the string, converting the double backslashes to a single backslash, and the result is

```
"${user.home}\foo.bat"
```

Then the `${user.home}` property is expanded and the end result is

```
"C:\users\cathy\foo.bat"
```

assuming the "`user.home`" value is `C:\users\cathy`. Of course, for platform independence, it would be better if the string was initially specified without any explicit slashes, i.e., using the `${/}` property instead, as in

```
"${user.home}${/}foo.bat"
```

General Expansion in Policy Files

Generalized forms of expansion are also supported in policy files. For example, permission names may contain a string of the following form:

```
${{protocol:protocol_data}}
```

If such a string occurs in a permission name, then the value in *protocol* determines the exact type of expansion that should occur, and *protocol_data* is used to help perform

the expansion. *protocol_data* may be empty, in which case the above string should simply take the form:

```
${{protocol}}
```

There are two protocols supported in the default policy file implementation:

1. `${{self}}`

   The protocol, **`self`**, denotes a replacement of the entire string, `${{self}}`, with one or more principal class/name pairs. The exact replacement performed depends upon the contents of the grant clause to which the permission belongs.

   If the grant clause does not contain any principal information, the permission will be ignored (permissions containing `${{self}}` in their target names are only valid in the context of a principal-based grant clause). For example, `BarPermission` will always be ignored in the following grant clause:

   ```
   grant codebase "www.example.com", signedby "duke" {
       permission BarPermission "... ${{self}} ...";
   };
   ```

   If the grant clause contains principal information, `${{self}}` will be replaced with that same principal information. For example, `${{self}}` in `BarPermission` will be replaced with **`javax.security.auth.x500.X500Principal "cn=Duke"`** in the following grant clause:

   ```
   grant principal javax.security.auth.x500.X500Principal "cn=Duke" {
       permission BarPermission "... ${{self}} ...";
   };
   ```

   If there is a comma-separated list of principals in the grant clause, then `${{self}}` will be replaced by the same comma-separated list or principals. In the case where both the principal class and name are wildcarded in the grant clause, `${{self}}` is replaced with all the principals associated with the `Subject` in the current `AccessControlContext`.

   The following example describes a scenario involving both **`self`** and Keystore Alias Replacement together:

   ```
   keystore "http://foo.example.com/blah/.keystore";

   grant principal "duke" {
       permission BarPermission "... ${{self}} ...";
   };
   ```

   In the above example, `"duke"` will first be expanded into **`javax.security.auth.x500.X500Principal "cn=Duke"`** assuming the X.509 certificate associated with the `KeyStore` alias, `"duke"`, has a subject distinguished name of `"cn=Duke"`. Next, `${{self}}` will be replaced with the same principal information that was just expanded in the grant clause: **`javax.security.auth.x500.X500Principal "cn=Duke"`**.

2. `${{alias:`*`alias_name`*`}}`

The protocol, **`alias`**, denotes a `java.security.KeyStore` alias substitution. The `KeyStore` used is the one specified in the Keystore Entry. *alias_name* represents an alias into the `KeyStore`. `${{alias:alias_name}}` is replaced with **`javax.security.auth.x500.X500Principal "DN"`**, where *DN* represents the subject distinguished name of the certificate belonging to *alias_name*. For example:

```
keystore "http://foo.example.com/blah/.keystore";

grant codebase "www.example.com" {
    permission BarPermission "... ${{alias:duke}} ...";
};
```

In the above example the X.509 certificate associated with the alias, *duke*, is retrieved from the `KeyStore`, *foo.example.com/blah/.keystore*. Assuming duke's certificate specifies "`o=dukeOrg, cn=duke`" as the subject distinguished name, then `${{alias:duke}}` is replaced with **`javax.security.auth.x500.X500Principal "o=dukeOrg, cn=duke"`**.

The permission entry is ignored under the following error conditions:

- The keystore entry is unspecified
- The *alias_name* is not provided
- The certificate for *alias_name* can not be retrieved
- The certificate retrieved is not an X.509 certificate

Appendix A: FilePermission Path Name Canonicalization Disabled By Default

A canonical path is a path that doesn't contain any links or shortcuts. Performing path name canonicalization in a `FilePermission` object can negatively affect performance.

Before JDK 9, path names were canonicalized when two `FilePermission` objects were compared. This allowed a program to access a file using a different name than the name that was granted to a `FilePermission` object in a policy file, as long as the object pointed to the same file. Because the canonicalization had to access the underlying file system, it could be quite slow.

In JDK 9, path name canonicalization is disabled by default. This means two `FilePermission` objects aren’t equal to each other if one uses an absolute path and the other uses a relative path, or one uses a symbolic link and the other uses a target, or one uses a Windows long name and the other uses a DOS-style 8.3 name. This is true even if they all point to the same file in the file system.

Therefore, if a path name is granted to a `FilePermission` object in a policy file, then the program should also access that file using the same path name style. For example, if the path name in the policy file is using a symbolic link, then the program should also use that symbolic link. Accessing the file with the target path name will fail the permission check.

Compatibility Layer

A compatibility layer has been added to ensure that granting a `FilePermission` object for a relative path will permit applications to access the file with an absolute path (and

conversely). This works for the default Policy provider and the Limited `doPrivileged` calls.

For example, a `FilePermission` object on a file with a relative path name of `"a"` no longer implies a `FilePermission` object on the same file with an absolute path name as `"/pwd/a"` (`"pwd"` is the current working directory). Granting code a `FilePermission` object to read `"a"` allows that code to also read `"/pwd/a"` when a Security Manager is enabled.

The compatibility layer doesn't cover translations between symbolic links and targets, or Windows long names and DOS-style 8.3 names, or any other different name forms that can be canonicalized to the same name.

Customizing Path Name Canonicalization

The system properties in Table 1-9 can be used to customize the `FilePermission` path name canonicalization. See How to Specify a java.lang.System Property.

Table 1-9 System Properties to Customize Path Name Canonicalization

System Property	Default Value	Description
`jdk.io.permissionsUseCanonicalPath`	false	The system property can be used to enable or disable path name canonicalization in the `FilePermission` object. • To disable `FilePermission` path name canonicalization, set `jdk.io.permissionsUseCanonicalPath=false`. • To enable `FilePermission` path name canonicalization, set `jdk.io.permissionsUseCanonicalPath=true`.

Table 1-9 (Cont.) System Properties to Customize Path Name Canonicalization

System Property	Default Value	Description
`jdk.security.filePermCompat`	false	The system property can be used to extend the compatibility layer to support third-party Policy implementations. • To disable the system property, set `jdk.security.filePermCompat=false`. The `FilePermission` for a relative path will permit applications to access the file with an absolute path for the default Policy provider and the Limited `doPrivileged` method. • To extend the compatibility layer to support third-party Policy implementations, set `jdk.security.filePermCompat=true`. The `FilePermission` for a relative path will permit applications to access the file with an absolute path for the default Policy provider, the Limited `doPrivileged` method, and for third-party Policy implementations.

Troubleshooting Security

To monitor security access, you can set the `java.security.debug` system property, which determines what trace messages are printed during execution.

To see a list of all debugging options, use the `help` setting:

```
java -Djava.security.debug=help
```

> **Note:**
>
> To use more than one option, separate options with a comma.
>
> JSSE also provides dynamic debug tracing support for SSL/TLS/DTLS troubleshooting. See Debugging Utilities.

The following table lists `java.security.debug` options and links to further information about each option:

Table 1-10 `java.security.debug` Options

Option	Description	Further Information
`all`	Turn on all the debugging options	None
`access`	Print all results from the `AccessController.checkPermission` method. You can use the following options with the `access` option: 1. `stack`: Include stack trace 2. `domain`: Dump all domains in context 3. `failure`: Before throwing exception, dump stack and domain that do not have permission You can use the following options with the `stack` and `domain` options: 1. `permission=<classname>`: Only dump output if specified permission is being checked 2. `codebase=<URL>`: Only dump output if specified codebase is being checked	Permissions in the JDK
`certpath`	Turns on debugging for the PKIX `CertPathValidator` and `CertPathBuilder` implementations. Use the `ocsp` option with the `certpath` option for OCSP protocol tracing. A hexadecimal dump of the OCSP request and response bytes is displayed.	PKI Programmer's Guide Overview
`combiner`	`SubjectDomainCombiner` debugging	Permissions in the JDK
`configfile`	JAAS (Java Authentication and Authorization Service) configuration file loading	Java Authentication and Authorization Service (JAAS) Reference Guide Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges
`configparser`	JAAS configuration file parsing	Java Authentication and Authorization Service (JAAS) Reference Guide Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges

Table 1-10 (Cont.) `java.security.debug` Options

Option	Description	Further Information
`gssloginconfig`	Java GSS (Generic Security Services) login configuration file debugging	Java Generic Security Services: (Java GSS) and Kerberos JAAS and Java GSS-API Tutorial `javax.security.auth.login.Configuration`: A Configuration object is responsible for specifying which `javax.net.ssl.SSLEngine` should be used for a particular application, and in what order the `LoginModules` should be invoked. Appendix B: JAAS Login Configuration File Advanced Security Programming in Java SE Authentication, Secure Communication and Single Sign-On
`jar`	JAR file verification	Verifying Signed JAR Files from The Java Tutorials
`jca`	JCA engine class debugging	Engine Classes and Algorithms
`keystore`	Keystore debugging	Keystores `KeyStore`
`logincontext`	`LoginContext` results	Java Authentication and Authorization Service (JAAS) Reference Guide Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges
`pkcs11`	PKCS11 session manager debugging	PKCS#11 Reference Guide
`pkcs11keystore`	PKCS11 KeyStore debugging	PKCS#11 Reference Guide
`pkcs12`	PKCS12 KeyStore debugging	None
`policy`	Loading and granting permissions with policy file	Set up the Policy File to Grant the Required Permissions (Controlling Applications) from The Java Tutorials Default Policy Implementation and Policy File Syntax

Table 1-10 (Cont.) `java.security.debug` Options

Option	Description	Further Information
`provider`	Security provider debugging The following options can be used with the provider option: `engine=<engines>` : The output is displayed only for a specified list of JCA engines. The supported values for *<engines>* are: • Cipher • KeyAgreement • KeyGenerator • KeyPairGenerator • KeyStore • Mac • MessageDigest • SecureRandom • Signature	Java Cryptography Architecture (JCA) Reference Guide
`scl`	Permissions that `SecureClassLoader` assigns	Permissions in the JDK
`securerandom`	SecureRandom debugging	The SecureRandom Class
`sunpkcs11`	SunPKCS11 provider debugging	PKCS#11 Reference Guide
`ts`	Timestamping debugging	None

2

Java Cryptography Architecture (JCA) Reference Guide

The Java Cryptography Architecture (JCA) is a major piece of the platform, and contains a "provider" architecture and a set of APIs for digital signatures, message digests (hashes), certificates and certificate validation, encryption (symmetric/asymmetric block/stream ciphers), key generation and management, and secure random number generation, to name a few.

Introduction to Java Cryptography Architecture

The Java platform strongly emphasizes security, including language safety, cryptography, public key infrastructure, authentication, secure communication, and access control.

The JCA is a major piece of the platform, and contains a "provider" architecture and a set of APIs for digital signatures, message digests (hashes), certificates and certificate validation, encryption (symmetric/asymmetric block/stream ciphers), key generation and management, and secure random number generation, to name a few. These APIs allow developers to easily integrate security into their application code. The architecture was designed around the following principles:

- **Implementation independence**: Applications do not need to implement security algorithms. Rather, they can request security services from the Java platform. Security services are implemented in providers (see Cryptographic Service Providers), which are plugged into the Java platform via a standard interface. An application may rely on multiple independent providers for security functionality.
- **Implementation interoperability**: Providers are interoperable across applications. Specifically, an application is not bound to a specific provider, and a provider is not bound to a specific application.
- **Algorithm extensibility**: The Java platform includes a number of built-in providers that implement a basic set of security services that are widely used today. However, some applications may rely on emerging standards not yet implemented, or on proprietary services. The Java platform supports the installation of custom providers that implement such services.

Other cryptographic communication libraries available in the JDK use the JCA provider architecture, but are described elsewhere. The JSSE components provides access to Secure Socket Layer (SSL), Transport Layer Security (TLS), and Datagram Transport Layer Security (DTLS) implementations; see Java Secure Socket Extension (JSSE) Reference Guide. You can use Java Generic Security Services (JGSS) (via Kerberos) APIs, and Simple Authentication and Security Layer (SASL) to securely exchange messages between communicating applications; see Introduction to JAAS and Java GSS-API Tutorials and Java SASL API Programming and Deployment Guide.

Notes on Terminology

- Prior to JDK 1.4, the JCE was an unbundled product, and as such, the JCA and JCE were regularly referred to as separate, distinct components. As JCE is now bundled in the JDK, the distinction is becoming less apparent. Since the JCE uses the same architecture as the JCA, the JCE should be more properly thought of as a part of the JCA.
- The JCA within the JDK includes two software components:
  - The framework that defines and supports cryptographic services for which providers supply implementations. This framework includes packages such as `java.security`, `javax.crypto`, `javax.crypto.spec`, and `javax.crypto.interfaces`.
  - The actual providers such as `Sun`, `SunRsaSign`, `SunJCE`, which contain the actual cryptographic implementations.

  Whenever a specific JCA provider is mentioned, it will be referred to explicitly by the provider's name.

> **WARNING:**
>
> The JCA makes it easy to incorporate security features into your application. However, this document does not cover the theory of security/cryptography beyond an elementary introduction to concepts necessary to discuss the APIs. This document also does not cover the strengths/weaknesses of specific algorithms, not does it cover protocol design. Cryptography is an advanced topic and one should consult a solid, preferably recent, reference in order to make best use of these tools.
> You should always understand what you are doing and why: DO NOT simply copy random code and expect it to fully solve your usage scenario. Many applications have been deployed that contain significant security or performance problems because the wrong tool or algorithm was selected.

JCA Design Principles

The JCA was designed around these principles:

- Implementation independence and interoperability
- Algorithm independence and extensibility

Implementation independence and algorithm independence are complementary; you can use cryptographic services, such as digital signatures and message digests, without worrying about the implementation details or even the algorithms that form the basis for these concepts. While complete algorithm-independence is not possible, the JCA provides standardized, algorithm-specific APIs. When implementation-independence is not desirable, the JCA lets developers indicate a specific implementation.

Algorithm independence is achieved by defining types of cryptographic "engines" (services), and defining classes that provide the functionality of these cryptographic engines. These classes are called *engine classes*, and examples are the

`MessageDigest`, `Signature`, `KeyFactory`, `KeyPairGenerator`, and `Cipher` classes.

Implementation independence is achieved using a "provider"-based architecture. The term Cryptographic Service Provider (CSP), which is used interchangeably with the term "provider," (see Cryptographic Service Providers) refers to a package or set of packages that implement one or more cryptographic services, such as digital signature algorithms, message digest algorithms, and key conversion services. A program may simply request a particular type of object (such as a `Signature` object) implementing a particular service (such as the DSA signature algorithm) and get an implementation from one of the installed providers. If desired, a program may instead request an implementation from a specific provider. Providers may be updated transparently to the application, for example when faster or more secure versions are available.

Implementation interoperability means that various implementations can work with each other, use each other's keys, or verify each other's signatures. This would mean, for example, that for the same algorithms, a key generated by one provider would be usable by another, and a signature generated by one provider would be verifiable by another.

Algorithm extensibility means that new algorithms that fit in one of the supported engine classes can be added easily.

Provider Architecture

Providers contain a package (or a set of packages) that supply concrete implementations for the advertised cryptographic algorithms.

Cryptographic Service Providers

`java.security.Provider` is the base class for all security providers. Each CSP contains an instance of this class which contains the provider's name and lists all of the security services/algorithms it implements. When an instance of a particular algorithm is needed, the JCA framework consults the provider's database, and if a suitable match is found, the instance is created.

Providers contain a package (or a set of packages) that supply concrete implementations for the advertised cryptographic algorithms. Each JDK installation has one or more providers installed and configured by default. Additional providers may be added statically or dynamically. Clients may configure their runtime environment to specify the provider *preference order*. The preference order is the order in which providers are searched for requested services when no specific provider is requested.

To use the JCA, an application simply requests a particular type of object (such as a `MessageDigest`) and a particular algorithm or service (such as the "SHA-256" algorithm), and gets an implementation from one of the installed providers. For example, the following statement requests a SHA-256 message digest from an installed provider:

```
md = MessageDigest.getInstance("SHA-256");
```

Alternatively, the program can request the objects from a specific provider. Each provider has a name used to refer to it. For example, the following statement requests a SHA-256 message digest from the provider named ProviderC:

```
md = MessageDigest.getInstance("SHA-256", "ProviderC");
```

The following figures illustrates requesting an SHA-256 message digest implementation. They show three different providers that implement various message digest algorithms (SHA-256, SHA-384, and SHA-512). The providers are ordered by preference from left to right (1-3). In Figure 2-1, an application requests a SHA-256 algorithm implementation **without** specifying a provider name. The providers are searched in preference order and the implementation from the first provider supplying that particular algorithm, ProviderB, is returned. In Figure 2-2, the application requests the SHA-256 algorithm implementation **from a specific provider**, ProviderC. This time, the implementation from ProviderC is returned, even though a provider with a higher preference order, ProviderB, also supplies an MD5 implementation.

Figure 2-1 Request SHA-256 Message Digest Implementation Without Specifying Provider

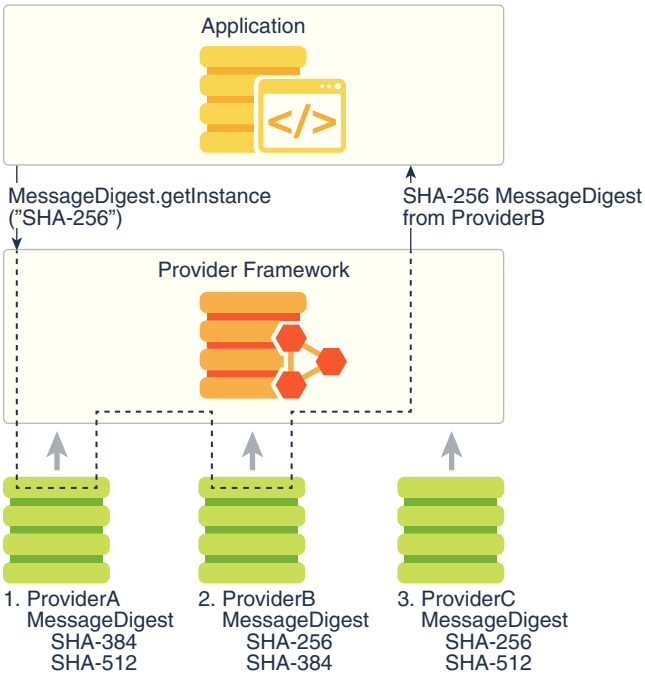

Figure 2-2 Request SHA-256 Message Digest with ProviderC

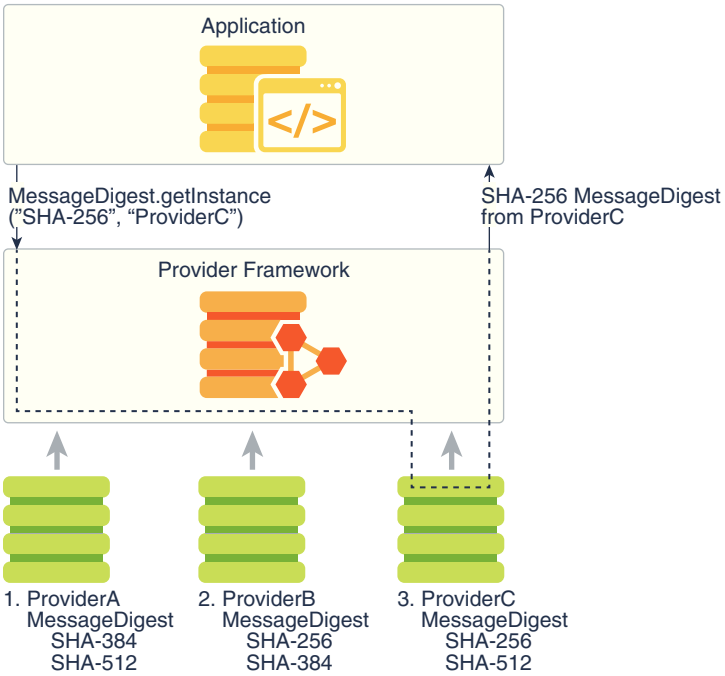

Cryptographic implementations in the JDK are distributed via several different providers (`Sun`, `SunJSSE`, `SunJCE`, `SunRsaSign`) primarily for historical reasons, but to a lesser extent by the type of functionality and algorithms they provide. Other Java runtime environments may not necessarily contain these providers, so applications should not request a provider-specific implementation unless it is known that a particular provider will be available.

The JCA offers a set of APIs that allow users to query which providers are installed and what services they support.

This architecture also makes it easy for end-users to add additional providers. Many third party provider implementations are already available. See The Provider Class for more information on how providers are written, installed, and registered.

How Providers Are Actually Implemented

Algorithm independence is achieved by defining a generic high-level Application Programming Interface (API) that all applications use to access a service type. *Implementation independence* is achieved by having all provider implementations conform to well-defined interfaces. Instances of engine classes are thus "backed" by implementation classes which have the same method signatures. Application calls are routed through the engine class and are delivered to the underlying backing implementation. The implementation handles the request and return the proper results.

The application API methods in each engine class are routed to the provider's implementations through classes that implement the corresponding Service Provider Interface (SPI). That is, for each engine class, there is a corresponding abstract SPI class which defines the methods that each cryptographic service provider's algorithm must implement. The name of each SPI class is the same as that of the corresponding engine class, followed by `Spi`. For example, the `Signature` engine class provides

access to the functionality of a digital signature algorithm. The actual provider implementation is supplied in a subclass of `SignatureSpi`. Applications call the engine class' API methods, which in turn call the SPI methods in the actual implementation.

Each SPI class is abstract. To supply the implementation of a particular type of service for a specific algorithm, a provider must subclass the corresponding SPI class and provide implementations for all the abstract methods.

For each engine class in the API, implementation instances are requested and instantiated by calling the `getInstance()` *factory method* in the engine class. A factory method is a static method that returns an instance of a class. The engine classes use the framework provider selection mechanism described above to obtain the actual backing implementation (SPI), and then creates the actual engine object. Each instance of the engine class encapsulates (as a private field) the instance of the corresponding SPI class, known as the SPI object. All API methods of an API object are declared final and their implementations invoke the corresponding SPI methods of the encapsulated SPI object.

To make this clearer, review Example 2-1 and Figure 2-3:

Example 2-1 Sample Code for Getting an Instance of an Engine Class

```
import javax.crypto.*;

Cipher c = Cipher.getInstance("AES");
c.init(ENCRYPT_MODE, key);
```

Figure 2-3 Application Retrieves “AES” Cipher Instance

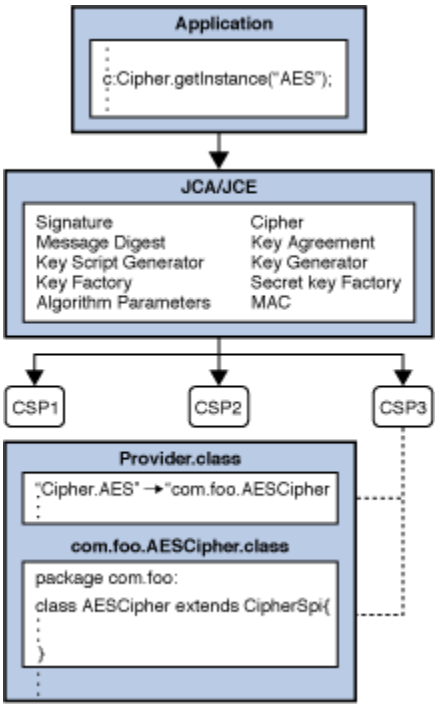

Here an application wants an "AES" `javax.crypto.Cipher` instance, and doesn't care which provider is used. The application calls the `getInstance()` factory methods of the `Cipher` engine class, which in turn asks the JCA framework to find the first provider instance that supports "AES". The framework consults each installed provider, and obtains the provider's instance of the `Provider` class. (Recall that the `Provider` class is a database of available algorithms.) The framework searches each provider, finally finding a suitable entry in CSP3. This database entry points to the implementation class `com.foo.AESCipher` which extends `CipherSpi`, and is thus suitable for use by the `Cipher` engine class. An instance of `com.foo.AESCipher` is created, and is encapsulated in a newly-created instance of `javax.crypto.Cipher`, which is returned to the application. When the application now does the `init()` operation on the `Cipher` instance, the `Cipher` engine class routes the request into the corresponding `engineInit()` backing method in the `com.foo.AESCipher` class.

Java Security Standard Algorithm Names lists the Standard Names defined for the Java environment. Other third-party providers may define their own implementations of these services, or even additional services.

Keystores

A database called a "keystore" can be used to manage a repository of keys and certificates. Keystores are available to applications that need data for authentication, encryption, or signing purposes.

Applications can access a keystore via an implementation of the `KeyStore` class, which is in the `java.security` package. As of JDK 9, the default and recommended keystore type (format) is "pkcs12", which is based on the RSA PKCS12 Personal Information Exchange Syntax Standard. Previously, the default keystore type was "jks", which is a proprietary format. Other keystore formats are available, such as "jceks", which is an alternate proprietary keystore format, and "pkcs11", which is based on the RSA PKCS11 Standard and supports access to cryptographic tokens such as hardware security modules and smartcards.

Applications can choose different keystore implementations from different providers, using the same provider mechanism described previously. See Key Management.

Engine Classes and Algorithms

An engine class provides the interface to a specific type of cryptographic service, independent of a particular cryptographic algorithm or provider.

The engines provides one of the following:

- cryptographic operations (encryption, digital signatures, message digests, etc.),
- generators or converters of cryptographic material (keys and algorithm parameters), or
- objects (keystores or certificates) that encapsulate the cryptographic data and can be used at higher layers of abstraction.

The following engine classes are available:

- `SecureRandom`: used to generate random or pseudo-random numbers.
- `MessageDigest`: used to calculate the message digest (hash) of specified data.

- `Signature`: initialized with keys, these are used to sign data and verify digital signatures.
- `Cipher`: initialized with keys, these used for encrypting/decrypting data. There are various types of algorithms: symmetric bulk encryption (e.g. AES), asymmetric encryption (e.g. RSA), and password-based encryption (e.g. PBE).
- Message Authentication Codes (MAC): like `MessageDigests`, these also generate hash values, but are first initialized with keys to protect the integrity of messages.
- `KeyFactory`: used to convert existing opaque cryptographic keys of type `Key` into key specifications (transparent representations of the underlying key material), and vice versa.
- `SecretKeyFactory`: used to convert existing opaque cryptographic keys of type `SecretKey` into key specifications (transparent representations of the underlying key material), and vice versa. `SecretKeyFactorys` are specialized `KeyFactorys` that create secret (symmetric) keys only.
- `KeyPairGenerator`: used to generate a new pair of public and private keys suitable for use with a specified algorithm.
- `KeyGenerator`: used to generate new secret keys for use with a specified algorithm.
- `KeyAgreement`: used by two or more parties to agree upon and establish a specific key to use for a particular cryptographic operation.
- `AlgorithmParameters`: used to store the parameters for a particular algorithm, including parameter encoding and decoding.
- `AlgorithmParameterGenerator` : used to generate a set of AlgorithmParameters suitable for a specified algorithm.
- `KeyStore`: used to create and manage a *keystore*. A keystore is a database of keys. Private keys in a keystore have a certificate chain associated with them, which authenticates the corresponding public key. A keystore also contains certificates from trusted entities.
- `CertificateFactory`: used to create public key certificates and Certificate Revocation Lists (CRLs).
- `CertPathBuilder`: used to build certificate chains (also known as certification paths).
- `CertPathValidator`: used to validate certificate chains.
- `CertStore`: used to retrieve `Certificates` and `CRLs` from a repository.

> **Note:**
>
> A *generator* creates objects with brand-new contents, whereas a *factory* creates objects from existing material (for example, an encoding).

Core Classes and Interfaces

The following are the core classes and interfaces provided in the JCA.

- `Provider` and `Security`

- `SecureRandom`, `MessageDigest`, `Signature`, `Cipher`, `Mac`, `KeyFactory`, `SecretKeyFactory`, `KeyPairGenerator`, `KeyGenerator`, `KeyAgreement`, `AlgorithmParameter`, `AlgorithmParameterGenerator`, `KeyStore`, `CertificateFactory`, and engine
- `Key Interface`, `KeyPair`
- `AlgorithmParameterSpec Interface`, `AlgorithmParameters`, `AlgorithmParameterGenerator`, and algorithm parameter specification interfaces and classes in the `java.security.spec` and `javax.crypto.spec` packages.
- `KeySpec Interface`, `EncodedKeySpec`, `PKCS8EncodedKeySpec`, and `X509EncodedKeySpec`.
- `SecretKeyFactory`, `KeyFactory`, `KeyPairGenerator`, `KeyGenerator`, `KeyAgreement`, and `KeyStore`.

> **Note:**
>
> See `CertPathBuilder`, `CertPathValidator`, and `CertStoreengine` classes in the Java PKI Programmer's Guide.

The guide will cover the most useful high-level classes first (Provider, Security, SecureRandom, MessageDigest, Signature, Cipher, and Mac), then delve into the various support classes. For now, it is sufficient to simply say that Keys (public, private, and secret) are generated and represented by the various JCA classes, and are used by the high-level classes as part of their operation.

This section shows the signatures of the main methods in each class and interface. Examples for some of these classes (MessageDigest, Signature, KeyPairGenerator, SecureRandom, KeyFactory, and key specification classes) are supplied in the corresponding Code Examples sections.

The complete reference documentation for the relevant Security API packages can be found in the package summaries:

- `java.security`
- `javax.crypto`
- `java.security.cert`
- `java.security.spec`
- `javax.crypto.spec`
- `java.security.interfaces`
- `javax.crypto.interfaces`

The Provider Class

The term "Cryptographic Service Provider" (used interchangeably with "provider" in this document) refers to a package or set of packages that supply a concrete implementation of a subset of the JDK Security API cryptography features. The `Provider` *class* is the interface to such a package or set of packages. It has methods for accessing the provider name, version number, and other information. Please note that in addition to registering implementations of cryptographic services, the `Provider`

class can also be used to register implementations of other security services that might get defined as part of the JDK Security API or one of its extensions.

To supply implementations of cryptographic services, an entity (e.g., a development group) writes the implementation code and creates a subclass of the `Provider` class. The constructor of the `Provider` subclass sets the values of various properties; the JDK Security API uses these values to look up the services that the provider implements. In other words, the subclass specifies the names of the classes implementing the services.

Figure 2-4 Provider Class

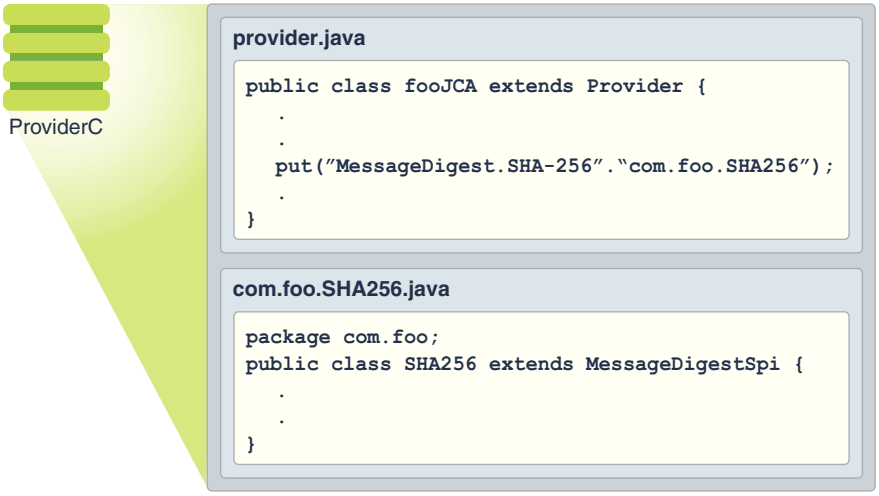

There are several types of services that can be implemented by provider packages; See Engine Classes and Algorithms.

The different implementations may have different characteristics. Some may be software-based, while others may be hardware-based. Some may be platform-independent, while others may be platform-specific. Some provider source code may be available for review and evaluation, while some may not. The JCA lets both end-users and developers decide what their needs are.

You can find information about how end-users install the cryptography implementations that fit their needs, and how developers request the implementations that fit theirs.

> **Note:**
>
> To implement a provider, see Steps to Implement and Integrate a Provider.

How Provider Implementations Are Requested and Supplied

For each engine class (see Engine Classes and Algorithms) in the API, a implementation instance is requested and instantiated by calling one of the `getInstance` methods on the engine class, specifying the name of the desired

algorithm and, optionally, the name of the provider (or the `Provider` class) whose implementation is desired.

```
static EngineClassName getInstance(String algorithm)
    throws NoSuchAlgorithmException

static EngineClassName getInstance(String algorithm, String provider)
    throws NoSuchAlgorithmException, NoSuchProviderException

static EngineClassName getInstance(String algorithm, Provider provider)
    throws NoSuchAlgorithmException
```

where

EngineClassName

is the desired engine type (MessageDigest/Cipher/etc). For example:

```
MessageDigest md = MessageDigest.getInstance("SHA-256");
KeyAgreement ka = KeyAgreement.getInstance("DH", "SunJCE");
```

return an instance of the "SHA-256" MessageDigest and "DH" KeyAgreement objects, respectively.

Java Security Standard Algorithm Names contains the list of names that have been standardized for use with the Java environment. Some providers may choose to also include alias names that also refer to the same algorithm. For example, the "SHA256" algorithm might be referred to as "SHA-256". Applications should use standard names instead of an alias, as not all providers may alias algorithm names in the same way.

> **Note:**
>
> The algorithm name is not case-sensitive. For example, all the following calls are equivalent:
>
> ```
> MessageDigest.getInstance("SHA256")
> MessageDigest.getInstance("sha256")
> MessageDigest.getInstance("sHa256")
> ```

If no provider is specified, `getInstance` searches the registered providers for an implementation of the requested cryptographic service associated with the named algorithm. In any given Java Virtual Machine (JVM), providers are installed in a given *preference order*, the order in which the provider list is searched if a specific provider is not requested. (See Installing Providers.) For example, suppose there are two providers installed in a JVM, `PROVIDER_1` and `PROVIDER_2`. Assume that:

- `PROVIDER_1` implements SHA-256 and DESede. `PROVIDER_1` has preference order 1 (the highest priority).

- `PROVIDER_2` implements SHA256withDSA, SHA-256, RC5, and RSA. `PROVIDER_2` has preference order 2.

Now let's look at three scenarios:

1. If we are looking for an SHA-256 implementation. Both providers supply such an implementation. The `PROVIDER_1` implementation is returned since `PROVIDER_1` has the highest priority and is searched first.
2. If we are looking for an SHA256withDSA signature algorithm, `PROVIDER_1` is first searched for it. No implementation is found, so `PROVIDER_2` is searched. Since an implementation is found, it is returned.
3. Suppose we are looking for a SHA256withRSA signature algorithm. Since no installed provider implements it, a `NoSuchAlgorithmException` is thrown.

The `getInstance` methods that include a provider argument are for developers who want to specify which provider they want an algorithm from. A federal agency, for example, will want to use a provider implementation that has received federal certification. Let's assume that the SHA256withDSA implementation from `PROVIDER_1` has not received such certification, while the DSA implementation of `PROVIDER_2` has received it.

A federal agency program would then have the following call, specifying `PROVIDER_2` since it has the certified implementation:

```
Signature dsa = Signature.getInstance("SHA256withDSA", "PROVIDER_2");
```

In this case, if `PROVIDER_2` was not installed, a `NoSuchProviderException` would be thrown, even if another installed provider implements the algorithm requested.

A program also has the option of getting a list of all the installed providers (using the `getProviders` method in The Security Class class) and choosing one from the list.

> **Note:**
>
> General purpose applications **SHOULD NOT** request cryptographic services from specific providers. Otherwise, applications are tied to specific providers which may not be available on other Java implementations. They also might not be able to take advantage of available optimized providers (for example hardware accelerators via PKCS11 or native OS implementations such as Microsoft's MSCAPI) that have a higher preference order than the specific requested provider.

Installing Providers

In order to be used, a cryptographic provider must first be installed, then registered either statically or dynamically. There are a variety of Sun providers shipped with this release (`SUN`, `SunJCE`, `SunJSSE`, `SunRsaSign`, etc.) that are already installed and registered. The following sections describe how to install and register additional providers.

All JDK providers are already installed and registered. However, if you require any third-party providers, see Step 8: Prepare for Testing from Steps to Implement and

Integrate a Provider for information about how to add providers to the class or module path, register providers (statically or dynamically), and add any required permissions.

Provider Class Methods

Each `Provider` class instance has a (currently case-sensitive) name, a version number, and a string description of the provider and its services.

You can query the `Provider` instance for this information by calling the following methods:

```
public String getName()
public double getVersion()
public String getInfo()
```

The Security Class

The `Security` class manages installed providers and security-wide properties. It only contains static methods and is never instantiated. The methods for adding or removing providers, and for setting `Security` properties, can only be executed by a trusted program. Currently, a "trusted program" is either

- A local application not running under a security manager, or
- An applet or application with permission to execute the specified method (see below).

The determination that code is considered trusted to perform an attempted action (such as adding a provider) requires that the applet is granted the proper permission(s) for that particular action. The policy configuration file(s) for a JDK installation specify what permissions (which types of system resource accesses) are allowed by code from specified code sources. (See below and Default Policy Implementation and Policy File Syntax and Java SE Platform Security Architecture.)

Code being executed is always considered to come from a particular "code source". The code source includes not only the location (URL) where the code originated from, but also a reference to any public key(s) corresponding to the private key(s) that may have been used to sign the code. Public keys in a code source are referenced by (symbolic) alias names from the user's .

In a policy configuration file, a code source is represented by two components: a code base (URL), and an alias name (preceded by `signedBy`), where the alias name identifies the keystore entry containing the public key that must be used to verify the code's signature.

Each "grant" statement in such a file grants a specified code source a set of permissions, specifying which actions are allowed.

Here is a sample policy configuration file:

```
grant codeBase "file:/home/sysadmin/", signedBy "sysadmin" {
    permission java.security.SecurityPermission "insertProvider";
    permission java.security.SecurityPermission "removeProvider";
    permission java.security.SecurityPermission "putProviderProperty.*";
};
```

This configuration file specifies that code loaded from a signed JAR file in the `/home/sysadmin/` directory on the local file system can add or remove providers or set provider properties. (Note that the signature of the JAR file can be verified using the public key referenced by the alias name `sysadmin` in the user's keystore.).

Either component of the code source (or both) may be missing. Here's an example of a configuration file where the `codeBase` is omitted:

```
grant signedBy "sysadmin" {
    permission java.security.SecurityPermission "insertProvider.*";
    permission java.security.SecurityPermission "removeProvider.*";
};
```

If this policy is in effect, code that comes in a JAR File signed by `/home/sysadmin/` directory on the local filesystem can add or remove providers. The code does not need to be signed.

An example where neither `codeBase` nor `signedBy` is included is:

```
grant {
    permission java.security.SecurityPermission "insertProvider.*";
    permission java.security.SecurityPermission "removeProvider.*";
};
```

Here, with both code source components missing, any code (regardless of where it originates, or whether or not it is signed, or who signed it) can add/remove providers. Obviously, this is definitely not recommended, as this grant could open a security hole. Untrusted code could install a Provider, thus affecting later code that is depending on a properly functioning implementation. (For example, a rogue `Cipher` object might capture and store the sensitive information it receives.)

Managing Providers

The following tables summarize the methods in the `Security` class you can use to query which `Providers` are installed, as well as to install or remove providers at runtime.

Querying Providers

Method	Description
`static Provider[] getProviders()`	Returns an array containing all the installed providers (technically, the `Provider` subclass for each package provider). The order of the `Providers` in the array is their preference order.
`static Provider getProvider (String providerName)`	Returns the `Provider` named `providerName`. It returns `null` if the `Provider` is not found.

Adding Providers

Method	Description
`static int addProvider(Provider provider)`	Adds a `Provider` to the end of the list of installed `Providers`. It returns the preference position in which the `Provider` was added, or `-1` if the `Provider` was not added because it was already installed.
`static int insertProviderAt (Provider provider, int position)`	Adds a new `Provider` at a specified position. If the given provider is installed at the requested position, the provider formerly at that position and all providers with a position greater than `position` are shifted up one position (towards the end of the list). This method returns the preference position in which the `Provider` was added, or `-1` if the `Provider` was not added because it was already installed.

Removing Providers

Method	Description
`static void removeProvider(String name)`	Removes the `Provider` with the specified name. It returns silently if the provider is not installed. When the specified provider is removed, all providers located at a position greater than where the specified provider was are shifted down one position (towards the head of the list of installed providers).

> **Note:**
>
> If you want to change the preference position of a provider, you must first remove it, and then insert it back in at the new preference position.

Security Properties

The `Security` class maintains a list of system-wide Security Properties. These properties are similar to the `System` properties, but are security-related. These properties can be set statically (through the `<`*`java-home`*`>/conf/security/java.security` file) or dynamically (using an API). See Step 8.1: Configure the Provider from Steps to Implement and Integrate a Provider. for an example of registering a provider statically with the `security.provider.`*`n`* Security Property. If you want to set properties dynamically, trusted programs can use the following methods:

```
static String getProperty(String key)
static void setProperty(String key, String datum)
```

> **Note:**
>
> The list of security providers is established during VM startup; therefore, the methods described above must be used to alter the provider list.

The SecureRandom Class

The SecureRandom class is an engine class (see Engine Classes and Algorithms) that provides cryptographically strong random numbers, either by accessing a pseudo-random number generator (PRNG), a deterministic algorithm that produces a pseudo-random sequence from an initial seed value, or by reading a native source of randomness (for example, `/dev/random` or a true random number generator). One example of a PRNG is the Deterministic Random Bits Generator (DRBG) as specified in NIST SP 800-90Ar1. Other implementations may produce true random numbers, and yet others may use a combination of both techniques. A cryptographically strong random number minimally complies with the statistical random number generator tests specified in FIPS 140-2, Security Requirements for Cryptographic Modules, section 4.9.1.

All Java SE implementations must indicate the strongest (most random) implementation of SecureRandom that they provide in the `securerandom.strongAlgorithms` property of the `java.security.Security` class. This implementation can be used when a particularly strong random value is required.

The `securerandom.drbg.config` property is used to specify the DRBG `SecureRandom` configuration and implementations in the SUN provider. The `securerandom.drbg.config` is a property of the `java.security.Security` class. Other DRBG implementations can also use the `securerandom.drbg.config` property.

Figure 2-5 SecureRandom class

Creating a SecureRandom Object

There are several ways to obtain an instance of `SecureRandom`:

- All Java SE implementations provide a default `SecureRandom` using the no-argument constructor: `new SecureRandom()`. This constructor traverses the list of registered security providers, starting with the most preferred provider, then returns a new `SecureRandom` object from the first provider that supports a `SecureRandom` random number generator (RNG) algorithm. If none of the providers support a RNG algorithm, then it returns a `SecureRandom` object that uses SHA1PRNG from the SUN provider.
- To get a specific implementation of `SecureRandom`, use one of the How Provider Implementations Are Requested and Supplied.
- Use the `getInstanceStrong()` method to obtain a strong `SecureRandom` implementation as defined by the `securerandom.strongAlgorithms` property of

the `java.security.Security` class. This property lists platform implementations that are suitable for generating important values.

Seeding or Re-Seeding the SecureRandom Object

The `SecureRandom` object is initialized with a random seed unless the call to `getInstance()` is followed by a call to one of the following `setSeed` methods.

```
void setSeed(byte[] seed)
void setSeed(long seed)
```

You must call `setSeed` before the first `nextBytes` call to prevent any environmental randomness.

The randomness of the bits produced by the `SecureRandom` object depends on the randomness of the seed bits

At any time a `SecureRandom` object may be re-seeded using one of the `setSeed` or `reseed` methods. The given seed for `setSeed` supplements, rather than replaces, the existing seed; therefore, repeated calls are guaranteed never to reduce randomness.

Using a SecureRandom Object

To get random bytes, a caller simply passes an array of any length, which is then filled with random bytes:

```
void nextBytes(byte[] bytes)
```

Generating Seed Bytes

If desired, it is possible to invoke the `generateSeed` method to generate a given number of seed bytes (to seed other random number generators, for example):

```
byte[] generateSeed(int numBytes)
```

The MessageDigest Class

The `MessageDigest` class is an engine class (see Engine Classes and Algorithms) designed to provide the functionality of cryptographically secure message digests such as SHA-256 or SHA-512. A cryptographically secure message digest takes arbitrary-sized input (a byte array), and generates a fixed-size output, called a *digest* or hash.

Figure 2-6 MessageDigest Class

For example, the SHA-256 algorithm produces a 32-byte digest, and SHA-512's is 64 bytes.

A digest has two properties:

- It should be computationally infeasible to find two messages that hash to the same value.
- The digest should not reveal anything about the input that was used to generate it.

Message digests are used to produce unique and reliable identifiers of data. They are sometimes called "checksums" or the "digital fingerprints" of the data. Changes to just one bit of the message should produce a different digest value.

Message digests have many uses and can determine when data has been modified, intentionally or not. Recently, there has been considerable effort to determine if there are any weaknesses in popular algorithms, with mixed results. When selecting a digest algorithm, one should always consult a recent reference to determine its status and appropriateness for the task at hand.

Creating a MessageDigest Object

Procedure to create a `MessageDigest` object.

- To compute a digest, create a message digest instance. The `MessageDigest` objects are obtained by using one of the `getInstance()` methods in the `MessageDigest` class. See How Provider Implementations Are Requested and Supplied.

  The factory method returns an initialized message digest object. It thus does not need further initialization.

Updating a Message Digest Object

Procedure to update the Message Digest object.

- To calculate the digest of some data, you have to supply the data to the initialized message digest object. It can be provided all at once, or in chunks. Pieces can be fed to the message digest by calling one of the `update` methods:

```
void update(byte input)
void update(byte[] input)
void update(byte[] input, int offset, int len)
```

Computing the Digest

Procedure to compute the digest using different types of `digest()` methods.

The data chunks have to be supplied by calls to `update` method. See Updating a Message Digest Object.

- The digest is computed using a call to one of the `digest` methods:

```
byte[] digest()
```

```
byte[] digest(byte[] input)
int digest(byte[] buf, int offset, int len)
```

1. The `byte[] digest()` method return the computed digest.
2. The `byte[] digest(byte[] input)` method does a final `update(input)` with the input byte array before calling `digest()`, which returns the digest byte array.
3. The `int digest(byte[] buf, int offset, int len)` method stores the computed digest in the provided buffer `buf`, starting at `offset`. `len` is the number of bytes in `buf` allotted for the digest, the method returns the number of bytes actually stored in `buf`. If there is not enough room in the buffer, the method will throw an exception.

See Computing a MessageDigest Object.

The Signature Class

The `Signature` class is an engine class (see Engine Classes and Algorithms(designed to provide the functionality of a cryptographic digital signature algorithm such as SHA256withDSA or SHA512withRSA. A cryptographically secure signature algorithm takes arbitrary-sized input and a private key and generates a relatively short (often fixed-size) string of bytes, called the *signature*, with the following properties:

- Only the owner of a private/public key pair is able to create a signature. It should be computationally infeasible for anyone having only the public key and a number of signatures to recover the private key.
- Given the public key corresponding to the private key used to generate the signature, it should be possible to verify the authenticity and integrity of the input.

Figure 2-7 Signature Class

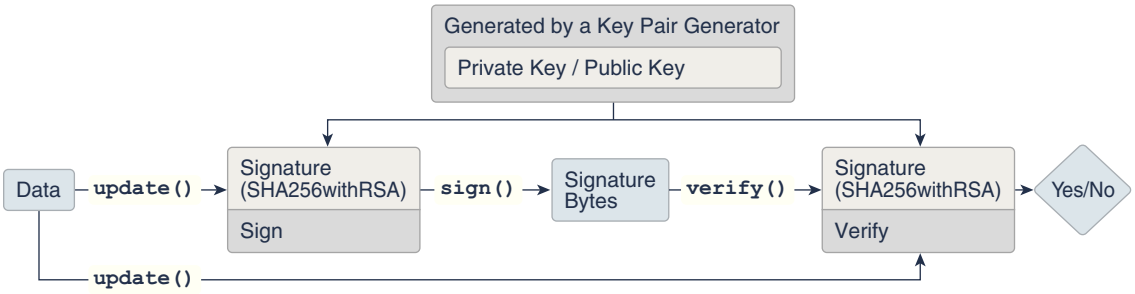

A `Signature` object is initialized for signing with a Private Key and is given the data to be signed. The resulting signature bytes are typically kept with the signed data. When verification is needed, another `Signature` object is created and initialized for verification and given the corresponding Public Key. The data and the signature bytes are fed to the signature object, and if the data and signature match, the `Signature` object reports success.

Even though a signature seems similar to a message digest, they have very different purposes in the type of protection they provide. In fact, algorithms such as "SHA256WithRSA" use the message digest "SHA256" to initially "compress" the large

data sets into a more manageable form, then sign the resulting 32 byte message digest with the "RSA" algorithm.

For an example for signing and verifying data, see Generating and Verifying a Signature Using Generated Keys.

Signature Object States

`Signature` objects are modal objects. This means that a `Signature` object is always in a given state, where it may only do one type of operation.

States are represented as final integer constants defined in their respective classes.

The three states a `Signature` object may have are:

- `UNINITIALIZED`
- `SIGN`
- `VERIFY`

When it is first created, a `Signature` object is in the `UNINITIALIZED` state. The `Signature` class defines two initialization methods, `initSign` and `initVerify`, which change the state to `SIGN` and `VERIFY` , respectively.

Creating a Signature Object

The first step for signing or verifying a signature is to create a `Signature` instance.

`Signature` objects are obtained by using one of the `Signature` `getInstance()` static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a Signature Object

A `Signature` object must be initialized before it is used. The initialization method depends on whether the object is going to be used for signing or for verification.

If it is going to be used for signing, the object must first be initialized with the private key of the entity whose signature is going to be generated. This initialization is done by calling the method:

```
final void initSign(PrivateKey privateKey)
```

This method puts the `Signature` object in the `SIGN` state. If instead the `Signature` object is going to be used for verification, it must first be initialized with the public key of the entity whose signature is going to be verified. This initialization is done by calling either of these methods:

```
final void initVerify(PublicKey publicKey)

final void initVerify(Certificate certificate)
```

This method puts the `Signature` object in the `VERIFY` state.

Signing with a Signature Object

If the `Signature` object has been initialized for signing (if it is in the `SIGN` state), the data to be signed can then be supplied to the object. This is done by making one or more calls to one of the `update` methods:

```
final void update(byte b)
final void update(byte[] data)
final void update(byte[] data, int off, int len)
```

Calls to the `update` method(s) should be made until all the data to be signed has been supplied to the `Signature` object.

To generate the signature, simply call one of the `sign` methods:

```
final byte[] sign()
final int sign(byte[] outbuf, int offset, int len)
```

The first method returns the signature result in a byte array. The second stores the signature result in the provided buffer *outbuf*, starting at *offset*. *len* is the number of bytes in *outbuf* allotted for the signature. The method returns the number of bytes actually stored.

Signature encoding is algorithm specific. See Java Security Standard Algorithm Names to know more about the use of ASN.1 encoding in the Java Cryptography Architecture.

A call to a `sign` method resets the signature object to the state it was in when previously initialized for signing via a call to `initSign`. That is, the object is reset and available to generate another signature with the same private key, if desired, via new calls to `update` and `sign`.

Alternatively, a new call can be made to `initSign` specifying a different private key, or to `initVerify` (to initialize the `Signature` object to verify a signature).

Verifying with a Signature Object

If the `Signature` object has been initialized for verification (if it is in the `VERIFY` state), it can then verify if an alleged signature is in fact the authentic signature of the data associated with it. To start the process, the data to be verified (as opposed to the signature itself) is supplied to the object. The data is passed to the object by calling one of the `update` methods:

```
final void update(byte b)
final void update(byte[] data)
final void update(byte[] data, int off, int len)
```

Calls to the `update` method(s) should be made until all the data to be verified has been supplied to the `Signature` object. The signature can now be verified by calling one of the `verify` methods:

```
final boolean verify(byte[] signature)

final boolean verify(byte[] signature, int offset, int length)
```

The argument must be a byte array containing the signature. This byte array would hold the signature bytes which were returned by a previous call to one of the `sign` methods.

The `verify` method returns a `boolean` indicating whether or not the encoded signature is the authentic signature of the data supplied to the `update` method(s).

A call to the `verify` method resets the signature object to its state when it was initialized for verification via a call to `initVerify`. That is, the object is reset and available to verify another signature from the identity whose public key was specified in the call to `initVerify`.

Alternatively, a new call can be made to `initVerify` specifying a different public key (to initialize the `Signature` object for verifying a signature from a different entity), or to `initSign` (to initialize the `Signature` object for generating a signature).

The Cipher Class

The `Cipher` class provides the functionality of a cryptographic cipher used for encryption and decryption. Encryption is the process of taking data (called *cleartext*) and a *key*, and producing data (*ciphertext*) meaningless to a third-party who does not know the key. Decryption is the inverse process: that of taking ciphertext and a key and producing cleartext.

Figure 2-8 The Cipher Class

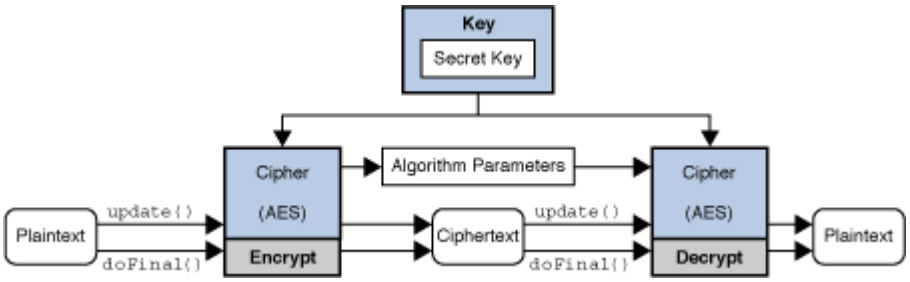

Symmetric vs. Asymmetric Cryptography

There are two major types of encryption: *symmetric* (also known as *secret key*), and *asymmetric* (or *public key cryptography*). In symmetric cryptography, the same secret key to both encrypt and decrypt the data. Keeping the key private is critical to keeping the data confidential. On the other hand, asymmetric cryptography uses a public/private key pair to encrypt data. Data encrypted with one key is decrypted with the other. A user first generates a public/private key pair, and then publishes the public key in a trusted database that anyone can access. A user who wishes to communicate

securely with that user encrypts the data using the retrieved public key. Only the holder of the private key will be able to decrypt. Keeping the private key confidential is critical to this scheme.

Asymmetric algorithms (such as RSA) are generally much slower than symmetric ones. These algorithms are not designed for efficiently protecting large amounts of data. In practice, asymmetric algorithms are used to exchange smaller secret keys which are used to initialize symmetric algorithms.

Stream vs. Block Ciphers

There are two major types of ciphers: *block* and *stream*. Block ciphers process entire blocks at a time, usually many bytes in length. If there is not enough data to make a complete input block, the data must be *padded*: that is, before encryption, dummy bytes must be added to make a multiple of the cipher's block size. These bytes are then stripped off during the decryption phase. The padding can either be done by the application, or by initializing a cipher to use a padding type such as "PKCS5PADDING". In contrast, stream ciphers process incoming data one small unit (typically a byte or even a bit) at a time. This allows for ciphers to process an arbitrary amount of data without padding.

Modes Of Operation

When encrypting using a simple block cipher, two identical blocks of plaintext will always produce an identical block of cipher text. Cryptanalysts trying to break the ciphertext will have an easier job if they note blocks of repeating text. A cipher mode of operation makes the ciphertext less predictable with output block alterations based on block position or the values of other ciphertext blocks. The first block will need an initial value, and this value is called the *initialization vector (IV)*. Since the IV simply alters the data before any encryption, the IV should be random but does not necessarily need to be kept secret. There are a variety of modes, such as CBC (Cipher Block Chaining), CFB (Cipher Feedback Mode), and OFB (Output Feedback Mode). ECB (Electronic Codebook Mode) is a mode in which there is no influence from block position or other ciphertext blocks. Because ECB ciphertexts are the same if they use the same plaintext/key, this mode is not typically suitable for cryptographic applications and should not be used.

Some algorithms such as AES and RSA allow for keys of different lengths, but others are fixed, such as 3DES. Encryption using a longer key generally implies a stronger resistance to message recovery. As usual, there is a trade off between security and time, so choose the key length appropriately.

Most algorithms use binary keys. Most humans do not have the ability to remember long sequences of binary numbers, even when represented in hexadecimal. Character passwords are much easier to recall. Because character passwords are generally chosen from a small number of characters (for example, [a-zA-Z0-9]), protocols such as "Password-Based Encryption" (PBE) have been defined which take character passwords and generate strong binary keys. In order to make the task of getting from password to key very time-consuming for an attacker (via so-called "rainbow table attacks" or "precomputed dictionary attacks" where common dictionary word->value mappings are precomputed), most PBE implementations will mix in a random number, known as a *salt*, to reduce the usefulness of precomputed tables.

Newer cipher modes such as Authenticated Encryption with Associated Data (AEAD) (for example, Galois/Counter Mode (GCM)) encrypt data and authenticate the resulting message simultaneously. Additional Associated Data (AAD) can be used during the calculation of the resulting AEAD tag (MAC), but this AAD data is not output as

ciphertext. (For example, some data might not need to be kept confidential, but should figure into the tag calculation to detect modifications.) The `Cipher.updateAAD()` methods can be used to include AAD in the tag calculations.

Using an AES Cipher with GCM Mode

AES Cipher with GCM is an AEAD Cipher which has different usage patterns than the non-AEAD ciphers. Apart from the regular data, it also takes AAD which is optional for encryption/decryption but AAD must be supplied before data for encryption/decryption. In addition, in order to use GCM securely, callers should not re-use key and IV combinations for encryption. This means that the cipher object should be explicitly re-initialized with a different set of parameters every time for each encryption operation.

Example 2-2 Sample Code for Using an AES Cipher with GCM Mode

```
    SecretKey myKey = ...
    byte[] myAAD = ...
    byte[] plainText = ...
        int myTLen = ...
        byte[] myIv = ...

    GCMParameterSpec myParams = new GCMParameterSpec(myTLen, myIv);
    Cipher c = Cipher.getInstance("AES/GCM/NoPadding");
    c.init(Cipher.ENCRYPT_MODE, myKey, myParams);

    // AAD is optional, if present, it must be supplied before any update/
doFinal calls.
    c.updateAAD(myAAD);  // if AAD is non-null
    byte[] cipherText = new byte[c.getOutputSize(plainText.length)];
    // conclusion of encryption operation
    int actualOutputLen = c.doFinal(plainText, 0, plainText.length,
cipherText);

    // To decrypt, same AAD and GCM parameters must be supplied
    c.init(Cipher.DECRYPT_MODE, myKey, myParams);
    c.updateAAD(myAAD);
    byte[] recoveredText = c.doFinal(cipherText, 0, actualOutputLen);

    // MUST CHANGE IV VALUE if the same key were to be used again for
encryption
         byte[] newIv = ...;
    myParams = new GCMParameterSpec(myTLen, newIv);
```

Creating a Cipher Object

`Cipher` objects are obtained by using one of the `Cipher` `getInstance()` static factory methods. See How Provider Implementations Are Requested and Supplied. Here, the algorithm name is slightly different than with other engine classes, in that it specifies not just an algorithm name, but a "transformation". A transformation is a string that describes the operation (or set of operations) to be performed on the given input to produce some output. A transformation always includes the name of a cryptographic algorithm (e.g., `AES`), and may be followed by a mode and padding scheme.

A transformation is of the form:

- "*algorithm/mode/padding*" or
- "*algorithm*"

For example, the following are valid transformations:

```
"AES/CBC/PKCS5Padding"

"AES"
```

If just a transformation name is specified, the system will determine if there is an implementation of the requested transformation available in the environment, and if there is more than one, returns there is a preferred one.

If both a transformation name and a package provider are specified, the system will determine if there is an implementation of the requested transformation in the package requested, and throw an exception if there is not.

It is recommended to use a transformation that fully specifies the algorithm, mode, and padding. By not doing so, the provider will use a default. For example, the SunJCE and SunPKCS11 providers use ECB as the default mode, and PKCS5Padding as the default padding for many symmetric ciphers.

This means that in the case of the `SunJCE` provider:

```
Cipher c1 = Cipher.getInstance("AES/ECB/PKCS5Padding");
```

and

```
Cipher c1 = Cipher.getInstance("AES");
```

are equivalent statements.

> **Note:**
>
> ECB mode is the easiest block cipher mode to use and is the default in the JDK and JRE. ECB works for single blocks of data when using different keys, but it absolutely should not be used for multiple data blocks. Other cipher modes such as Cipher Block Chaining (CBC) or Galois/Counter Mode (GCM) are more appropriate.

Using modes such as CFB and OFB, block ciphers can encrypt data in units smaller than the cipher's actual block size. When requesting such a mode, you may optionally specify the number of bits to be processed at a time by appending this number to the mode name as shown in the "*AES/CFB8/NoPadding*" and "*AES/OFB32/PKCS5Padding*" transformations. If no such number is specified, a provider-specific default is used. (For example, the `SunJCE` provider uses a default of 128 bits for AES.) Thus, block ciphers can be turned into byte-oriented stream ciphers by using an 8 bit mode such as CFB8 or OFB8.

Java Security Standard Algorithm Names contains a list of standard names that can be used to specify the algorithm name, mode, and padding scheme components of a transformation.

The objects returned by factory methods are uninitialized, and must be initialized before they become usable.

Initializing a Cipher Object

A Cipher object obtained via `getInstance` must be initialized for one of four modes, which are defined as final integer constants in the `Cipher` class. The modes can be referenced by their symbolic names, which are shown below along with a description of the purpose of each mode:

ENCRYPT_MODE
Encryption of data.

DECRYPT_MODE
Decryption of data.

WRAP_MODE
Wrapping a `java.security.Key` into bytes so that the key can be securely transported.

UNWRAP_MODE
Unwrapping of a previously wrapped key into a `java.security.Key` object.

Each of the Cipher initialization methods takes an operational mode parameter (`opmode`), and initializes the Cipher object for that mode. Other parameters include the key (`key`) or certificate containing the key (`certificate`), algorithm parameters (`params`), and a source of randomness (`random`).

To initialize a Cipher object, call one of the following `init` methods:

```
public void init(int opmode, Key key);

public void init(int opmode, Certificate certificate);

public void init(int opmode, Key key, SecureRandom random);

public void init(int opmode, Certificate certificate,
                 SecureRandom random);

public void init(int opmode, Key key,
                 AlgorithmParameterSpec params);

public void init(int opmode, Key key,
                 AlgorithmParameterSpec params, SecureRandom random);

public void init(int opmode, Key key,
                 AlgorithmParameters params);

public void init(int opmode, Key key,
                 AlgorithmParameters params, SecureRandom random);
```

If a Cipher object that requires parameters (e.g., an initialization vector) is initialized for encryption, and no parameters are supplied to the `init` method, the underlying cipher implementation is supposed to supply the required parameters itself, either by generating random parameters or by using a default, provider-specific set of parameters.

However, if a Cipher object that requires parameters is initialized for decryption, and no parameters are supplied to the `init` method, an `InvalidKeyException` or `InvalidAlgorithmParameterException` exception will be raised, depending on the `init` method that has been used.

See Managing Algorithm Parameters.

The same parameters that were used for encryption must be used for decryption.

Note that when a Cipher object is initialized, it loses all previously-acquired state. In other words, initializing a Cipher is equivalent to creating a new instance of that Cipher, and initializing it. For example, if a Cipher is first initialized for decryption with a given key, and then initialized for encryption, it will lose any state acquired while in decryption mode.

Encrypting and Decrypting Data

Data can be encrypted or decrypted in one step (*single-part operation*) or in multiple steps (*multiple-part operation*). A multiple-part operation is useful if you do not know in advance how long the data is going to be, or if the data is too long to be stored in memory all at once.

To encrypt or decrypt data in a single step, call one of the `doFinal` methods:

```
public byte[] doFinal(byte[] input);

public byte[] doFinal(byte[] input, int inputOffset, int inputLen);

public int doFinal(byte[] input, int inputOffset,
                   int inputLen, byte[] output);

public int doFinal(byte[] input, int inputOffset,
                   int inputLen, byte[] output, int outputOffset)
```

To encrypt or decrypt data in multiple steps, call one of the `update` methods:

```
public byte[] update(byte[] input);

public byte[] update(byte[] input, int inputOffset, int inputLen);

public int update(byte[] input, int inputOffset, int inputLen,
                  byte[] output);

public int update(byte[] input, int inputOffset, int inputLen,
                  byte[] output, int outputOffset)
```

A multiple-part operation must be terminated by one of the above `doFinal` methods (if there is still some input data left for the last step), or by one of the following `doFinal` methods (if there is no input data left for the last step):

```
public byte[] doFinal();

public int doFinal(byte[] output, int outputOffset);
```

All the `doFinal` methods take care of any necessary padding (or unpadding), if padding (or unpadding) has been requested as part of the specified transformation.

A call to `doFinal` resets the Cipher object to the state it was in when initialized via a call to `init`. That is, the Cipher object is reset and available to encrypt or decrypt (depending on the operation mode that was specified in the call to `init`) more data.

Wrapping and Unwrapping Keys

Wrapping a key enables secure transfer of the key from one place to another.

The `wrap/unwrap` API makes it more convenient to write code since it works with key objects directly. These methods also enable the possibility of secure transfer of hardware-based keys.

To **wrap** a Key, first initialize the Cipher object for WRAP_MODE, and then call the following:

```
public final byte[] wrap(Key key);
```

If you are supplying the wrapped key bytes (the result of calling `wrap`) to someone else who will unwrap them, be sure to also send additional information the recipient will need in order to do the `unwrap`:

- The name of the key algorithm.
- The type of the wrapped key (one of `Cipher.SECRET_KEY`, `Cipher.PRIVATE_KEY`, or `Cipher.PUBLIC_KEY`).

The key algorithm name can be determined by calling the `getAlgorithm` method from the Key interface:

```
public String getAlgorithm();
```

To **unwrap** the bytes returned by a previous call to `wrap`, first initialize a Cipher object for UNWRAP_MODE, then call the following:

```
public final Key unwrap(byte[] wrappedKey,
                        String wrappedKeyAlgorithm,
                        int wrappedKeyType));
```

Here, `wrappedKey` is the bytes returned from the previous call to wrap, `wrappedKeyAlgorithm` is the algorithm associated with the wrapped key, and

`wrappedKeyType` is the type of the wrapped key. This must be one of `Cipher.SECRET_KEY`, `Cipher.PRIVATE_KEY`, or `Cipher.PUBLIC_KEY`.

Managing Algorithm Parameters

The parameters being used by the underlying `Cipher` implementation, which were either explicitly passed to the `init` method by the application or generated by the underlying implementation itself, can be retrieved from the `Cipher` object by calling its `getParameters` method, which returns the parameters as a `java.security.AlgorithmParameters` object (or `null` if no parameters are being used). If the parameter is an initialization vector (IV), it can also be retrieved by calling the `getIV` method.

In the following example, a `Cipher` object implementing password-based encryption (PBE) is initialized with just a key and no parameters. However, the selected algorithm for password-based encryption requires two parameters - a *salt* and an *iteration count*. Those will be generated by the underlying algorithm implementation itself. The application can retrieve the generated parameters from the `Cipher` object, see Example 2-3.

The same parameters that were used for encryption must be used for decryption. They can be instantiated from their encoding and used to initialize the corresponding Cipher object for decryption, see Example 2-4.

If you did not specify any parameters when you initialized a Cipher object, and you are not sure whether or not the underlying implementation uses any parameters, you can find out by simply calling the `getParameters` method of your Cipher object and checking the value returned. A return value of `null` indicates that no parameters were used.

The following cipher algorithms implemented by the `SunJCE` provider use parameters:

- AES, DES-EDE, and Blowfish, when used in feedback (i.e., CBC, CFB, OFB, or PCBC) mode, use an initialization vector (IV). The `javax.crypto.spec.IvParameterSpec` class can be used to initialize a `Cipher` object with a given IV. In addition, CTR and GCM modes require an IV.
- PBE `Cipher` algorithms use a set of parameters, comprising a salt and an iteration count. The `javax.crypto.spec.PBEParameterSpec` class can be used to initialize a `Cipher` object implementing a PBE algorithm (for example: PBEWithHmacSHA256AndAES_256) with a given salt and iteration count.

Note that you do not have to worry about storing or transferring any algorithm parameters for use by the decryption operation if you use the The SealedObject Class class. This class attaches the parameters used for sealing (encryption) to the encrypted object contents, and uses the same parameters for unsealing (decryption).

Example 2-3 Sample Code for Retrieving Parameters from the Cipher Object

The application can retrieve the generated parameters for encryption from the Cipher object as follows:

```
import javax.crypto.*;
import java.security.AlgorithmParameters;

// get cipher object for password-based encryption
Cipher c = Cipher.getInstance("PBEWithHmacSHA256AndAES_256");
```

```
    // initialize cipher for encryption, without supplying
    // any parameters. Here, "myKey" is assumed to refer
    // to an already-generated key.
    c.init(Cipher.ENCRYPT_MODE, myKey);

    // encrypt some data and store away ciphertext
    // for later decryption
    byte[] cipherText = c.doFinal("This is just an example".getBytes());

    // retrieve parameters generated by underlying cipher
    // implementation
    AlgorithmParameters algParams = c.getParameters();

    // get parameter encoding and store it away
    byte[] encodedAlgParams = algParams.getEncoded();
```

Example 2-4 Sample Code for Initializing the Cipher Object for Decryption

The same parameters that were used for encryption must be used for decryption. They can be instantiated from their encoding and used to initialize the corresponding Cipher object for decryption as follows:

```
    import javax.crypto.*;
    import java.security.AlgorithmParameters;

    // get parameter object for password-based encryption
    AlgorithmParameters algParams;
    algParams =
AlgorithmParameters.getInstance("PBEWithHmacSHA256AndAES_256");

    // initialize with parameter encoding from above
    algParams.init(encodedAlgParams);

    // get cipher object for password-based encryption
    Cipher c = Cipher.getInstance("PBEWithHmacSHA256AndAES_256");

    // initialize cipher for decryption, using one of the
    // init() methods that takes an AlgorithmParameters
    // object, and pass it the algParams object from above
    c.init(Cipher.DECRYPT_MODE, myKey, algParams);
```

Cipher Output Considerations

Some of the `update` and `doFinal` methods of Cipher allow the caller to specify the output buffer into which to encrypt or decrypt the data. In these cases, it is important to pass a buffer that is large enough to hold the result of the encryption or decryption operation.

The following method in Cipher can be used to determine how big the output buffer should be:

```
public int getOutputSize(int inputLen)
```

Other Cipher-based Classes

There are some helper classes which internally use `Ciphers` to provide easy access to common cipher uses.

Topics

The Cipher Stream Classes

The SealedObject Class

The Cipher Stream Classes

The `CipherInputStream` and `CipherOutputStream` classes are Cipher stream classes.

The CipherInputStream Class

This class is a `FilterInputStream` that encrypts or decrypts the data passing through it. It is composed of an `InputStream`. `CipherInputStream` represents a secure input stream into which a `Cipher` object has been interposed. The `read` methods of `CipherInputStream` return data that are read from the underlying `InputStream` but have additionally been processed by the embedded `Cipher` object. The `Cipher` object must be fully initialized before being used by a CipherInputStream.

For example, if the embedded `Cipher` has been initialized for decryption, the `CipherInputStream` will attempt to decrypt the data it reads from the underlying `InputStream` before returning them to the application.

This class adheres strictly to the semantics, especially the failure semantics, of its ancestor classes `java.io.FilterInputStream` and `java.io.InputStream`. This class has exactly those methods specified in its ancestor classes, and overrides them all, so that the data are additionally processed by the embedded cipher. Moreover, this class catches all exceptions that are not thrown by its ancestor classes. In particular, the `skip(long)` method skips only data that has been processed by the `Cipher`.

It is crucial for a programmer using this class not to use methods that are not defined or overridden in this class (such as a new method or constructor that is later added to one of the super classes), because the design and implementation of those methods are unlikely to have considered security impact with regard to `CipherInputStream`. See Example 2-5 for its usage, suppose `cipher1` has been initialized for encryption. The program reads and encrypts the content from the file `/tmp/a.txt` and then stores the result (the encrypted bytes) in `/tmp/b.txt`.

Example 2-6 demonstrates how to easily connect several instances of `CipherInputStream` and `FileInputStream`. In this example, assume that `cipher1` and `cipher2` have been initialized for encryption and decryption (with corresponding keys), respectively. The program copies the content from file `/tmp/a.txt` to `/tmp/b.txt`, except that the content is first encrypted and then decrypted back when it is read from `/tmp/a.txt`. Of course since this program simply encrypts text and decrypts it

back right away, it's actually not very useful except as a simple way of illustrating chaining of `CipherInputStreams`.

Note that the read methods of the `CipherInputStream` will block until data is returned from the underlying cipher. If a block cipher is used, a full block of cipher text will have to be obtained from the underlying `InputStream`.

Example 2-5 Sample Code for Using CipherInputStream and FileInputStream

The code below demonstrates how to use a `CipherInputStream` containing that cipher and a `FileInputStream` in order to encrypt input stream data:

```
try (FileInputStream fis = new FileInputStream("/tmp/a.txt");
CipherInputStream cis = new CipherInputStream(fis, cipher1);
FileOutputStream fos = new FileOutputStream("/tmp/b.txt")) {
    byte[] b = new byte[8];
    int i = cis.read(b);
    while (i != -1) {
        fos.write(b, 0, i);
        i = cis.read(b);
    }
}
```

Example 2-6 Sample Code for Connecting CipherInputStream and FileInputStream

The following example demonstrates how to easily connect several instances of `CipherInputStream` and `FileInputStream`. In this example, assume that `cipher1` and `cipher2` have been initialized for encryption and decryption (with corresponding keys), respectively:

```
try (FileInputStream fis = new FileInputStream("/tmp/a.txt");
        CipherInputStream cis1 = new CipherInputStream(fis, cipher1);
        CipherInputStream cis2 = new CipherInputStream(cis1, cipher2);
        FileOutputStream fos = new FileOutputStream("/tmp/b.txt")) {
    byte[] b = new byte[8];
    int i = cis2.read(b);
    while (i != -1) {
        fos.write(b, 0, i);
        i = cis2.read(b);
    }
}
```

The CipherOutputStream Class

This class is a `FilterOutputStream` that encrypts or decrypts the data passing through it. It is composed of an `OutputStream`, or one of its subclasses, and a `Cipher`. `CipherOutputStream` represents a secure output stream into which a `Cipher` object has been interposed. The `write` methods of `CipherOutputStream` first process the data with the embedded `Cipher` object before writing them out to the underlying `OutputStream`. The `Cipher` object must be fully initialized before being used by a `CipherOutputStream`.

For example, if the embedded `Cipher` has been initialized for encryption, the `CipherOutputStream` will encrypt its data, before writing them out to the underlying output stream.

This class adheres strictly to the semantics, especially the failure semantics, of its ancestor classes `java.io.OutputStream` and `java.io.FilterOutputStream`. This class has exactly those methods specified in its ancestor classes, and overrides them all, so that all data are additionally processed by the embedded cipher. Moreover, this class catches all exceptions that are not thrown by its ancestor classes.

It is crucial for a programmer using this class not to use methods that are not defined or overridden in this class (such as a new method or constructor that is later added to one of the super classes), because the design and implementation of those methods are unlikely to have considered security impact with regard to `CipherOutputStream`.

See Example 2-7 , for its usage, suppose `cipher1` has been initialized for encryption. The program reads the content from the file `/tmp/a.txt`, then encrypts and stores the result (the encrypted bytes) in `/tmp/b.txt`.

Example 2-7 demonstrates how to easily connect several instances of `CipherOutputStream` and `FileOutputStream`. In this example, assume that `cipher1` and `cipher2` have been initialized for decryption and encryption (with corresponding keys), respectively. The program copies the content from file `/tmp/a.txt` to `/tmp/b.txt`, except that the content is first encrypted and then decrypted back before it is written to `/tmp/b.txt`.

One thing to keep in mind when using *block* cipher algorithms is that a full block of plaintext data must be given to the `CipherOutputStream` before the data will be encrypted and sent to the underlying output stream.

There is one other important difference between the `flush` and `close` methods of this class, which becomes even more relevant if the encapsulated `Cipher` object implements a block cipher algorithm with padding turned on:

- `flush` flushes the underlying `OutputStream` by forcing any buffered output bytes that have already been processed by the encapsulated Cipher object to be written out. Any bytes buffered by the encapsulated `Cipher` object and waiting to be processed by it will **not** be written out.
- `close` closes the underlying `OutputStream` and releases any system resources associated with it. It invokes the `doFinal` method of the encapsulated `Cipher` object, causing any bytes buffered by it to be processed and written out to the underlying stream by calling its `flush` method.

Example 2-7 Sample Code for Using CipherOutputStream and FileOutputStream

```
CipherOutputStreamFileOutputStream

try (FileInputStream fis = new FileInputStream("/tmp/a.txt");
        FileOutputStream fos = new FileOutputStream("/tmp/b.txt");
        CipherOutputStream cos = new CipherOutputStream(fos, cipher1)) {
    byte[] b = new byte[8];
    int i = fis.read(b);
    while (i != -1) {
        cos.write(b, 0, i);
        i = fis.read(b);
    }
```

```
    cos.flush();
}
```

Example 2-8 Sample Code for Connecting CipherOutputStream and FileOutputStream

```
CipherOutputStreamFileOutputStreamcipher1cipher2

try (FileInputStream fis = new FileInputStream("/tmp/a.txt");
        FileOutputStream fos = new FileOutputStream("/tmp/b.txt");
        CipherOutputStream cos1 = new CipherOutputStream(fos, cipher1);
        CipherOutputStream cos2 = new CipherOutputStream(cos1, cipher2)) {
    byte[] b = new byte[8];
    int i = fis.read(b);
    while (i != -1) {
        cos2.write(b, 0, i);
        i = fis.read(b);
    }
    cos2.flush();
}
```

The SealedObject Class

This class enables a programmer to create an object and protect its confidentiality with a cryptographic algorithm.

Given any object that implements the `java.io.Serializable` interface, one can create a `SealedObject` that encapsulates the original object, in serialized format (i.e., a "deep copy"), and seals (encrypts) its serialized contents, using a cryptographic algorithm such as AES, to protect its confidentiality. The encrypted content can later be decrypted (with the corresponding algorithm using the correct decryption key) and de-serialized, yielding the original object.

A typical usage is illustrated in the following code segment: In order to seal an object, you create a `SealedObject` from the object to be sealed and a fully initialized `Cipher` object that will encrypt the serialized object contents. In this example, the String "This is a secret" is sealed using the AES algorithm. Note that any algorithm parameters that may be used in the sealing operation are stored inside of `SealedObject`:

```
// create Cipher object
// NOTE: sKey is assumed to refer to an already-generated
// secret AES key.
Cipher c = Cipher.getInstance("AES");
c.init(Cipher.ENCRYPT_MODE, sKey);

// do the sealing
SealedObject so = new SealedObject("This is a secret", c);
```

The original object that was sealed can be recovered in two different ways:

- by using a `Cipher` object that has been initialized with the exact same algorithm, key, padding scheme, etc., that were used to seal the object:

```
c.init(Cipher.DECRYPT_MODE, sKey);
try {
    String s = (String)so.getObject(c);
} catch (Exception e) {
    // do something
};
```

  This approach has the advantage that the party who unseals the sealed object does not require knowledge of the decryption key. For example, after one party has initialized the cipher object with the required decryption key, it could hand over the cipher object to another party who then unseals the sealed object.

- by using the appropriate decryption key (since AES is a symmetric encryption algorithm, we use the same key for sealing and unsealing):

```
try {
    String s = (String)so.getObject(sKey);
} catch (Exception e) {
    // do something
};
```

  In this approach, the `getObject` method creates a cipher object for the appropriate decryption algorithm and initializes it with the given decryption key and the algorithm parameters (if any) that were stored in the sealed object. This approach has the advantage that the party who unseals the object does not need to keep track of the parameters (e.g., the IV) that were used to seal the object.

The Mac Class

Similar to a `MessageDigest`, a Message Authentication Code (MAC) provides a way to check the integrity of information transmitted over or stored in an unreliable medium, but includes a secret key in the calculation.

Only someone with the proper key will be able to verify the received message. Typically, message authentication codes are used between two parties that share a secret key in order to validate information transmitted between these parties.

Figure 2-9 The Mac Class

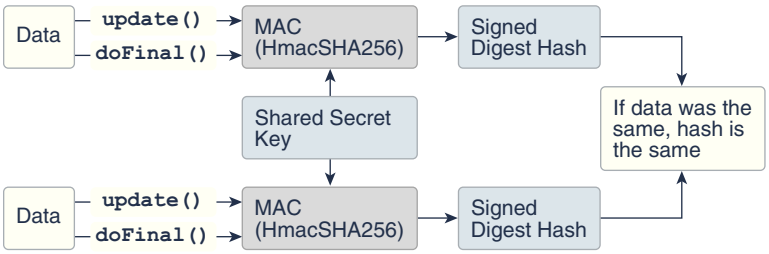

A MAC mechanism that is based on cryptographic hash functions is referred to as HMAC. HMAC can be used with any cryptographic hash function, e.g., SHA-256, in combination with a secret shared key.

The `Mac` class provides the functionality of a Message Authentication Code (MAC). See HMAC-SHA256 Example.

Creating a Mac Object

`Mac` objects are obtained by using one of the `Mac` `getInstance()` static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a Mac Object

A Mac object is always initialized with a (secret) key and may optionally be initialized with a set of parameters, depending on the underlying MAC algorithm.

To initialize a Mac object, call one of its `init` methods:

```
public void init(Key key);

public void init(Key key, AlgorithmParameterSpec params);
```

You can initialize your Mac object with any (secret-)key object that implements the `javax.crypto.SecretKey` interface. This could be an object returned by `javax.crypto.KeyGenerator.generateKey()`, or one that is the result of a key agreement protocol, as returned by `javax.crypto.KeyAgreement.generateSecret()`, or an instance of `javax.crypto.spec.SecretKeySpec`.

With some MAC algorithms, the (secret-)key algorithm associated with the (secret-)key object used to initialize the Mac object does not matter (this is the case with the HMAC-MD5 and HMAC-SHA1 implementations of the `SunJCE` provider). With others, however, the (secret-)key algorithm does matter, and an `InvalidKeyException` is thrown if a (secret-)key object with an inappropriate (secret-)key algorithm is used.

Computing a MAC

A MAC can be computed in one step (*single-part operation*) or in multiple steps (*multiple-part operation*). A multiple-part operation is useful if you do not know in advance how long the data is going to be, or if the data is too long to be stored in memory all at once.

To compute the MAC of some data in a single step, call the following `doFinal` method:

```
public byte[] doFinal(byte[] input);
```

To compute the MAC of some data in multiple steps, call one of the `update` methods:

```
public void update(byte input);

public void update(byte[] input);
```

```
public void update(byte[] input, int inputOffset, int inputLen);
```

A multiple-part operation must be terminated by the above `doFinal` method (if there is still some input data left for the last step), or by one of the following `doFinal` methods (if there is no input data left for the last step):

```
public byte[] doFinal();

public void doFinal(byte[] output, int outOffset);
```

Key Interfaces

The `java.security.Key` interface is the top-level interface for all opaque keys. It defines the functionality shared by all opaque key objects.

To this point, we have focused the high-level uses of the JCA without getting lost in the details of what keys are and how they are generated/represented. It is now time to turn our attention to keys.

An *opaque* key representation is one in which you have no direct access to the key material that constitutes a key. In other words: "opaque" gives you limited access to the key--just the three methods defined by the `Key` interface (see below): `getAlgorithm`, `getFormat`, and `getEncoded`.

This is in contrast to a *transparent* representation, in which you can access each key material value individually, through one of the `get` methods defined in the corresponding `KeySpec` interface (see The KeySpec Interface).

All opaque keys have three characteristics:

An Algorithm
The key algorithm for that key. The key algorithm is usually an encryption or asymmetric operation algorithm (such as `AES`, `DSA` or `RSA`), which will work with those algorithms and with related algorithms (such as `SHA256withRSA`). The name of the algorithm of a key is obtained using this method:

```
String getAlgorithm()
```

An Encoded Form
The external encoded form for the key used when a standard representation of the key is needed outside the Java Virtual Machine, as when transmitting the key to some other party. The key is encoded according to a standard format (such as X.509 or PKCS8), and is returned using the method:

```
byte[] getEncoded()
```

A Format
The name of the format of the encoded key. It is returned by the method:

```
String getFormat()
```

Keys are generally obtained through key generators such as the `KeyGenerator` class and the `KeyPairGenerator` class, certificates, key specifications (see the The KeySpec Interface) using a `KeyFactory`, or a `Keystore` implementation accessing a keystore database used to manage keys. It is possible to parse encoded keys, in an algorithm-dependent manner, using a `KeyFactory`.

It is also possible to parse certificates, using a `CertificateFactory`.

Here is a list of interfaces which extend the `Key` interface in the `java.security.interfaces` and `javax.crypto.interfaces` packages:

- `SecretKey`
  - `PBEKey`
- `PrivateKey`
  - `DHPrivateKey`
  - `DSAPrivateKey`
  - `ECPrivateKey`
  - `RSAMultiPrimePrivateCrtKey`
  - `RSAPrivateCrtKey`
  - `RSAPrivateKey`
- `PublicKey`
  - `DHPublicKey`
  - `DSAPublicKey`
  - `ECPublicKey`
  - `RSAPublicKey`

The PublicKey and PrivateKey Interfaces

The `PublicKey` and `PrivateKey` interfaces (which both extend the `Key` interface) are methodless interfaces, used for type-safety and type-identification.

The KeyPair Class

The `KeyPair` class is a simple holder for a key pair (a public key and a private key).

It has two public methods, one for returning the private key, and the other for returning the public key:

```
PrivateKey getPrivate()
PublicKey getPublic()
```

Key Specification Interfaces and Classes

`Key` objects and key specifications (`KeySpecs`) are two different representations of key data. `Ciphers` use `Key` objects to initialize their encryption algorithms, but keys may need to be converted into a more portable format for transmission or storage.

A *transparent* representation of keys means that you can access each key material value individually, through one of the `get` methods defined in the corresponding specification class. For example, `DSAPrivateKeySpec` defines `getX`, `getP`, `getQ`, and `getG` methods, to access the private key `x`, and the DSA algorithm parameters used to calculate the key: the prime `p`, the sub-prime `q`, and the base `g`. If the key is stored on a hardware device, its specification may contain information that helps identify the key on the device.

This representation is contrasted with an *opaque* representation, as defined by the Key Interfaces interface, in which you have no direct access to the key material fields. In other words, an "opaque" representation gives you limited access to the key--just the three methods defined by the `Key` interface: `getAlgorithm`, `getFormat`, and `getEncoded`.

A key may be specified in an algorithm-specific way, or in an algorithm-independent encoding format (such as ASN.1). For example, a DSA private key may be specified by its components `x`, `p`, `q`, and `g` (see `DSAPrivateKeySpec`), or it may be specified using its DER encoding (see `PKCS8EncodedKeySpec`).

The The KeyFactory Class and The SecretKeyFactory Class classes can be used to convert between opaque and transparent key representations (that is, between `Keys` and `KeySpecs`, assuming that the operation is possible. (For example, private keys on smart cards might not be able leave the card. Such `Keys` are not convertible.)

In the following sections, we discuss the key specification interfaces and classes in the `java.security.spec` package.

The KeySpec Interface

This interface contains no methods or constants. Its only purpose is to group and provide type safety for all key specifications. All key specifications must implement this interface.

The KeySpec Subinterfaces

Like the `Key` interface, there are a similar set of `KeySpec` interfaces.

- `SecretKeySpec`
- `EncodedKeySpec`
  - `PKCS8EncodedKeySpec`
  - `X509EncodedKeySpec`
- `DESKeySpec`
- `DESedeKeySpec`
- `PBEKeySpec`
- `DHPrivateKeySpec`

- DSAPrivateKeySpec
- ECPrivateKeySpec
- RSAPrivateKeySpec
  - RSAMultiPrimePrivateCrtKeySpec
  - RSAPrivateCrtKeySpec
- DHPublicKeySpec
- DSAPublicKeySpec
- ECPublicKeySpec
- RSAPublicKeySpec

The EncodedKeySpec Class

This abstract class (which implements the The KeySpec Interface interface) represents a public or private key in encoded format. Its `getEncoded` method returns the encoded key:

```
abstract byte[] getEncoded();
```

and its `getFormat` method returns the name of the encoding format:

```
abstract String getFormat();
```

See the next sections for the concrete implementations `PKCS8EncodedKeySpec` and `X509EncodedKeySpec`.

The PKCS8EncodedKeySpec Class

This class, which is a subclass of `EncodedKeySpec`, represents the DER encoding of a private key, according to the format specified in the PKCS8 standard.

Its `getEncoded` method returns the key bytes, encoded according to the PKCS8 standard. Its `getFormat` method returns the string "PKCS#8".

The X509EncodedKeySpec Class

This class, which is a subclass of `EncodedKeySpec`, represents the DER encoding of a public key, according to the format specified in the X.509 standard.

Its `getEncoded` method returns the key bytes, encoded according to the X.509 standard. Its `getFormat` method returns the string "X.509".

Generators and Factories

Newcomers to Java and the JCA APIs in particular sometimes do not grasp the distinction between generators and factories.

Figure 2-10 Generators and Factories

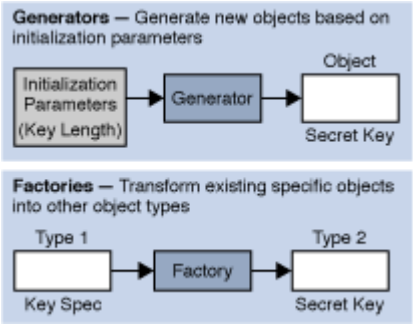

Generators are used to **generate brand new objects**. Generators can be initialized in either an algorithm-dependent or algorithm-independent way. For example, to create a Diffie-Hellman (DH) keypair, an application could specify the necessary P and G values, or the generator could simply be initialized with the appropriate key length, and the generator will select appropriate P and G values. In both cases, the generator will produce brand new keys based on the parameters.

On the other hand, factories are used to **convert data from one existing object type to another**. For example, an application might have available the components of a DH private key and can package them as a The KeySpec Interface, but needs to convert them into a PrivateKey object that can be used by a `KeyAgreement` object, or vice-versa. Or they might have the byte array of a certificate, but need to use a `CertificateFactory` to convert it into a `X509Certificate` object. Applications use factory objects to do the conversion.

The KeyFactory Class

The `KeyFactory` class is an Engine Classes and Algorithms designed to perform conversions between opaque cryptographic Key Interfaces and Key Specification Interfaces and Classes (transparent representations of the underlying key material).

Figure 2-11 KeyFactory Class

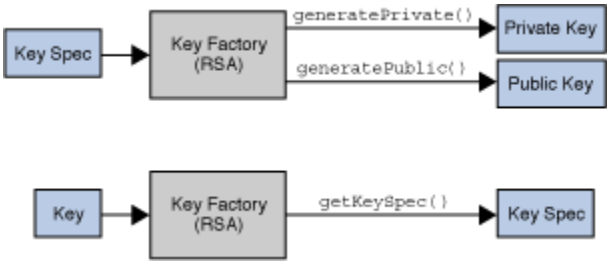

Key factories are bi-directional. They allow you to build an opaque key object from a given key specification (key material), or to retrieve the underlying key material of a key object in a suitable format.

Multiple compatible key specifications can exist for the same key. For example, a DSA public key may be specified by its components `y`, `p`, `q`, and `g` (see

`java.security.spec.DSAPublicKeySpec`), or it may be specified using its DER encoding according to the X.509 standard (see The X509EncodedKeySpec Class).

A key factory can be used to translate between compatible key specifications. Key parsing can be achieved through translation between compatible key specifications, e.g., when you translate from `X509EncodedKeySpec` to `DSAPublicKeySpec`, you basically parse the encoded key into its components. For an example, see the end of the Generating/Verifying Signatures Using Key Specifications and KeyFactory section.

Creating a KeyFactory Object

`KeyFactory` objects are obtained by using one of the `KeyFactory`getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Converting Between a Key Specification and a Key Object

If you have a key specification for a public key, you can obtain an opaque `PublicKey` object from the specification by using the `generatePublic` method:

```
PublicKey generatePublic(KeySpec keySpec)
```

Similarly, if you have a key specification for a private key, you can obtain an opaque `PrivateKey` object from the specification by using the `generatePrivate` method:

```
PrivateKey generatePrivate(KeySpec keySpec)
```

Converting Between a Key Object and a Key Specification

If you have a `Key` object, you can get a corresponding key specification object by calling the `getKeySpec` method:

```
KeySpec getKeySpec(Key key, Class keySpec)
```

`keySpec` identifies the specification class in which the key material should be returned. It could, for example, be `DSAPublicKeySpec.class` , to indicate that the key material should be returned in an instance of the `DSAPublicKeySpec` class. See Generating/Verifying Signatures Using Key Specifications and KeyFactory.

The SecretKeyFactory Class

The `SecretKeyFactory` class represents a factory for secret keys. Unlike the `KeyFactory` class (see The KeyFactory Class), a `javax.crypto.SecretKeyFactory` object operates only on secret (symmetric) keys, whereas a `java.security.KeyFactory` object processes the public and private key components of a key pair.

Figure 2-12 SecretKeyFactory Class

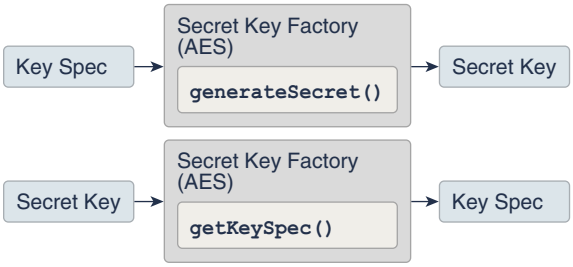

Key factories are used to convert Key Interfaces (opaque cryptographic keys of type `java.security.Key`) into Key Specification Interfaces and Classes (transparent representations of the underlying key material in a suitable format), and vice versa.

Objects of type `java.security.Key`, of which `java.security.PublicKey`, `java.security.PrivateKey`, and `javax.crypto.SecretKey` are subclasses, are opaque key objects, because you cannot tell how they are implemented. The underlying implementation is provider-dependent, and may be software or hardware based. Key factories allow providers to supply their own implementations of cryptographic keys.

For example, if you have a key specification for a Diffie-Hellman public key, consisting of the public value `y`, the prime modulus `p`, and the base `g`, and you feed the same specification to Diffie-Hellman key factories from different providers, the resulting `PublicKey` objects will most likely have different underlying implementations.

A provider should document the key specifications supported by its secret key factory. For example, the `SecretKeyFactory` for DES keys supplied by the `SunJCE` provider supports `DESKeySpec` as a transparent representation of DES keys, the `SecretKeyFactory` for DES-EDE keys supports `DESedeKeySpec` as a transparent representation of DES-EDE keys, and the `SecretKeyFactory` for PBE supports `PBEKeySpec` as a transparent representation of the underlying password.

The following is an example of how to use a `SecretKeyFactory` to convert secret key data into a `SecretKey` object, which can be used for a subsequent `Cipher` operation:

```
// Note the following bytes are not realistic secret key data
// bytes but are simply supplied as an illustration of using data
// bytes (key material) you already have to build a DESedeKeySpec.

byte[] desEdeKeyData = getKeyData();
DESedeKeySpec desEdeKeySpec = new DESedeKeySpec(desEdeKeyData);
SecretKeyFactory keyFactory = SecretKeyFactory.getInstance("DESede");
SecretKey secretKey = keyFactory.generateSecret(desEdeKeySpec);
```

In this case, the underlying implementation of `SecretKey` is based on the provider of `KeyFactory`.

An alternative, provider-independent way of creating a functionally equivalent `SecretKey` object from the same key material is to use the

`javax.crypto.spec.SecretKeySpec` class, which implements the `javax.crypto.SecretKey` interface:

```
byte[] aesKeyData = getKeyData();
SecretKeySpec secretKey = new SecretKeySpec(aesKeyData, "AES");
```

Creating a SecretKeyFactory Object

`SecretKeyFactory` objects are obtained by using one of the `SecretKeyFactory` `getInstance()` static factory methods. See How Provider Implementations Are Requested and Supplied.

Converting Between a Key Specification and a Secret Key Object

If you have a key specification for a secret key, you can obtain an opaque `SecretKey` object from the specification by using the `generateSecret` method:

```
SecretKey generateSecret(KeySpec keySpec)
```

Converting Between a Secret Key Object and a Key Specification

If you have a `SecretKey` object, you can get a corresponding key specification object by calling the `getKeySpec` method:

```
KeySpec getKeySpec(Key key, Class keySpec)
```

`keySpec` identifies the specification class in which the key material should be returned. It could, for example, be `DESKeySpec.class`, to indicate that the key material should be returned in an instance of the `DESKeySpec` class.

The KeyPairGenerator Class

The `KeyPairGenerator` class is an engine class (see Engine Classes and Algorithms) used to generate pairs of public and private keys.

Figure 2-13 KeyPairGenerator Class

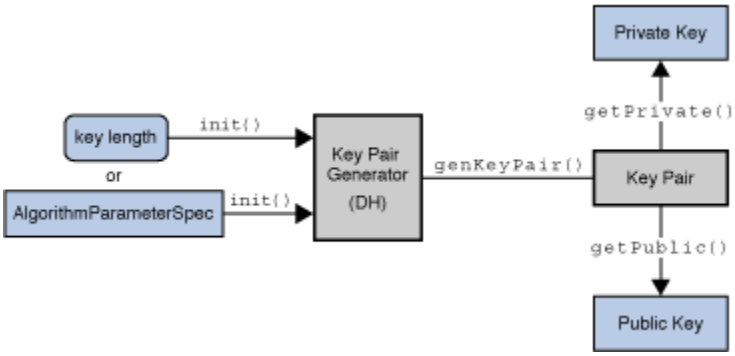

There are two ways to generate a key pair: in an algorithm-independent manner, and in an algorithm-specific manner. The only difference between the two is the initialization of the object.

See Generating a Pair of Keys for examples of calls to the methods documented below.

Creating a KeyPairGenerator

All key pair generation starts with a `KeyPairGenerator`. `KeyPairGenerator` objects are obtained by using one of the `KeyPairGenerator` `getInstance()` static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a KeyPairGenerator

A key pair generator for a particular algorithm creates a public/private key pair that can be used with this algorithm. It also associates algorithm-specific parameters with each of the generated keys.

A key pair generator needs to be initialized before it can generate keys. In most cases, algorithm-independent initialization is sufficient. But in other cases, algorithm-specific initialization can be used.

Algorithm-Independent Initialization

All key pair generators share the concepts of a keysize and a source of randomness. The keysize is interpreted differently for different algorithms. For example, in the case of the DSA algorithm, the keysize corresponds to the length of the modulus. (See Java Security Standard Algorithm Names for information about the keysizes for specific algorithms.)

An `initialize` method takes two universally shared types of arguments:

```
void initialize(int keysize, SecureRandom random)
```

Another `initialize` method takes only a `keysize` argument; it uses a system-provided source of randomness:

```
void initialize(int keysize)
```

Since no other parameters are specified when you call the above algorithm-independent `initialize` methods, it is up to the provider what to do about the algorithm-specific parameters (if any) to be associated with each of the keys.

If the algorithm is a "DSA" algorithm, and the modulus size (keysize) is 512, 768, 1024, 2048, or 3072, then the `SUN` provider uses a set of precomputed values for the `p`, `q`, and `g` parameters. If the modulus size is not one of the above values, the `SUN` provider creates a new set of parameters. Other providers might have precomputed parameter sets for more than just the three modulus sizes mentioned above. Still others might not have a list of precomputed parameters at all and instead always create new parameter sets.

Algorithm-Specific Initialization

For situations where a set of algorithm-specific parameters already exists (such as "community parameters" in DSA), there are two `initialize` methods that have an The

AlgorithmParameterSpec Interface argument. One also has a `SecureRandom` argument, while the source of randomness is system-provided for the other:

```
void initialize(AlgorithmParameterSpec params,
                SecureRandom random)

void initialize(AlgorithmParameterSpec params)
```

See Generating a Pair of Keys.

Generating a Key Pair

The procedure for generating a key pair is always the same, regardless of initialization (and of the algorithm). You always call the following method from `KeyPairGenerator`:

```
KeyPair generateKeyPair()
```

Multiple calls to `generateKeyPair` will yield different key pairs.

The KeyGenerator Class

A key generator is used to generate secret keys for symmetric algorithms.

Figure 2-14 The KeyGenerator Class

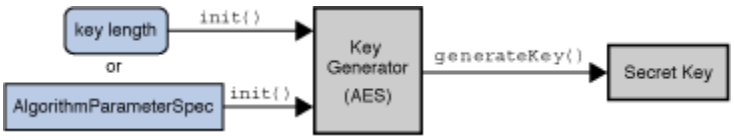

Creating a `KeyGenerator`

`KeyGenerator` objects are obtained by using one of the `KeyGenerator` `getInstance()` static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a KeyGenerator Object

A key generator for a particular symmetric-key algorithm creates a symmetric key that can be used with that algorithm. It also associates algorithm-specific parameters (if any) with the generated key.

There are two ways to generate a key: in an algorithm-independent manner, and in an algorithm-specific manner. The only difference between the two is the initialization of the object:

- **Algorithm-Independent Initialization**
  All key generators share the concepts of a *keysize* and a *source of randomness*. There is an `init` method that takes these two universally shared types of arguments. There is also one that takes just a `keysize` argument, and uses a

system-provided source of randomness, and one that takes just a source of randomness:

```
public void init(SecureRandom random);

public void init(int keysize);

public void init(int keysize, SecureRandom random);
```

Since no other parameters are specified when you call the above algorithm-independent `init` methods, it is up to the provider what to do about the algorithm-specific parameters (if any) to be associated with the generated key.

- **Algorithm-Specific Initialization**
  For situations where a set of algorithm-specific parameters already exists, there are two `init` methods that have an `AlgorithmParameterSpec` argument. One also has a `SecureRandom` argument, while the source of randomness is system-provided for the other:

```
public void init(AlgorithmParameterSpec params);

public void init(AlgorithmParameterSpec params, SecureRandom
random);
```

In case the client does not explicitly initialize the KeyGenerator (via a call to an `init` method), each provider must supply (and document) a default initialization.

Creating a Key

The following method generates a secret key:

```
public SecretKey generateKey();
```

The KeyAgreement Class

Key agreement is a protocol by which 2 or more parties can establish the same cryptographic keys, without having to exchange any secret information.

Figure 2-15 The KeyAgreement Class

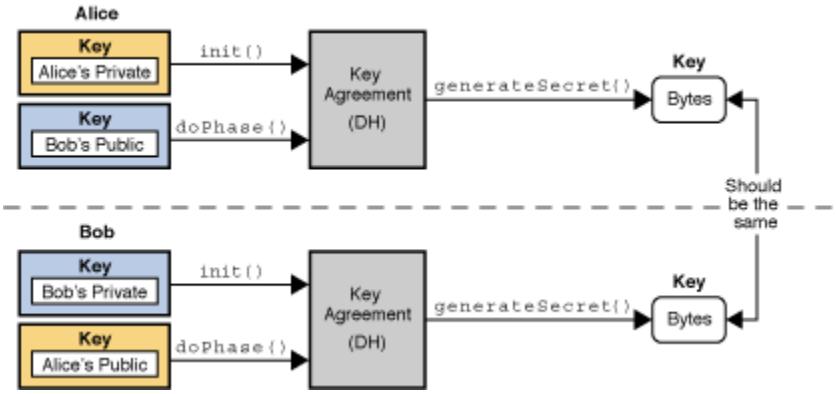

Each party initializes their key agreement object with their private key, and then enters the public keys for each party that will participate in the communication. In most cases, there are just two parties, but algorithms such as Diffie-Hellman allow for multiple parties (3 or more) to participate. When all the public keys have been entered, each `KeyAgreement` object will generate (agree upon) the same key.

The KeyAgreement class provides the functionality of a key agreement protocol. The keys involved in establishing a shared secret are created by one of the key generators (`KeyPairGenerator` or `KeyGenerator`), a `KeyFactory`, or as a result from an intermediate phase of the key agreement protocol.

Creating a KeyAgreement Object

Each party involved in the key agreement has to create a KeyAgreement object. `KeyAgreement` objects are obtained by using one of the `KeyAgreement` `getInstance()` static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a KeyAgreement Object

You initialize a KeyAgreement object with your private information. In the case of Diffie-Hellman, you initialize it with your Diffie-Hellman private key. Additional initialization information may contain a source of randomness and/or a set of algorithm parameters. Note that if the requested key agreement algorithm requires the specification of algorithm parameters, and only a key, but no parameters are provided to initialize the KeyAgreement object, the key must contain the required algorithm parameters. (For example, the Diffie-Hellman algorithm uses a prime modulus `p` and a base generator `g` as its parameters.)

To initialize a KeyAgreement object, call one of its `init` methods:

```
public void init(Key key);

public void init(Key key, SecureRandom random);

public void init(Key key, AlgorithmParameterSpec params);
```

```
public void init(Key key, AlgorithmParameterSpec params,
                 SecureRandom random);
```

Executing a KeyAgreement Phase

Every key agreement protocol consists of a number of phases that need to be executed by each party involved in the key agreement.

To execute the next phase in the key agreement, call the `doPhase` method:

```
public Key doPhase(Key key, boolean lastPhase);
```

The `key` parameter contains the key to be processed by that phase. In most cases, this is the public key of one of the other parties involved in the key agreement, or an intermediate key that was generated by a previous phase. `doPhase` may return an intermediate key that you may have to send to the other parties of this key agreement, so they can process it in a subsequent phase.

The `lastPhase` parameter specifies whether or not the phase to be executed is the last one in the key agreement: A value of `FALSE` indicates that this is not the last phase of the key agreement (there are more phases to follow), and a value of `TRUE` indicates that this is the last phase of the key agreement and the key agreement is completed, i.e., `generateSecret` can be called next.

In the example of Diffie-Hellman Key Exchange between Two Parties , you call `doPhase` once, with `lastPhase` set to `TRUE`. In the example of Diffie-Hellman between three parties, you call `doPhase` twice: the first time with `lastPhase` set to `FALSE`, the 2nd time with `lastPhase` set to `TRUE`.

Generating the Shared Secret

After each party has executed all the required key agreement phases, it can compute the shared secret by calling one of the `generateSecret` methods:

```
public byte[] generateSecret();

public int generateSecret(byte[] sharedSecret, int offset);

public SecretKey generateSecret(String algorithm);
```

Key Management

A database called a "keystore" can be used to manage a repository of keys and certificates. (A *certificate* is a digitally signed statement from one entity, saying that the public key of some other entity has a particular value.)

Keystore Location

The user keystore is by default stored in a file named `.keystore` in the user's home directory, as determined by the `user.home` system property whose default value depends on the operating system:

- Solaris, Linux, and MacOS: `/home/`*`username`*`/`

- Windows: `C:\Users\username\`

Of course, keystore files can be located as desired. In some environments, it may make sense for multiple keystores to exist. For example, one keystore might hold a user's private keys, and another might hold certificates used to establish trust relationships.

In addition to the user's keystore, the JDK also maintains a system-wide keystore which is used to store trusted certificates from a variety of Certificate Authorities (CA's). These CA certificates can be used to help make trust decisions. For example, in SSL/TLS/DTLS when the `SunJSSE` provider is presented with certificates from a remote peer, the default trustmanager will consult one of the following files to determine if the connection is to be trusted:

- Solaris, Linux, and MacOS: `<java-home>/lib/security/cacerts`
- Windows: `<java-home>\lib\security\cacerts`

Instead of using the system-wide `cacerts` keystore, applications can set up and use their own keystores, or even use the user keystore described above.

Keystore Implementation

The `KeyStore` class supplies well-defined interfaces to access and modify the information in a keystore. It is possible for there to be multiple different concrete implementations, where each implementation is that for a particular *type* of keystore.

Currently, there are two command-line tools that make use of `KeyStore`: **`keytool`** and **`jarsigner`**. It is also used by the `Policy` reference implementation when it processes policy files specifying the permissions (allowed accesses to system resources) to be granted to code from various sources. Since `KeyStore` is publicly available, JDK users can write additional security applications that use it.

Applications can choose different *types* of keystore implementations from different providers, using the `getInstance` factory method in the `KeyStore` class. A keystore type defines the storage and data format of the keystore information, and the algorithms used to protect private keys in the keystore and the integrity of the keystore itself. Keystore implementations of different types are not compatible.

The default keystore implementation is "pkcs12". This is a cross-platform keystore based on the RSA PKCS12 Personal Information Exchange Syntax Standard. This standard is primarily meant for storing or transporting a user's private keys, certificates, and miscellaneous secrets. Arbitrary attributes can be associated with individual entries in a PKCS12 keystore.

```
keystore.type=pkcs12
```

To have tools and other applications use a different default keystore implementation, you can change that line to specify a different type.

Some applications, such as `keytool`, also let you override the default keystore type (via the `-storetype` command-line parameter).

> **Note:**
>
> Keystore type designations are case-insensitive. For example, "jks" would be considered the same as "JKS".

PKCS12 is the default and recommened keystore type. However, there are three other types of keystores that come with the JDK implementation.

1. **"jceks"** is an alternate proprietary keystore format to "jks" that uses Password-Based Encryption with Triple-DES.
   The "jceks" implementation can parse and convert a "jks" keystore file to the "jceks" format. You may upgrade your keystore of type "jks" to a keystore of type "jceks" by changing the password of a private-key entry in your keystore and specifying "`-storetype jceks`" as the keystore type. To apply the cryptographically strong(er) key protection supplied to a private key named "signkey" in your default keystore, use the following command, which will prompt you for the old and new key passwords:

   ```
   keytool -keypass -alias signkey -storetype jceks
   ```

   See keytool in *Java Platform, Standard Edition Tools Reference* .
2. **"jks"** is another option. It implements the keystore as a file, utilizing a proprietary keystore type (format). It protects each private key with its own individual password, and also protects the integrity of the entire keystore with a (possibly different) password.
3. **"dks"** is a domain keystore. It is a collection of keystores presented as a single logical keystore. The keystores that comprise a given domain are specified by configuration data whose syntax is described in `DomainLoadStoreParameter`.

Keystore implementations are provider-based. If you want to write your own KeyStore implementations, see How to Implement a Provider in the Java Cryptography Architecture.

The KeyStore Class

The `KeyStore` class supplies well-defined interfaces to access and modify the information in a keystore.

The `KeyStore` class is an Engine Classes and Algorithms.

Figure 2-16 KeyStore Class

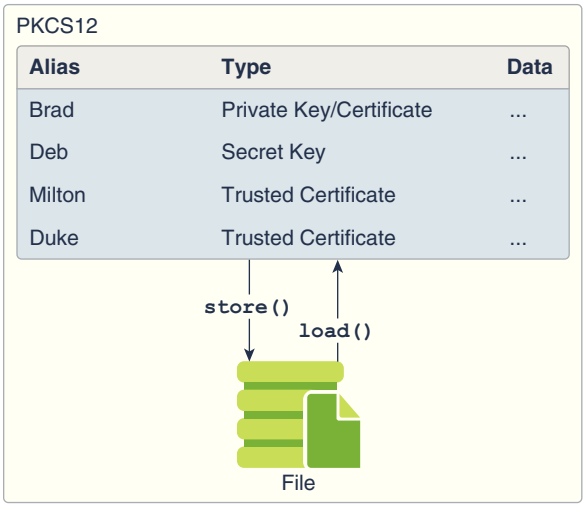

This class represents an in-memory collection of keys and certificates. `KeyStore` manages two types of entries:

- **Key Entry**: This type of keystore entry holds very sensitive cryptographic key information, which must be protected from unauthorized access. Typically, a key stored in this type of entry is a secret key, or a private key accompanied by the certificate chain authenticating the corresponding public key.

  Private keys and certificate chains are used by a given entity for self-authentication using digital signatures. For example, software distribution organizations digitally sign JAR files as part of releasing and/or licensing software.

- **Trusted Certificate Entry**: This type of entry contains a single public key certificate belonging to another party. It is called a *trusted certificate* because the keystore owner trusts that the public key in the certificate indeed belongs to the identity identified by the *subject* (owner) of the certificate.

  This type of entry can be used to authenticate other parties.

Each entry in a keystore is identified by an "alias" string. In the case of private keys and their associated certificate chains, these strings distinguish among the different ways in which the entity may authenticate itself. For example, the entity may authenticate itself using different certificate authorities, or using different public key algorithms.

Whether keystores are persistent, and the mechanisms used by the keystore if it is persistent, are not specified here. This convention allows use of a variety of techniques for protecting sensitive (e.g., private or secret) keys. Smart cards or other integrated cryptographic engines (SafeKeyper) are one option, and simpler mechanisms such as files may also be used (in a variety of formats).

The main `KeyStore` methods are described below.

Creating a KeyStore Object

KeyStore objects are obtained by using one of the KeyStore getInstance() method. See How Provider Implementations Are Requested and Supplied.

Loading a Particular Keystore into Memory

Before a `KeyStore` object can be used, the actual keystore data must be loaded into memory via the `load` method:

```
final void load(InputStream stream, char[] password)
```

The optional password is used to check the integrity of the keystore data. If no password is supplied, no integrity check is performed.

To create an empty keystore, you pass `null` as the `InputStream` argument to the `load` method.

A DKS keystore is loaded by passing a `DomainLoadStoreParameter` to the alternative load method:

```
final void load(KeyStore.LoadStoreParameter param)
```

Getting a List of the Keystore Aliases

All keystore entries are accessed via unique *aliases*. The `aliases` method returns an enumeration of the alias names in the keystore:

```
final Enumeration aliases()
```

Determining Keystore Entry Types

As stated in the KeyStore class, there are two different types of entries in a keystore. The following methods determine whether the entry specified by the given alias is a key/certificate or a trusted certificate entry, respectively:

```
final boolean isKeyEntry(String alias)
final boolean isCertificateEntry(String alias)
```

Adding/Setting/Deleting Keystore Entries

The `setCertificateEntry` method assigns a certificate to a specified alias:

```
final void setCertificateEntry(String alias, Certificate cert)
```

If `alias` doesn't exist, a trusted certificate entry with that alias is created. If `alias` exists and identifies a trusted certificate entry, the certificate associated with it is replaced by `cert`.

The `setKeyEntry` methods add (if `alias` doesn't yet exist) or set key entries:

```
final void setKeyEntry(String alias,
```

```
                        Key key,
                        char[] password,
                        Certificate[] chain)

final void setKeyEntry(String alias,
                        byte[] key,
                        Certificate[] chain)
```

In the method with `key` as a byte array, it is the bytes for a key in protected format. For example, in the keystore implementation supplied by the `SUN` provider, the `key` byte array is expected to contain a protected private key, encoded as an `EncryptedPrivateKeyInfo` as defined in the PKCS8 standard. In the other method, the `password` is the password used to protect the key.

The `deleteEntry` method deletes an entry:

```
final void deleteEntry(String alias)
```

PKCS #12 keystores support entries containing arbitrary attributes. Use the `PKCS12Attribute` class to create the attributes. When creating the new keystore entry use a constructor method that accepts attributes. Finally, use the following method to add the entry to the keystore:

```
final void setEntry(String alias, Entry entry,
                    ProtectionParameter protParam)
```

Getting Information from the Keystore

The `getKey` method returns the key associated with the given alias. The key is recovered using the given password:

```
final Key getKey(String alias, char[] password)
```

The following methods return the certificate, or certificate chain, respectively, associated with the given alias:

```
final Certificate getCertificate(String alias)
final Certificate[] getCertificateChain(String alias)
```

You can determine the name (`alias`) of the first entry whose certificate matches a given certificate via the following:

```
final String getCertificateAlias(Certificate cert)
```

PKCS #12 keystores support entries containing arbitrary attributes. Use the following method to retrieve an entry that may contain attributes:

```
final Entry getEntry(String alias, ProtectionParameter protParam)
```

and then use the `KeyStore.Entry.getAttributes` method to extract such attributes and use the methods of the `KeyStore.Entry.Attribute` interface to examine them.

Saving the KeyStore

The in-memory keystore can be saved via the `store` method:

```
final void store(OutputStream stream, char[] password)
```

The password is used to calculate an integrity checksum of the keystore data, which is appended to the keystore data.

A DKS keystore is stored by passing a `DomainLoadStoreParameter` to the alternative store method:

```
final void store(KeyStore.LoadStoreParameter param)
```

Algorithm Parameters Classes

Like `Keys` and `Keyspecs`, an algorithm's initialization parameters are represented by either `AlgorithmParameters` or `AlgorithmParameterSpecs`.

Depending on the use situation, algorithms can use the parameters directly, or the parameters might need to be converted into a more portable format for transmission or storage.

A *transparent* representation of a set of parameters (via `AlgorithmParameterSpec`) means that you can access each parameter value in the set individually. You can access these values through one of the `get` methods defined in the corresponding specification class (e.g., `DSAParameterSpec` defines `getP`, `getQ`, and `getG` methods, to access `p`, `q`, and `g`, respectively).

In contrast, the The AlgorithmParameters Class class supplies an *opaque* representation, in which you have no direct access to the parameter fields. You can only get the name of the algorithm associated with the parameter set (via `getAlgorithm`) and some kind of encoding for the parameter set (via `getEncoded`).

The AlgorithmParameterSpec Interface

`AlgorithmParameterSpec` is an interface to a transparent specification of cryptographic parameters. This interface contains no methods or constants. Its only purpose is to group (and provide type safety for) all parameter specifications. All parameter specifications must implement this interface.

The following are the algorithm parameter specification interfaces and classes in the `java.security.spec` and `javax.crypto.spec` packages:

- `DHParameterSpec`
- `DHGenParameterSpec`
- `DSAParameterSpec`
- `ECGenParameterSpec`
- `ECParameterSpec`
- `GCMParameterSpec`
- `IvParameterSpec`
- `MGF1ParameterSpec`
- `OAEPParameterSpec`
- `PBEParameterSpec`
- `PSSParameterSpec`
- `RC2ParameterSpec`
- `RC5ParameterSpec`
- `RSAKeyGenParameterSpec`

The following algorithm parameter specs are used specifically for XML digital signatures.

- `Interface C14NMethodParameterSpec`
- `Interface DigestMethodParameterSpec`
- `Interface SignatureMethodParameterSpec`
- `Interface TransformParameterSpec`
- `Interface ExcC14NParameterSpec`
- `Interface HMACParameterSpec`
- `Interface XPathFilter2ParameterSpec`
- `Interface XPathFilterParameterSpec`
- `XSLTTransformParameterSpec`

The AlgorithmParameters Class

The `AlgorithmParameters` class provides an opaque representation of cryptographic parameters.

The AlgorithmParameters Class

The `AlgorithmParameters` class is an Engine Classes and Algorithms .You can initialize the `AlgorithmParameters` class using a specific `AlgorithmParameterSpec` object, or by encoding the parameters in a recognized format. You can retrieve the resulting specification with the `getParameterSpec` method (see the following section).

Creating an AlgorithmParameters Object

`AlgorithmParameters` objects are obtained by using one of the `AlgorithmParameters` getInstance() static factory methods. For more information, see How Provider Implementations Are Requested and Supplied.

Initializing an AlgorithmParameters Object

Once an `AlgorithmParameters` object is instantiated, it must be initialized via a call to `init`, using an appropriate parameter specification or parameter encoding:

```
void init(AlgorithmParameterSpec paramSpec)
void init(byte[] params)
void init(byte[] params, String format)
```

In these `init` methods, `params` is an array containing the encoded parameters, and `format` is the name of the decoding format. In the `init` method with a `params` argument but no `format` argument, the primary decoding format for parameters is used. The primary decoding format is ASN.1, if an ASN.1 specification for the parameters exists.

Obtaining the Encoded Parameters

A byte encoding of the parameters represented in an `AlgorithmParameters` object may be obtained via a call to `getEncoded`:

```
byte[] getEncoded()
```

This method returns the parameters in their primary encoding format. The primary encoding format for parameters is ASN.1, if an ASN.1 specification for this type of parameters exists.

If you want the parameters returned in a specified encoding format, use

```
byte[] getEncoded(String format)
```

If `format` is null, the primary encoding format for parameters is used, as in the other `getEncoded` method.

Converting an AlgorithmParameters Object to a Transparent Specification

A transparent parameter specification for the algorithm parameters may be obtained from an `AlgorithmParameters` object via a call to `getParameterSpec`:

```
AlgorithmParameterSpec getParameterSpec(Class paramSpec)
```

`paramSpec` identifies the specification class in which the parameters should be returned. The specification class could be, for example, `DSAParameterSpec.class` to

indicate that the parameters should be returned in an instance of the `DSAParameterSpec` class. (This class is in the `java.security.spec` package.)

The AlgorithmParameterGenerator Class

The `AlgorithmParameterGenerator` class is an Engine Classes and Algorithms used to generate a set of **brand-new** parameters suitable for a certain algorithm (the algorithm is specified when an `AlgorithmParameterGenerator` instance is created). This object is used when you do not have an existing set of algorithm parameters, and want to generate one from scratch.

Creating an AlgorithmParameterGenerator Object

`AlgorithmParameterGenerator` objects are obtained by using one of the `AlgorithmParameterGenerator` getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing an AlgorithmParameterGenerator Object

The `AlgorithmParameterGenerator` object can be initialized in two different ways: an algorithm-independent manner or an algorithm-specific manner.

The algorithm-independent approach uses the fact that all parameter generators share the concept of a "size" and a source of randomness. The measure of size is universally shared by all algorithm parameters, though it is interpreted differently for different algorithms. For example, in the case of parameters for the DSA algorithm, "size" corresponds to the size of the prime modulus, in bits. (See Java Security Standard Algorithm Names to know more about the sizes for specific algorithms.) When using this approach, algorithm-specific parameter generation values--if any--default to some standard values. One `init` method that takes these two universally shared types of arguments:

```
void init(int size, SecureRandom random);
```

Another `init` method takes only a `size` argument and uses a system-provided source of randomness:

```
void init(int size)
```

A third approach initializes a parameter generator object using algorithm-specific semantics, which are represented by a set of algorithm-specific parameter generation values supplied in an `AlgorithmParameterSpec` object:

```
void init(AlgorithmParameterSpec genParamSpec,
                          SecureRandom random)

void init(AlgorithmParameterSpec genParamSpec)
```

To generate Diffie-Hellman system parameters, for example, the parameter generation values usually consist of the size of the prime modulus and the size of the random exponent, both specified in number of bits.

Generating Algorithm Parameters

Once you have created and initialized an AlgorithmParameterGenerator object, you can use the generateParameters method to generate the algorithm parameters:

```
AlgorithmParameters generateParameters()
```

The CertificateFactory Class

The `CertificateFactory` class defines the functionality of a certificate factory, which is used to generate certificate and certificate revocation list (CRL) objects from their encoding.

The `CertificateFactory` class is an Engine Classes and Algorithms.

A certificate factory for X.509 must return certificates that are an instance of `java.security.cert.X509Certificate`, and CRLs that are an instance of `java.security.cert.X509CRL`.

Creating a CertificateFactory Object

`CertificateFactory` objects are obtained by using one of the `getInstance()` static factory methods. For more information, see How Provider Implementations Are Requested and Supplied.

Generating Certificate Objects

To generate a certificate object and initialize it with the data read from an input stream, use the `generateCertificate` method:

```
final Certificate generateCertificate(InputStream inStream)
```

To return a (possibly empty) collection view of the certificates read from a given input stream, use the `generateCertificates` method:

```
final Collection generateCertificates(InputStream inStream)
```

Generating CRL Objects

To generate a certificate revocation list (CRL) object and initialize it with the data read from an input stream, use the `generateCRL` method:

```
final CRL generateCRL(InputStream inStream)
```

To return a (possibly empty) collection view of the CRLs read from a given input stream, use the `generateCRLs` method:

```
final Collection generateCRLs(InputStream inStream)
```

Generating CertPath Objects

The certificate path builder and validator for PKIX is defined by the Internet X.509 Public Key Infrastructure Certificate and CRL Profile, RFC 5280.

A certificate store implementation for retrieving certificates and CRLs from Collection and LDAP directories, using the PKIX LDAP V2 Schema is also available from the IETF as RFC 2587.

To generate a `CertPath` object and initialize it with data read from an input stream, use one of the following `generateCertPath` methods (with or without specifying the encoding to be used for the data):

```
final CertPath generateCertPath(InputStream inStream)

final CertPath generateCertPath(InputStream inStream,
                                String encoding)
```

To generate a `CertPath` object and initialize it with a list of certificates, use the following method:

```
final CertPath generateCertPath(List certificates)
```

To retrieve a list of the `CertPath` encoding supported by this certificate factory, you can call the `getCertPathEncodings` method:

```
final Iterator getCertPathEncodings()
```

The default encoding will be listed first.

How the JCA Might Be Used in a SSL/TLS Implementation

With an understanding of the JCA classes, consider how these classes might be combined to implement an advanced network protocol like SSL/TLS.

The SSL/TLS Overview section in the TLS and DTLS Protocols describes at a high level how the protocols work. As asymmetric (public key) cipher operations are much slower than symmetric operations (secret key), public key cryptography is used to establish secret keys which are then used to protect the actual application data. Vastly simplified, the SSL/TLS handshake involves exchanging initialization data, performing some public key operations to arrive at a secret key, and then using that key to encrypt further traffic.

> **Note:**
>
> The details presented here simply show how some of the above classes might be employed. This section will not present sufficient information for building a SSL/TLS implementation. For more information, see Java Secure Socket Extension (JSSE) Reference Guide and RFC 5246: The Transport Layer Security (TLS) Protocol, Version 1.2.

Assume that this SSL/TLS implementation will be made available as a JSSE provider. A concrete implementation of the `Provider` class is first written that will eventually be registered in the `Security` class' list of providers. This provider mainly provides a mapping from algorithm names to actual implementation classes. (that is: "SSLContext.TLS"->"com.foo.TLSImpl") When an application requests an "TLS" instance (via `SSLContext.getInstance("TLS")`), the provider's list is consulted for the requested algorithm, and an appropriate instance is created.

Before discussing details of the actual handshake, a quick review of some of the JSSE's architecture is needed. The heart of the JSSE architecture is the `SSLContext`. The context eventually creates end objects (`SSLSocket` and `SSLEngine`) which actually implement the SSL/TLS protocol. `SSLContexts` are initialized with two callback classes, `KeyManager` and `TrustManager`, which allow applications to first select authentication material to send and second to verify credentials sent by a peer.

A JSSE `KeyManager` is responsible for choosing which credentials to present to a peer. Many algorithms are possible, but a common strategy is to maintain a RSA or DSA public/private key pair along with a `X509Certificate` in a `KeyStore` backed by a disk file. When a `KeyStore` object is initialized and loaded from the file, the file's raw bytes are converted into `PublicKey` and `PrivateKey` objects using a `KeyFactory`, and a certificate chain's bytes are converted using a `CertificateFactory`. When a credential is needed, the `KeyManager` simply consults this `KeyStore` object and determines which credentials to present.

A `KeyStore`'s contents might have originally been created using a utility such as `keytool`. `keytool` creates a RSA or DSA `KeyPairGenerator` and initializes it with an appropriate keysize. This generator is then used to create a `KeyPair` which `keytool` would store along with the newly-created certificate in the `KeyStore`, which is eventually written to disk.

A JSSE `TrustManager` is responsible for verifying the credentials received from a peer. There are many ways to verify credentials: one of them is to create a `CertPath` object, and let the JDK's built-in Public Key Infrastructure (PKI) framework handle the validation. Internally, the CertPath implementation might create a `Signature` object, and use that to verify that the each of the signatures in the certificate chain.

With this basic understanding of the architecture, we can look at some of the steps in the SSL/TLS handshake. The client begins by sending a ClientHello message to the server. The server selects a ciphersuite to use, and sends that back in a ServerHello message, and begins creating JCA objects based on the suite selection. We'll use server-only authentication in the following examples.

Figure 2-17 SSL/TLS Messages

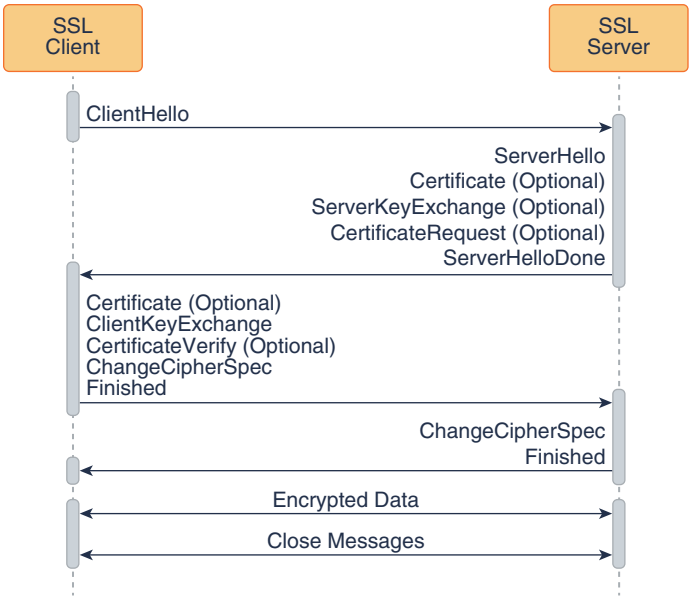

Server-only authentication is described in the following examples. The examples are vastly simplified, but gives an idea of how the JSSE classes might be combined to create a higher level protocol:

Example 2-9 SSL/TLS Server Uses a RSA-based ciphersuite Such as TLS_RSA_WITH_AES_128_CBC_SHA

```
KeyManagerTrustManagerSecureRandomCipherPublicKeyPrivateKeyCipher
```

Example 2-10 Choose an Ephemeral Diffie-Hellman Key Agreement Algorithm Along with the DSA Signature Algorithm such as TLS_DHE_DSS_WITH_AES_128_CBC_SHA

```
KeyPairGeneratorKeyFactoryDHPublicKeySpecKeyAgreementPrivateKeyServerKeyEx
changeKeyFactoryKeyAgreement
```

Once the actual encryption keys have been established, the secret key is used to initialize a symmetric `Cipher` object, and this cipher is used to protect all data in transit. To help determine if the data has been modified, a `MessageDigest` is created and receives a copy of the data destined for the network. When the packet is complete, the digest (hash) is appended to data, and the entire packet is encrypted by the `Cipher`. If a block cipher such as AES is used, the data must be padded to make a complete block. On the remote side, the steps are simply reversed.

Cryptographic Strength Configuration

You can configure the cryptographic strength of the Java Cryptography Extension (JCE) architecture using jurisdiction policy files (see Jurisdiction Policy File Format) and the security properties file.

Prior to Oracle Java JDK 9, the default cryptographic strength allowed by Oracle implementations was “strong but limited” (for example AES keys limited to 128 bits). To remove this restriction, administrators could download and install a separate

"unlimited strength" Jurisdiction Policy Files bundle. The Jurisdiction Policy File mechanism was reworked for JDK 9. It now allows for much more flexible configuration. The Oracle JDK now ships with a default value of "unlimited" rather than "limited". As always, administrators and users must still continue to follow all import/export guidelines for their geographical locations. The active cryptographic strength is now determined using a Security Property (typically set in the `java.security` properties file), in combination with the jurisdiction policy files found in the configuration directory.

All the necessary JCE policy files to provide either unlimited cryptographic strength or strong but limited cryptographic strength are bundled with the JDK.

Cryptographic Strength Settings

Each directory under `<java_home>/conf/security/policy` represents a set of policy configurations defined by the jurisdiction policy files that they contain. You activate a particular cryptographic strength setting represented by the policy files in a directory by setting the `crypto.policy` Security Property (configured in the file `<java_home>/conf/security/java.security`) to point to that directory.

The JDK comes bundled with two such directories, `limited` and `unlimited`, each containing a number of policy files. By default, the `crypto.policy` Security Property is set to:

```
crypto.policy = unlimited
```

The overall value is the intersection of the files contained within the directory. These policy files settings are VM-wide, and affect all applications running on this VM. If you want to override cryptographic strength at the application level, see How to Make Applications Exempt from Cryptographic Restrictions.

Unlimited Directory Contents

The `unlimited` directory contains the following policy files:

- `<java_home>/conf/security/unlimited/default_US_export.policy`

  ```
  // Default US Export policy file.
  grant {
  // There is no restriction to any algorithms.
       permission javax.crypto.CryptoAllPermission;
  };
  ```

  > **Note:**
  >
  > As there are no current restrictions on export of cryptography from the United States, the `default_US_export.policy` file is set with no restrictions.

- `<java_home>/conf/security/unlimited/default_local.policy`

  ```
  // Country specific policy file for countries with no limits on crypto
  strength.
  grant {
  ```

```
// There is no restriction to any algorithms.
     permission javax.crypto.CryptoAllPermission;
};
```

> **Note:**
>
> Depending on the country, there may be local restrictions, but as this policy file is located in the `unlimited` directory, there are no restrictions listed here.

To select unlimited cryptographic strength as defined in these two files set `crypto.policy = unlimited` in the file `<java_home>/conf/security/java.security`.

Limited Directory Contents

The `limited` directory currently contains the following policy files:

- `<java_home>/conf/security/limited/default_US_export.policy`

```
// Default US Export policy file.
grant {
// There is no restriction to any algorithms.
     permission javax.crypto.CryptoAllPermission;
};
```

> **Note:**
>
> Even though this is in the `limited` directory, as there are no current restrictions on export of cryptography from the United States, the `default_US_export.policy` file is set with no restrictions.

- `<java_home>/conf/security/limited/default_local.policy`

```
// Some countries have import limits on crypto strength. This policy
file
// is worldwide importable.

grant {
    permission javax.crypto.CryptoPermission "DES", 64;
    permission javax.crypto.CryptoPermission "DESede", *;
    permission javax.crypto.CryptoPermission "RC2", 128,

"javax.crypto.spec.RC2ParameterSpec", 128;
    permission javax.crypto.CryptoPermission "RC4", 128;
    permission javax.crypto.CryptoPermission "RC5", 128,
          "javax.crypto.spec.RC5ParameterSpec", *, 12, *;
    permission javax.crypto.CryptoPermission "RSA", *;
    permission javax.crypto.CryptoPermission *, 128;
};
```

> **Note:**
>
> This local policy file shows the default restrictions. It should be allowed by any country, including those that have import restrictions, but please obtain legal guidance.

- `<java_home>/conf/security/limited/exempt_local.policy`

```
// Some countries have import limits on crypto strength, but may allow
for
// these exemptions if the exemption mechanism is used.

grant {
    // There is no restriction to any algorithms if KeyRecovery is
enforced.
    permission javax.crypto.CryptoPermission *, "KeyRecovery";

    // There is no restriction to any algorithms if KeyEscrow is
enforced.
    permission javax.crypto.CryptoPermission *, "KeyEscrow";

    // There is no restriction to any algorithms if KeyWeakening is
enforced.
    permission javax.crypto.CryptoPermission *, "KeyWeakening";
};
```

> **Note:**
>
> Countries that have import restrictions should use “limited”, but these restrictions could be relaxed if the exemption mechanism can be employed. See How to Make Applications Exempt from Cryptographic Restrictions. Please obtain legal guidance for your situation.

Custom Cryptographic Strength Settings

To set up restrictions to cryptographic strength that are different than the settings in the policy files in the `limited` or `unlimited` directory, you can create a new directory, parallel with `limited` and `unlimited`, and place your policy files there. For example, you may create a directory called `custom`. In this `custom` directory you include the files `default_*export.policy` and/or `exempt_*local.policy`.

To select cryptographic strength as defined in the files in the `custom` directory, set `crypto.policy = custom` in the file `<java_home>/conf/security/java.security`.

Jurisdiction Policy File Format

JCA represents its jurisdiction policy files as Java-style policy files with corresponding permission statements. As described in Cryptographic Strength Configuration, a Java policy file specifies what permissions are allowed for code from specified code sources. A permission represents access to a system resource. In the case of JCA,

the "resources" are cryptography algorithms, and code sources do not need to be specified, because the cryptographic restrictions apply to all code.

A jurisdiction policy file consists of a very basic "grant entry" containing one or more "permission entries."

```
grant {
    <permission entries>;
};
```

The format of a permission entry in a jurisdiction policy file is:

```
permission <crypto permission class name>
    [<alg_name>
        [
            [, <exemption mechanism name>]
            [, <maxKeySize>
                [, <AlgorithmParameterSpec class name>,
                       <parameters for constructing an
AlgorithmParameterSpec object>
                ]
            ]
        ]
    ];
```

A sample jurisdiction policy file that includes restricting the AES algorithm to maximum key sizes of 128 bits is:

```
grant {
    permission javax.crypto.CryptoPermission "AES", 128;
    // ...
};
```

A permission entry must begin with the word `permission`. Items that appear in a permission entry must appear in the specified order. An entry is terminated with a semicolon. Case is unimportant for the identifiers (`grant`, `permission`) but is significant for the *<crypto permission class name>* or for any string that is passed in as a value. An asterisk (`*`) can be used as a wildcard for any permission entry option. For example, an asterisk for an *<alg_name>* option means "all algorithms."

The following table describes a permission entry's options:

Table 2-1 Permission Entry Options

Option	Description
<crypto permission class name>	Specific permission class name, such as `javax.crypto.CryptoPermission`. Required. A crypto permission class reflects the ability of an application to use certain algorithms with certain key sizes in certain environments. There are two crypto permission classes: `CryptoPermission` and `CryptoAllPermission`. The special `CryptoAllPermission` class implies all cryptography-related permissions, that is, it specifies that there are no cryptography-related restrictions.
<alg_name>	Quoted string specifying the standard name of a cryptography algorithm, such as "AES" or "RSA". Optional.
<exemption mechanism name>	Quoted string indicating an exemption mechanism which, if enforced, enables a reduction in cryptographic restrictions. Optional. Exemption mechanism names that can be used include "KeyRecovery" "KeyEscrow", and "KeyWeakening".
<maxKeySize>	Integer specifying the maximum key size (in bits) allowed for the specified algorithm. Optional.
<AlgorithmParameterSpec class name>	Class name that specifies the strength of the algorithm. Optional. For some algorithms, it may not be sufficient to specify the algorithm strength in terms of just a key size. For example, in the case of the "RC5" algorithm, the number of rounds must also be considered. For algorithms whose strength needs to be expressed as more than a key size, use this option to specify the `AlgorithmParameterSpec` class name that does this (such as `javax.crypto.spec.RC5ParameterSpec` for the "RC5" algorithm).
<parameters for constructing an AlgorithmParameterSpec object>	List of parameters for constructing the specified `AlgorithmParameterSpec` object. Required if *<AlgorithmParameterSpec class name>* has been specified and requires parameters.

How to Make Applications Exempt from Cryptographic Restrictions

> **NOT_SUPPORTED:**
>
> This section should be ignored by most application developers. It is only for people whose applications may be exported to those few countries whose governments mandate cryptographic restrictions, if it is desired that such applications have fewer cryptographic restrictions than those mandated.

By default, an application can use cryptographic algorithms of any strength. However, due to import control restrictions by the governments of a few countries, you may have to limit those algorithms' strength. The JCA framework includes an ability to enforce restrictions regarding the maximum strengths of cryptographic algorithms available to applications in different jurisdiction contexts (locations). You specify these restrictions in jurisdiction policy files. For more information about jurisdiction policy files and how to create and configure them, see Cryptographic Strength Configuration.

It is possible that the governments of some or all such countries may allow certain applications to become exempt from some or all cryptographic restrictions. For example, they may consider certain types of applications as "special" and thus exempt. Or they may exempt any application that utilizes an "exemption mechanism," such as key recovery. Applications deemed to be exempt could get access to stronger cryptography than that allowed for non-exempt applications in such countries.

In order for an application to be recognized as "exempt" at runtime, it must meet the following conditions:

- It must have a permission policy file bundled with it in a JAR file. The permission policy file specifies what cryptography-related permissions the application has, and under what conditions (if any).
- The JAR file containing the application and the permission policy file must have been signed using a code-signing certificate issued after the application was accepted as exempt.

Below are sample steps required in order to make an application exempt from some cryptographic restrictions. This is a basic outline that includes information about what is required by JCA in order to recognize and treat applications as being exempt. You will need to know the exemption requirements of the particular country or countries in which you would like your application to be able to be run but whose governments require cryptographic restrictions. You will also need to know the requirements of a JCA framework vendor that has a process in place for handling exempt applications. Consult such a vendor for further information.

> **Note:**
>
> The `SunJCE` provider does not supply an implementation of the `ExemptionMechanismSpi` class

1. Write and Compile Your Application Code
2. Create a Permission Policy File Granting Appropriate Cryptographic Permissions
3. Prepare for Testing
   a. Apply for Government Approval From the Government Mandating Restrictions.

   b. Get a Code-Signing Certificate
   c. Bundle the Application and Permission Policy File into a JAR file
   d. Step 7.1: Get a Code-Signing Certificate
   e. Set Up Your Environment Like That of a User in a Restricted Country
   f. (only for applications using exemption mechanisms) Install a Provider Implementing the Exemption Mechanism Specified by the entry in the Permission Policy File
4. Test Your Application
5. Apply for U.S. Government Export Approval If Required
6. Deploy Your Application

Special Code Requirements for Applications that Use Exemption Mechanisms

When an application has a permission policy file associated with it (in the same JAR file) and that permission policy file specifies an exemption mechanism, then when the `Cipher` `getInstance` method is called to instantiate a `Cipher`, the JCA code searches the installed providers for one that implements the specified exemption mechanism. If it finds such a provider, JCA instantiates an `ExemptionMechanism` API object associated with the provider's implementation, and then associates the `ExemptionMechanism` object with the `Cipher` returned by `getInstance`.

After instantiating a Cipher, and prior to initializing it (via a call to the Cipher `init` method), your code must call the following `Cipher` method:

```
public ExemptionMechanism getExemptionMechanism()
```

This call returns the `ExemptionMechanism` object associated with the Cipher. You must then initialize the exemption mechanism implementation by calling the following method on the returned `ExemptionMechanism`:

```
public final void init(Key key)
```

The argument you supply should be the same as the argument of the same types that you will subsequently supply to a Cipher `init` method.

Once you have initialized the `ExemptionMechanism`, you can proceed as usual to initialize and use the Cipher.

Permission Policy Files

In order for an application to be recognized at runtime as being "exempt" from some or all cryptographic restrictions, it must have a permission policy file bundled with it in a JAR file. The permission policy file specifies what cryptography-related permissions the application has, and under what conditions (if any).

The format of a permission entry in a permission policy file that accompanies an exempt application is the same as the format for a jurisdiction policy file downloaded with the JDK, which is:

```
permission <crypto permission class name>
    [<alg_name>
```

```
        [
            [, <exemption mechanism name>]
            [, <maxKeySize>
                [, <AlgorithmParameterSpec class name>,
                       <parameters for constructing an
AlgorithmParameterSpec object>
                ]
            ]
        ]
    ];
```

See Jurisdiction Policy File Format.

Permission Policy Files for Exempt Applications

Some applications may be allowed to be completely unrestricted. Thus, the permission policy file that accompanies such an application usually just needs to contain the following:

```
grant {
    // There are no restrictions to any algorithms.
    permission javax.crypto.CryptoAllPermission;
};
```

If an application just uses a single algorithm (or several specific algorithms), then the permission policy file could simply mention that algorithm (or algorithms) explicitly, rather than granting `CryptoAllPermission`.

For example, if an application just uses the Blowfish algorithm, the permission policy file doesn't have to grant `CryptoAllPermission` to all algorithms. It could just specify that there is no cryptographic restriction if the Blowfish algorithm is used. In order to do this, the permission policy file would look like the following:

```
grant {
    permission javax.crypto.CryptoPermission "Blowfish";
};
```

Permission Policy Files for Applications Exempt Due to Exemption Mechanisms

If an application is considered "exempt" if an exemption mechanism is enforced, then the permission policy file that accompanies the application must specify one or more exemption mechanisms. At run time, the application will be considered exempt if any of those exemption mechanisms is enforced. Each exemption mechanism must be specified in a permission entry that looks like the following:

```
    // No algorithm restrictions if specified
    // exemption mechanism is enforced.
    permission javax.crypto.CryptoPermission *,
        "<ExemptionMechanismName>";
```

where *<ExemptionMechanismName>* specifies the name of an exemption mechanism. The list of possible exemption mechanism names includes:

- `KeyRecovery`

- `KeyEscrow`
- `KeyWeakening`

As an example, suppose your application is exempt if either key recovery or key escrow is enforced. Then your permission policy file should contain the following:

```
grant {
    // No algorithm restrictions if KeyRecovery is enforced.
    permission javax.crypto.CryptoPermission *, "KeyRecovery";

    // No algorithm restrictions if KeyEscrow is enforced.
    permission javax.crypto.CryptoPermission *, "KeyEscrow";
};
```

> **Note:**
>
> Permission entries that specify exemption mechanisms should *not* also specify maximum key sizes. The allowed key sizes are actually determined from the installed exempt jurisdiction policy files, as described in the next section.

How Bundled Permission Policy Files Affect Cryptographic Permissions

At runtime, when an application instantiates a `Cipher` (via a call to its `getInstance` method) and that application has an associated permission policy file, JCA checks to see whether the permission policy file has an entry that applies to the algorithm specified in the `getInstance` call. If it does, and the entry grants `CryptoAllPermission` or does not specify that an exemption mechanism must be enforced, it means there is no cryptographic restriction for this particular algorithm.

If the permission policy file has an entry that applies to the algorithm specified in the `getInstance` call and the entry *does* specify that an exemption mechanism must be enforced, then the exempt jurisdiction policy file(s) are examined. If the exempt permissions include an entry for the relevant algorithm and exemption mechanism, and that entry is implied by the permissions in the permission policy file bundled with the application, and if there is an implementation of the specified exemption mechanism available from one of the registered providers, then the maximum key size and algorithm parameter values for the `Cipher` are determined from the exempt permission entry.

If there is no exempt permission entry implied by the relevant entry in the permission policy file bundled with the application, or if there is no implementation of the specified exemption mechanism available from any of the registered providers, then the application is only allowed the standard default cryptographic permissions.

Standard Names

The Standard Names document contains information about the algorithm specifications.

Java Security Standard Algorithm Names describes the standard names for algorithms, certificate and keystore types that the JDK Security API requires and uses.

It also contains more information about the algorithm specifications. Specific provider information can be found in JDK Providers Documentation.

Cryptographic implementations in the JDK are distributed through several different providers primarily for historical reasons (`Sun`, `SunJSSE`, `SunJCE`, `SunRsaSign`). Note these providers may not be available on all JDK implementations, and therefore, truly portable applications should call `getInstance()` without specifying specific providers. Applications specifying a particular provider may not be able to take advantage of native providers tuned for an underlying operating environment (such as PKCS or Microsoft's CAPI).

The `SunPKCS11` provider itself does not contain any cryptographic algorithms, but instead, directs requests into an underlying PKCS11 implementation. Consult PKCS#11 Reference Guide and the underlying PKCS11 implementation to determine if a desired algorithm will be available through the PKCS11 provider. Likewise, on Windows systems, the `SunMSCAPI` provider does not provide any cryptographic functionality, but instead routes requests to the underlying Operating System for handling.

Packaging Your Application

You can package an application in three different kinds of modules:

- Named or explicit module: A module that appears on the module path and contains module configuration information in the `module-info.class` file.
- Automatic module: A module that appears on the module path, but does not contain module configuration information in a `module-info.class` file (essentially a "regular" JAR file).
- Unnamed module: A module that appears on the class path. It may or may not have a `module-info.class` file; this file is ignored.

It is recommended that you package your applications in named modules as they provide better performance, stronger encapsulation, and simpler configuration. They also offer greater flexibility; you can use them with non-modular JDKs or even as unnamed modules by specifying them in a modular JDK's class path.

For more information about modules, see The State of the Module System and JEP 261: Module System

Additional JCA Code Samples

These examples illustrate use of several JCA mechanisms. See also Sample Programs for Diffie-Hellman Key Exchange, AES/GCM, and HMAC-SHA256

Topics

Computing a MessageDigest Object

Generating a Pair of Keys

Generating and Verifying a Signature Using Generated Keys

Generating/Verifying Signatures Using Key Specifications and KeyFactory

Determining If Two Keys Are Equal

Reading Base64-Encoded Certificates

Parsing a Certificate Reply

Using Encryption

Using Password-Based Encryption

Computing a MessageDigest Object

An example describing the procedure to compute a MessageDigest object.

1. Create the `MessageDigest` object, as in the following example:

```
MessageDigest sha = MessageDigest.getInstance("SHA-256");
```

   This call assigns a properly initialized message digest object to the `sha` variable. The implementation implements the Secure Hash Algorithm (SHA-256), as defined in the National Institute for Standards and Technology's (NIST) FIPS 180-4 document.

2. Suppose we have three byte arrays, `i1`, `i2` and `i3`, which form the total input whose message digest we want to compute. This digest (or "hash") could be calculated via the following calls:

```
sha.update(i1);
sha.update(i2);
sha.update(i3);
byte[] hash = sha.digest();
```

3. Optional: An equivalent alternative series of calls would be:

```
sha.update(i1);
sha.update(i2);
byte[] hash = sha.digest(i3);
```

   After the message digest has been calculated, the message digest object is automatically reset and ready to receive new data and calculate its digest. All former state (i.e., the data supplied to `update` calls) is lost.

Example 2-11 Hash Implementations Through Cloning

Some hash implementations may support intermediate hashes through cloning. Suppose we want to calculate separate hashes for:

- `i1`
- `i1 and i2`
- `i1, i2, and i3`

The following is one way to calculate these hashes; however, this code works only if the SHA-256 implementation is cloneable:

```
/* compute the hash for i1 */
sha.update(i1);
byte[] i1Hash = sha.clone().digest();

/* compute the hash for i1 and i2 */
sha.update(i2);
byte[] i12Hash = sha.clone().digest();

/* compute the hash for i1, i2 and i3 */
sha.update(i3);
byte[] i123hash = sha.digest();
```

Example 2-12 Determine if the Hash Implementation is Cloneable or not Cloneable

```
MessageDigest

try {
    // try and clone it
    /* compute the hash for i1 */
    sha.update(i1);
    byte[] i1Hash = sha.clone().digest();
    // ...
    byte[] i123hash = sha.digest();
} catch (CloneNotSupportedException cnse) {
    // do something else, such as the code shown below
}
```

Example 2-13 Compute Intermediate Digests if the Hash Implementation is not Cloneable

```
MessageDigest md1 = MessageDigest.getInstance("SHA-256");
MessageDigest md2 = MessageDigest.getInstance("SHA-256");
MessageDigest md3 = MessageDigest.getInstance("SHA-256");

byte[] i1Hash = md1.digest(i1);

md2.update(i1);
byte[] i12Hash = md2.digest(i2);

md3.update(i1);
md3.update(i2);
byte[] i123Hash = md3.digest(i3);
```

Generating a Pair of Keys

In this example we will generate a public-private key pair for the algorithm named "DSA" (Digital Signature Algorithm), and use this keypair in future examples. We will

generate keys with a 2048-bit modulus. We don't care which provider supplies the algorithm implementation.

Creating the Key Pair Generator

The first step is to get a key pair generator object for generating keys for the DSA algorithm:

```
KeyPairGenerator keyGen = KeyPairGenerator.getInstance("DSA");
```

Initializing the Key Pair Generator

The next step is to initialize the key pair generator. In most cases, algorithm-independent initialization is sufficient, but in some cases, algorithm-specific initialization is used.

Algorithm-Independent Initialization

All key pair generators share the concepts of a keysize and a source of randomness. The `KeyPairGenerator` class initialization methods at a minimum needs a keysize. If the source of randomness is not explicitly provided, a `SecureRandom` implementation of the highest-priority installed provider will be used. Thus to generate keys with a keysize of 2048, simply call:

```
keyGen.initialize(2048);
```

The following code illustrates how to use a specific, additionally seeded `SecureRandom` object:

```
SecureRandom random = SecureRandom.getInstance("DRBG", "SUN");
random.setSeed(userSeed);
keyGen.initialize(2048, random);
```

Since no other parameters are specified when you call the above algorithm-independent `initialize` method, it is up to the provider what to do about the algorithm-specific parameters (if any) to be associated with each of the keys. The provider may use precomputed parameter values or may generate new values.

Algorithm-Specific Initialization

For situations where a set of algorithm-specific parameters already exists (such as "community parameters" in DSA), there are two `initialize` methods that have an `AlgorithmParameterSpec` argument. Suppose your key pair generator is for the "DSA" algorithm, and you have a set of DSA-specific parameters, `p`, `q`, and `g`, that you would like to use to generate your key pair. You could execute the following code to initialize your key pair generator (recall that `DSAParameterSpec` is an `AlgorithmParameterSpec`):

```
DSAParameterSpec dsaSpec = new DSAParameterSpec(p, q, g);
keyGen.initialize(dsaSpec);
```

Generating the Pair of Keys

The final step is actually generating the key pair. No matter which type of initialization was used (algorithm-independent or algorithm-specific), the same code is used to generate the `KeyPair`:

```
KeyPair pair = keyGen.generateKeyPair();
```

Generating and Verifying a Signature Using Generated Keys

Examples of generating and verifying a signature using generated keys.

The following signature generation and verification examples use the `KeyPair` generated in the Generating a Pair of Keys .

Generating a Signature

We first create a Signature Class object:

```
Signature dsa = Signature.getInstance("SHA256withDSA");
```

Next, using the key pair generated in the key pair example, we initialize the object with the private key, then sign a byte array called `data`.

```
/* Initializing the object with a private key */
 PrivateKey priv = pair.getPrivate();
 dsa.initSign(priv);

/* Update and sign the data */
 dsa.update(data);
 byte[] sig = dsa.sign();
```

Verifying a Signature

Verifying the signature is straightforward. (Note that here we also use the key pair generated in the key pair example.)

```
/* Initializing the object with the public key */
PublicKey pub = pair.getPublic();
dsa.initVerify(pub);

/* Update and verify the data */
dsa.update(data);
boolean verifies = dsa.verify(sig);
System.out.println("signature verifies: " + verifies);
```

Generating/Verifying Signatures Using Key Specifications and KeyFactory

Sample code that is used to generate and verify signatures using key specifications and `KeyFactory`.

Suppose that, rather than having a public/private key pair (as, for example, was generated in the Generating a Pair of Keys above), you simply have the components of your DSA private key: `x` (the private key), `p` (the prime), `q` (the sub-prime), and `g` (the base).

Further suppose you want to use your private key to digitally sign some data, which is in a byte array named `someData`. You would do the following steps, which also illustrate creating a key specification and using a key factory to obtain a `PrivateKey` from the key specification (`initSign` requires a `PrivateKey`):

```
DSAPrivateKeySpec dsaPrivKeySpec = new DSAPrivateKeySpec(x, p, q, g);

KeyFactory keyFactory = KeyFactory.getInstance("DSA");
PrivateKey privKey = keyFactory.generatePrivate(dsaPrivKeySpec);

Signature sig = Signature.getInstance("SHA256withDSA");
sig.initSign(privKey);
sig.update(someData);
byte[] signature = sig.sign();
```

Suppose Alice wants to use the data you signed. In order for her to do so, and to verify your signature, you need to send her three things:

1. The data
2. The signature
3. The public key corresponding to the private key you used to sign the data

You can store the `someData` bytes in one file, and the `signature` bytes in another, and send those to Alice.

For the public key, assume, as in the signing example above, you have the components of the DSA public key corresponding to the DSA private key used to sign the data. Then you can create a `DSAPublicKeySpec` from those components:

```
DSAPublicKeySpec dsaPubKeySpec = new DSAPublicKeySpec(y, p, q, g);
```

You still need to extract the key bytes so that you can put them in a file. To do so, you can first call the `generatePublic` method on the DSA key factory already created in the example above:

```
PublicKey pubKey = keyFactory.generatePublic(dsaPubKeySpec);
```

Then you can extract the (encoded) key bytes via the following:

```
byte[] encKey = pubKey.getEncoded();
```

You can now store these bytes in a file, and send it to Alice along with the files containing the data and the signature.

Now, assume Alice has received these files, and she copied the data bytes from the data file to a byte array named `data`, the signature bytes from the signature file to a byte array named `signature`, and the encoded public key bytes from the public key file to a byte array named `encodedPubKey`.

Alice can now execute the following code to verify the signature. The code also illustrates how to use a key factory in order to instantiate a DSA public key from its encoding (`initVerify` requires a `PublicKey`).

```
X509EncodedKeySpec pubKeySpec = new X509EncodedKeySpec(encodedPubKey);

KeyFactory keyFactory = KeyFactory.getInstance("DSA");
PublicKey pubKey = keyFactory.generatePublic(pubKeySpec);

Signature sig = Signature.getInstance("SHA256withDSA");
sig.initVerify(pubKey);
sig.update(data);
sig.verify(signature);
```

Note:

In the above, Alice needed to generate a `PublicKey` from the encoded key bits, since `initVerify` requires a `PublicKey` . Once she has a `PublicKey`, she could also use the `KeyFactorygetKeySpec` method to convert it to a `DSAPublicKeySpec` so that she can access the components, if desired, as in:

```
DSAPublicKeySpec dsaPubKeySpec =
    (DSAPublicKeySpec)keyFactory.getKeySpec(pubKey,
DSAPublicKeySpec.class);
```

Now she can access the DSA public key components `y`, `p`, `q`, and `g` through the corresponding "get" methods on the `DSAPublicKeySpec` class (`getY`, `getP`, `getQ`, and `getG`).

Generating Random Numbers

The following code sample illustrates generating random numbers configured with different security strengths using a DRBG implementation of the `SecureRandom` class:

```
    SecureRandom drbg;
    byte[] buffer = new byte[32];

    // Any DRBG can be provided
    drbg = SecureRandom.getInstance("DRBG");
    drbg.nextBytes(buffer);

    SecureRandomParameters params = drbg.getParameters();
    if (params instanceof DrbgParameters.Instantiation) {
        DrbgParameters.Instantiation ins = (DrbgParameters.Instantiation)
params;
        if (ins.getCapability().supportsReseeding()) {
            drbg.reseed();
        }
    }

    // The following call requests a weak DRBG instance. It is only
    // guaranteed to support 112 bits of security strength.
    drbg = SecureRandom.getInstance("DRBG",
        DrbgParameters.instantiation(112, NONE, null));

    // Both the next two calls will likely fail, because drbg could be
    // instantiated with a smaller strength with no prediction resistance
    // support.
    drbg.nextBytes(buffer,
        DrbgParameters.nextBytes(256, false, "more".getBytes()));
    drbg.nextBytes(buffer,
        DrbgParameters.nextBytes(112, true, "more".getBytes()));

    // The following call requests a strong DRBG instance, with a
    // personalization string. If it successfully returns an instance,
    // that instance is guaranteed to support 256 bits of security strength
    // with prediction resistance available.
    drbg = SecureRandom.getInstance("DRBG", DrbgParameters.instantiation(
        256, PR_AND_RESEED, "hello".getBytes()));

    // Prediction resistance is not requested in this single call,
    // but an additional input is used.
    drbg.nextBytes(buffer,
        DrbgParameters.nextBytes(-1, false, "more".getBytes()));

    // Same for this call.
    drbg.reseed(DrbgParameters.reseed(false, "extra".getBytes()));
```

Determining If Two Keys Are Equal

Example code for determining if two keys are equal.

In many cases you would like to know if two keys are equal; however, the default method `java.lang.Object.equals` may not give the desired result. The most provider-independent approach is to compare the encoded keys. If this comparison isn't appropriate (for example, when comparing an `RSAPrivateKey` and an `RSAPrivateCrtKey`), you should compare each component.

The following code demonstrates this idea:

```
static boolean keysEqual(Key key1, Key key2) {
    if (key1.equals(key2)) {
       return true;
    }

    if (Arrays.equals(key1.getEncoded(), key2.getEncoded())) {
       return true;
    }

 // More code for different types of keys here.
 // For example, the following code can check if
 // an RSAPrivateKey and an RSAPrivateCrtKey are equal:
 // if ((key1 instanceof RSAPrivateKey) &&
 //     (key2 instanceof RSAPrivateKey)) {
 //     if ((key1.getModulus().equals(key2.getModulus())) &&
 //         (key1.getPrivateExponent().equals(
 //                                       key2.getPrivateExponent()))) {
 //         return true;
 //     }
 // }

     return false;
 }
```

Reading Base64-Encoded Certificates

The following example reads a file with Base64-encoded certificates, which are each bounded at the beginning by

```
-----BEGIN CERTIFICATE-----
```

and at the end by

```
-----END CERTIFICATE-----
```

We convert the `FileInputStream` (which does not support `mark` and `reset`) to a `ByteArrayInputStream` (which supports those methods), so that each call to

`generateCertificate` consumes only one certificate, and the read position of the input stream is positioned to the next certificate in the file:

```
 try (FileInputStream fis = new FileInputStream(filename);
     BufferedInputStream bis = new BufferedInputStream(fis)) {
     CertificateFactory cf = CertificateFactory.getInstance("X.509");
     while (bis.available() > 0) {
         Certificate cert = cf.generateCertificate(bis);
         System.out.println(cert.toString());
     }
}
```

Parsing a Certificate Reply

The following example parses a PKCS7-formatted certificate reply stored in a file and extracts all the certificates from it:

```
try (FileInputStream fis = new FileInputStream(filename)) {
   CertificateFactory cf = CertificateFactory.getInstance("X.509");

   Collection<? extends Certificate> c = cf.generateCertificates(fis);
   for (Certificate cert : c) {
       System.out.println(cert);
   }

   // Or use the aggregate operations below for the above for-loop
   // c.stream().forEach(e -> System.out.println(e));
}
```

Using Encryption

This section takes the user through the process of generating a key, creating and initializing a cipher object, encrypting a file, and then decrypting it. Throughout this example, we use the Advanced Encryption Standard (AES).

Generating a Key

To create an AES key, we have to instantiate a `KeyGenerator` for AES. We do not specify a provider, because we do not care about a particular AES key generation implementation. Since we do not initialize the `KeyGenerator`, a system-provided source of randomness and a default keysize will be used to create the AES key:

```
KeyGenerator keygen = KeyGenerator.getInstance("AES");
keygen.init(128);
SecretKey aesKey = keygen.generateKey();
```

After the key has been generated, the same KeyGenerator object can be re-used to create further keys.

Creating a Cipher

The next step is to create a `Cipher` instance. To do this, we use one of the `getInstance` factory methods of the `Cipher` class. We must specify the name of the requested transformation, which includes the following components, separated by slashes (/):

- the algorithm name
- the mode (optional)
- the padding scheme (optional)

In this example, we create an AES cipher in Cipher Block Chaining mode, with PKCS5-style padding. We do not specify a provider, because we do not care about a particular implementation of the requested transformation.

The standard algorithm name for AES is "AES", the standard name for the Cipher Block Chaining mode is "CBC", and the standard name for PKCS5-style padding is "PKCS5Padding":

```
Cipher aesCipher;

// Create the cipher
aesCipher = Cipher.getInstance("AES/CBC/PKCS5Padding");
```

We use the generated `aesKey` from above to initialize the `Cipher` object for encryption:

```
// Initialize the cipher for encryption
aesCipher.init(Cipher.ENCRYPT_MODE, aesKey);

// Our cleartext
byte[] cleartext = "This is just an example".getBytes();

// Encrypt the cleartext
byte[] ciphertext = aesCipher.doFinal(cleartext);

// Retrieve the parameters used during encryption to properly
// initialize the cipher for decryption
AlgorithmParameters params = aesCipher.getParameters();

// Initialize the same cipher for decryption
aesCipher.init(Cipher.DECRYPT_MODE, aesKey, params);

// Decrypt the ciphertext
byte[] cleartext1 = aesCipher.doFinal(ciphertext);
```

`cleartext` and `cleartext1` are identical.

Using Password-Based Encryption

In this example, we prompt the user for a password from which we derive an encryption key.

It would seem logical to collect and store the password in an object of type `java.lang.String`. However, here's the caveat: Objects of type `String` are immutable, i.e., there are no methods defined that allow you to change (overwrite) or zero out the contents of a `String` after usage. This feature makes `String` objects unsuitable for storing security sensitive information such as user passwords. You should always collect and store security sensitive information in a char array instead. For that reason, the `javax.crypto.spec.PBEKeySpec` class takes (and returns) a password as a char array.

In order to use Password-Based Encryption (PBE) as defined in PKCS5, we have to specify a *salt* and an *iteration count*. The same salt and iteration count that are used for encryption must be used for decryption. Newer PBE algorithms use an iteration count of at least 1000.

```
    PBEKeySpec pbeKeySpec;
    PBEParameterSpec pbeParamSpec;
    SecretKeyFactory keyFac;

    // Salt
    byte[] salt = new SecureRandom().nextBytes(salt);

    // Iteration count
    int count = 1000;

    // Create PBE parameter set
    pbeParamSpec = new PBEParameterSpec(salt, count);

    // Prompt user for encryption password.
    // Collect user password as char array, and convert
    // it into a SecretKey object, using a PBE key
    // factory.
    char[] password = System.console.readPassword("Enter encryption
password: ");
    pbeKeySpec = new PBEKeySpec(password);
    keyFac = SecretKeyFactory.getInstance("PBEWithHmacSHA256AndAES_256");
    SecretKey pbeKey = keyFac.generateSecret(pbeKeySpec);

    // Create PBE Cipher
    Cipher pbeCipher = Cipher.getInstance("PBEWithHmacSHA256AndAES_256");

    // Initialize PBE Cipher with key and parameters
    pbeCipher.init(Cipher.ENCRYPT_MODE, pbeKey, pbeParamSpec);

    // Our cleartext
    byte[] cleartext = "This is another example".getBytes();

    // Encrypt the cleartext
    byte[] ciphertext = pbeCipher.doFinal(cleartext);
```

Sample Programs for Diffie-Hellman Key Exchange, AES/GCM, and HMAC-SHA256

The following are sample programs for Diffie-Hellman key exchange, AES/GCM, and HMAC-SHA256.

Topics

Diffie-Hellman Key Exchange between Two Parties

Diffie-Hellman Key Exchange between Three Parties

AES/GCM Example

HMAC-SHA256 Example

Diffie-Hellman Key Exchange between Two Parties

The program runs the Diffie-Hellman key agreement protocol between two parties.

```
/*
 * Copyright (c) 1997, 2017, Oracle and/or its affiliates. All rights
reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 *   - Redistributions of source code must retain the above copyright
 *     notice, this list of conditions and the following disclaimer.
 *
 *   - Redistributions in binary form must reproduce the above copyright
 *     notice, this list of conditions and the following disclaimer in the
 *     documentation and/or other materials provided with the distribution.
 *
 *   - Neither the name of Oracle nor the names of its
 *     contributors may be used to endorse or promote products derived
 *     from this software without specific prior written permission.
 *
 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS
 * IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO,
 * THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
 * PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE COPYRIGHT OWNER OR
 * CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
 * EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
 * PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR
 * PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
 * LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
 * NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
 * SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
 */
import java.io.*;
```

```
import java.math.BigInteger;
import java.security.*;
import java.security.spec.*;
import java.security.interfaces.*;
import javax.crypto.*;
import javax.crypto.spec.*;
import javax.crypto.interfaces.*;
import com.sun.crypto.provider.SunJCE;

public class DHKeyAgreement2 {
    private DHKeyAgreement2() {}
    public static void main(String argv[]) throws Exception {

        /*
         * Alice creates her own DH key pair with 2048-bit key size
         */
        System.out.println("ALICE: Generate DH keypair ...");
        KeyPairGenerator aliceKpairGen =
KeyPairGenerator.getInstance("DH");
        aliceKpairGen.initialize(2048);
        KeyPair aliceKpair = aliceKpairGen.generateKeyPair();

        // Alice creates and initializes her DH KeyAgreement object
        System.out.println("ALICE: Initialization ...");
        KeyAgreement aliceKeyAgree = KeyAgreement.getInstance("DH");
        aliceKeyAgree.init(aliceKpair.getPrivate());

        // Alice encodes her public key, and sends it over to Bob.
        byte[] alicePubKeyEnc = aliceKpair.getPublic().getEncoded();

        /*
         * Let's turn over to Bob. Bob has received Alice's public key
         * in encoded format.
         * He instantiates a DH public key from the encoded key material.
         */
        KeyFactory bobKeyFac = KeyFactory.getInstance("DH");
        X509EncodedKeySpec x509KeySpec = new
X509EncodedKeySpec(alicePubKeyEnc);

        PublicKey alicePubKey = bobKeyFac.generatePublic(x509KeySpec);

        /*
         * Bob gets the DH parameters associated with Alice's public key.
         * He must use the same parameters when he generates his own key
         * pair.
         */
        DHParameterSpec dhParamFromAlicePubKey =
((DHPublicKey)alicePubKey).getParams();

        // Bob creates his own DH key pair
        System.out.println("BOB: Generate DH keypair ...");
        KeyPairGenerator bobKpairGen = KeyPairGenerator.getInstance("DH");
        bobKpairGen.initialize(dhParamFromAlicePubKey);
        KeyPair bobKpair = bobKpairGen.generateKeyPair();
```

```
// Bob creates and initializes his DH KeyAgreement object
System.out.println("BOB: Initialization ...");
KeyAgreement bobKeyAgree = KeyAgreement.getInstance("DH");
bobKeyAgree.init(bobKpair.getPrivate());

// Bob encodes his public key, and sends it over to Alice.
byte[] bobPubKeyEnc = bobKpair.getPublic().getEncoded();

/*
 * Alice uses Bob's public key for the first (and only) phase
 * of her version of the DH
 * protocol.
 * Before she can do so, she has to instantiate a DH public key
 * from Bob's encoded key material.
 */
KeyFactory aliceKeyFac = KeyFactory.getInstance("DH");
x509KeySpec = new X509EncodedKeySpec(bobPubKeyEnc);
PublicKey bobPubKey = aliceKeyFac.generatePublic(x509KeySpec);
System.out.println("ALICE: Execute PHASE1 ...");
aliceKeyAgree.doPhase(bobPubKey, true);

/*
 * Bob uses Alice's public key for the first (and only) phase
 * of his version of the DH
 * protocol.
 */
System.out.println("BOB: Execute PHASE1 ...");
bobKeyAgree.doPhase(alicePubKey, true);

/*
 * At this stage, both Alice and Bob have completed the DH key
 * agreement protocol.
 * Both generate the (same) shared secret.
 */
try {
    byte[] aliceSharedSecret = aliceKeyAgree.generateSecret();
    int aliceLen = aliceSharedSecret.length;
    byte[] bobSharedSecret = new byte[aliceLen];
    int bobLen;
} catch (ShortBufferException e) {
    System.out.println(e.getMessage());
}        // provide output buffer of required size
bobLen = bobKeyAgree.generateSecret(bobSharedSecret, 0);
System.out.println("Alice secret: " +
        toHexString(aliceSharedSecret));
System.out.println("Bob secret: " +
        toHexString(bobSharedSecret));
if (!java.util.Arrays.equals(aliceSharedSecret, bobSharedSecret))
    throw new Exception("Shared secrets differ");
System.out.println("Shared secrets are the same");

/*
 * Now let's create a SecretKey object using the shared secret
 * and use it for encryption. First, we generate SecretKeys for the
 * "AES" algorithm (based on the raw shared secret data) and
```

```
         * Then we use AES in CBC mode, which requires an initialization
         * vector (IV) parameter. Note that you have to use the same IV
         * for encryption and decryption: If you use a different IV for
         * decryption than you used for encryption, decryption will fail.
         *
         * If you do not specify an IV when you initialize the Cipher
         * object for encryption, the underlying implementation will
generate
         * a random one, which you have to retrieve using the
         * javax.crypto.Cipher.getParameters() method, which returns an
         * instance of java.security.AlgorithmParameters. You need to
transfer
         * the contents of that object (e.g., in encoded format, obtained
via
         * the AlgorithmParameters.getEncoded() method) to the party who
will
         * do the decryption. When initializing the Cipher for decryption,
         * the (reinstantiated) AlgorithmParameters object must be
explicitly
         * passed to the Cipher.init() method.
         */
        System.out.println("Use shared secret as SecretKey object ...");
        SecretKeySpec bobAesKey = new SecretKeySpec(bobSharedSecret, 0,
16, "AES");
        SecretKeySpec aliceAesKey = new SecretKeySpec(aliceSharedSecret,
0, 16, "AES");

        /*
         * Bob encrypts, using AES in CBC mode
         */
        Cipher bobCipher = Cipher.getInstance("AES/CBC/PKCS5Padding");
        bobCipher.init(Cipher.ENCRYPT_MODE, bobAesKey);
        byte[] cleartext = "This is just an example".getBytes();
        byte[] ciphertext = bobCipher.doFinal(cleartext);

        // Retrieve the parameter that was used, and transfer it to Alice
in
        // encoded format
        byte[] encodedParams = bobCipher.getParameters().getEncoded();

        /*
         * Alice decrypts, using AES in CBC mode
         */

        // Instantiate AlgorithmParameters object from parameter encoding
        // obtained from Bob
        AlgorithmParameters aesParams =
AlgorithmParameters.getInstance("AES");
        aesParams.init(encodedParams);
        Cipher aliceCipher = Cipher.getInstance("AES/CBC/PKCS5Padding");
        aliceCipher.init(Cipher.DECRYPT_MODE, aliceAesKey, aesParams);
        byte[] recovered = aliceCipher.doFinal(ciphertext);
        if (!java.util.Arrays.equals(cleartext, recovered))
            throw new Exception("AES in CBC mode recovered text is " +
                    "different from cleartext");
```

```
        System.out.println("AES in CBC mode recovered text is "
                "same as cleartext");
    }

    /*
     * Converts a byte to hex digit and writes to the supplied buffer
     */
    private static void byte2hex(byte b, StringBuffer buf) {
        char[] hexChars = { '0', '1', '2', '3', '4', '5', '6', '7', '8',
                '9', 'A', 'B', 'C', 'D', 'E', 'F' };
        int high = ((b & 0xf0) >> 4);
        int low = (b & 0x0f);
        buf.append(hexChars[high]);
        buf.append(hexChars[low]);
    }

    /*
     * Converts a byte array to hex string
     */
    private static String toHexString(byte[] block) {
        StringBuffer buf = new StringBuffer();
        int len = block.length;
        for (int i = 0; i < len; i++) {
            byte2hex(block[i], buf);
            if (i < len-1) {
                buf.append(":");
            }
        }
        return buf.toString();
    }
}
```

Diffie-Hellman Key Exchange between Three Parties

The program runs the Diffie-Hellman key agreement protocol between 3 parties.

```
/*
 * Copyright (c) 1997, 2017, Oracle and/or its affiliates. All rights
reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 *   - Redistributions of source code must retain the above copyright
 *     notice, this list of conditions and the following disclaimer.
 *
 *   - Redistributions in binary form must reproduce the above copyright
 *     notice, this list of conditions and the following disclaimer in the
 *     documentation and/or other materials provided with the distribution.
 *
 *   - Neither the name of Oracle nor the names of its
 *     contributors may be used to endorse or promote products derived
```

```
 *     from this software without specific prior written permission.
 *
 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS
 * IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO,
 * THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
 * PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE COPYRIGHT OWNER OR
 * CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
 * EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
 * PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR
 * PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
 * LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
 * NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
 * SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
 */
import java.security.*;
import java.security.spec.*;
import javax.crypto.*;
import javax.crypto.spec.*;
import javax.crypto.interfaces.*;
   /*
    * This program executes the Diffie-Hellman key agreement protocol
between
    * 3 parties: Alice, Bob, and Carol using a shared 2048-bit DH
parameter.
    */
    public class DHKeyAgreement3 {
        private DHKeyAgreement3() {}
        public static void main(String argv[]) throws Exception {
        // Alice creates her own DH key pair with 2048-bit key size
            System.out.println("ALICE: Generate DH keypair ...");
            KeyPairGenerator aliceKpairGen =
KeyPairGenerator.getInstance("DH");
            aliceKpairGen.initialize(2048);
            KeyPair aliceKpair = aliceKpairGen.generateKeyPair();
        // This DH parameters can also be constructed by creating a
        // DHParameterSpec object using agreed-upon values
            DHParameterSpec dhParamShared =
((DHPublicKey)aliceKpair.getPublic()).getParams();
        // Bob creates his own DH key pair using the same params
            System.out.println("BOB: Generate DH keypair ...");
            KeyPairGenerator bobKpairGen =
KeyPairGenerator.getInstance("DH");
            bobKpairGen.initialize(dhParamShared);
            KeyPair bobKpair = bobKpairGen.generateKeyPair();
        // Carol creates her own DH key pair using the same params
            System.out.println("CAROL: Generate DH keypair ...");
            KeyPairGenerator carolKpairGen =
KeyPairGenerator.getInstance("DH");
            carolKpairGen.initialize(dhParamShared);
            KeyPair carolKpair = carolKpairGen.generateKeyPair();
        // Alice initialize
            System.out.println("ALICE: Initialize ...");
            KeyAgreement aliceKeyAgree = KeyAgreement.getInstance("DH");
            aliceKeyAgree.init(aliceKpair.getPrivate());
```

```
        // Bob initialize
            System.out.println("BOB: Initialize ...");
            KeyAgreement bobKeyAgree = KeyAgreement.getInstance("DH");
            bobKeyAgree.init(bobKpair.getPrivate());
        // Carol initialize
            System.out.println("CAROL: Initialize ...");
            KeyAgreement carolKeyAgree = KeyAgreement.getInstance("DH");
            carolKeyAgree.init(carolKpair.getPrivate());
        // Alice uses Carol's public key
            Key ac = aliceKeyAgree.doPhase(carolKpair.getPublic(), false);
        // Bob uses Alice's public key
            Key ba = bobKeyAgree.doPhase(aliceKpair.getPublic(), false);
        // Carol uses Bob's public key
            Key cb = carolKeyAgree.doPhase(bobKpair.getPublic(), false);
        // Alice uses Carol's result from above
            aliceKeyAgree.doPhase(cb, true);
        // Bob uses Alice's result from above
            bobKeyAgree.doPhase(ac, true);
        // Carol uses Bob's result from above
            carolKeyAgree.doPhase(ba, true);
        // Alice, Bob and Carol compute their secrets
            byte[] aliceSharedSecret = aliceKeyAgree.generateSecret();
            System.out.println("Alice secret: " +
toHexString(aliceSharedSecret));
            byte[] bobSharedSecret = bobKeyAgree.generateSecret();
            System.out.println("Bob secret: " +
toHexString(bobSharedSecret));
            byte[] carolSharedSecret = carolKeyAgree.generateSecret();
            System.out.println("Carol secret: " +
toHexString(carolSharedSecret));
        // Compare Alice and Bob
            if (!java.util.Arrays.equals(aliceSharedSecret,
bobSharedSecret))
                throw new Exception("Alice and Bob differ");
            System.out.println("Alice and Bob are the same");
        // Compare Bob and Carol
            if (!java.util.Arrays.equals(bobSharedSecret,
carolSharedSecret))
                throw new Exception("Bob and Carol differ");
            System.out.println("Bob and Carol are the same");
        }
    /*
     * Converts a byte to hex digit and writes to the supplied buffer
     */
        private static void byte2hex(byte b, StringBuffer buf) {
            char[] hexChars = { '0', '1', '2', '3', '4', '5', '6', '7',
'8',
                                '9', 'A', 'B', 'C', 'D', 'E', 'F' };
            int high = ((b & 0xf0) >> 4);
            int low = (b & 0x0f);
            buf.append(hexChars[high]);
            buf.append(hexChars[low]);
        }
    /*
     * Converts a byte array to hex string
```

```
    */
       private static String toHexString(byte[] block) {
           StringBuffer buf = new StringBuffer();
           int len = block.length;
           for (int i = 0; i < len; i++) {
               byte2hex(block[i], buf);
               if (i < len-1) {
                   buf.append(":");
               }
           }
           return buf.toString();
       }
   }
```

AES/GCM Example

The following is a sample program to demonstrate AES/GCM usage to encrypt/decrypt data.

```
/*
 * Copyright (c) 2017, Oracle and/or its affiliates. All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 *   - Redistributions of source code must retain the above copyright
 *     notice, this list of conditions and the following disclaimer.
 *
 *   - Redistributions in binary form must reproduce the above copyright
 *     notice, this list of conditions and the following disclaimer in the
 *     documentation and/or other materials provided with the distribution.
 *
 *   - Neither the name of Oracle nor the names of its
 *     contributors may be used to endorse or promote products derived
 *     from this software without specific prior written permission.
 *
 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS
 * IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO,
 * THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
 * PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE COPYRIGHT OWNER OR
 * CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
 * EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
 * PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR
 * PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
 * LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
 * NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
 * SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
 */

import java.security.AlgorithmParameters;
import java.util.Arrays;
import javax.crypto.*;
```

```
public class AESGCMTest {

    public static void main(String[] args) throws Exception {
        // Slightly longer than 1 AES block (128 bits) to show PADDING
        // is "handled" by GCM.
        byte[] data = {
            0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07,
            0x08, 0x09, 0x0a, 0x0b, 0x0c, 0x0d, 0x0e, 0x0f,
            0x10};

        // Create a 128-bit AES key.
        KeyGenerator kg = KeyGenerator.getInstance("AES");
        kg.init(128);
        SecretKey key = kg.generateKey();

        // Obtain a AES/GCM cipher to do the enciphering. Must obtain
        // and use the Parameters for successful decryption.
        Cipher encCipher = Cipher.getInstance("AES/GCM/NOPADDING");
        encCipher.init(Cipher.ENCRYPT_MODE, key);
        byte[] enc = encCipher.doFinal(data);
        AlgorithmParameters ap = encCipher.getParameters();

        // Obtain a similar cipher, and use the parameters.
        Cipher decCipher = Cipher.getInstance("AES/GCM/NOPADDING");
        decCipher.init(Cipher.DECRYPT_MODE, key, ap);
        byte[] dec = decCipher.doFinal(enc);

        if (Arrays.compare(data, dec) != 0) {
            throw new Exception("Original data != decrypted data");
        }
    }
}
```

HMAC-SHA256 Example

The following is a sample program that demonstrates how to generate a secret-key object for HMAC-SHA256, and initialize a HMAC-SHA256 object with it.

Example 2-14 Generate a Secret-key Object for HMAC-SHA256

```
/*
 * Copyright (c) 1997, 2017, Oracle and/or its affiliates. All rights
reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 *   - Redistributions of source code must retain the above copyright
 *     notice, this list of conditions and the following disclaimer.
 *
 *   - Redistributions in binary form must reproduce the above copyright
 *     notice, this list of conditions and the following disclaimer in the
```

```
 *     documentation and/or other materials provided with the distribution.
 *
 *   - Neither the name of Oracle nor the names of its
 *     contributors may be used to endorse or promote products derived
 *     from this software without specific prior written permission.
 *
 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS
 * IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
TO,
 * THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
 * PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE COPYRIGHT OWNER OR
 * CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
 * EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
 * PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR
 * PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
 * LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
 * NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
 * SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
 */

import java.security.*;
import javax.crypto.*;

    /**
     * This program demonstrates how to generate a secret-key object for
     * HMACSHA256, and initialize an HMACSHA256 object with it.
     */

    public class initMac {

        public static void main(String[] args) throws Exception {

            // Generate secret key for HmacSHA256
            KeyGenerator kg = KeyGenerator.getInstance("HmacSHA256");
            SecretKey sk = kg.generateKey();

            // Get instance of Mac object implementing HmacSHA256, and
            // initialize it with the above secret key
            Mac mac = Mac.getInstance("HmacSHA256");
            mac.init(sk);
            byte[] result = mac.doFinal("Hi There".getBytes());
        }
    }
```

3

How to Implement a Provider in the Java Cryptography Architecture

This document describes what you need to do in order to integrate your provider into Java SE so that algorithms and other services can be found when Java Security API clients request them.

Who Should Read This Document

Programmers that only need to use the Java Security APIs (see Core Classes and Interfaces in Java Cryptography Architecture (JCA) Reference Guide) to access existing cryptography algorithms and other services do *not* need to read this document.

This document is intended for experienced programmers wishing to create their own provider packages supplying cryptographic service implementations. It documents what you need to do in order to integrate your provider into Java so that your algorithms and other services can be found when Java Security API clients request them.

Notes on Terminology

Throughout this document, the terms *JCA* by itself refers to the JCA framework. Whenever this document notes a specific JCA provider, it will be referred to explicitly by the provider name.

- Prior to JDK 1.4, the JCE was an unbundled product, and as such, the JCA and JCE were regularly referred to as separate, distinct components. As JCE is now bundled in JDK, the distinction is becoming less apparent. Since the JCE uses the same architecture as the JCA, the JCE should be more properly thought of as a subset of the JCA.
- The JCA within the JDK includes two software components:
  - the framework that defines and supports cryptographic services for which providers supply implementations. This framework includes packages such as `java.security`, `javax.crypto`, `javax.crypto.spec`, and `javax.crypto.interfaces`.
  - the actual providers such as *Sun*, *SunRsaSign*, *SunJCE*, which contain the actual cryptographic implementations.

Introduction to Implementing Providers

The Java platform defines a set of APIs spanning major security areas, including cryptography, public key infrastructure, authentication, secure communication, and

access control. These APIs allow developers to easily integrate security into their application code. They were designed around the following principles:

- **Implementation independence**: Applications do not need to implement security themselves. Rather, they can request security services from the Java platform. Security services are implemented in providers (see below), which are plugged into the Java platform via a standard interface. An application may rely on multiple independent providers for security functionality.
- **Implementation interoperability**: Providers are interoperable across applications. Specifically, an application is not bound to a specific provider, and a provider is not bound to a specific application.
- **Algorithm extensibility**: The Java platform includes a number of built-in providers that implement a basic set of security services that are widely used today. However, some applications may rely on emerging standards not yet implemented, or on proprietary services. The Java platform supports the installation of custom providers that implement such services.

A Cryptographic Service Provider (provider) refers to a package (or a set of packages) that supply a concrete implementation of a subset of the cryptography aspects of the JDK Security API.

The `java.security.Provider` class encapsulates the notion of a security provider in the Java platform. It specifies the provider's name and lists the security services it implements. Multiple providers may be configured at the same time, and are listed in order of preference. When a security service is requested, the highest priority provider that implements that service is selected. See Security Providers, which illustrates how a provider selects a requested security service.

Engine Classes and Corresponding Service Provider Interface Classes

An engine class defines a cryptographic service in an abstract fashion (without a concrete implementation). A cryptographic service is always associated with a particular algorithm or type.

A *cryptographic service* either provides cryptographic operations (like those for digital signatures or message digests, ciphers or key agreement protocols); generates or supplies the cryptographic material (keys or parameters) required for cryptographic operations; or generates data objects (keystores or certificates) that encapsulate cryptographic keys (which can be used in a cryptographic operation) in a secure fashion.

For example, here are four engine classes:

- `Signature` class provides access to the functionality of a digital signature algorithm.
- A DSA `KeyFactory` class supplies a DSA private or public key (from its encoding or transparent specification) in a format usable by the initSign or initVerify methods, respectively, of a DSA Signature object.
- `Cipher` class provides access to the functionality of an encryption algorithm (such as AES)
- `KeyAgreement` class provides access to the functionality of a key agreement protocol (such as Diffie-Hellman)

The Java Cryptography Architecture encompasses the classes comprising the Security package that relate to cryptography, including the engine classes. Users of the API request and utilize instances of the engine classes to carry out corresponding operations. The JDK defines the following engine classes:

- `MessageDigest` - used to calculate the message digest (hash) of specified data.
- `Signature` - used to sign data and verify digital signatures.
- `KeyPairGenerator` - used to generate a pair of public and private keys suitable for a specified algorithm.
- `KeyFactory` - used to convert opaque cryptographic keys of type `Key` into *key specifications* (transparent representations of the underlying key material), and vice versa.
- `KeyStore` - used to create and manage a *keystore*. A keystore is a database of keys. Private keys in a keystore have a certificate chain associated with them, which authenticates the corresponding public key. A keystore also contains certificates from trusted entities.
- `CertificateFactory` - used to create public key certificates and Certificate Revocation Lists (CRLs).
- `AlgorithmParameters` - used to manage the parameters for a particular algorithm, including parameter encoding and decoding.
- `AlgorithmParameterGenerator` - used to generate a set of parameters suitable for a specified algorithm.
- `SecureRandom` - used to generate random or pseudo-random numbers.
- `Cipher` - used to encrypt or decrypt some specified data.
- `KeyAgreement` - used to execute a key agreement (key exchange) protocol between 2 or more parties.
- `KeyGenerator` - used to generate a secret (symmetric) key suitable for a specified algorithm.
- `Mac`: used to compute the *message authentication code* of some specified data.
- `SecretKeyFactory` - used to convert opaque cryptographic keys of type `SecretKey` into *key specifications* (transparent representations of the underlying key material), and vice versa.
- `CertPathBuilder` - used to create public key certificates and Certificate Revocation Lists (CRLs).
- `CertPathValidator` - used to validate certificate chains.
- `CertStore` - used to retrieve Certificates and CRLs from a repository.
- `ExemptionMechanism` - used to provide the functionality of an exemption mechanism such as *key recovery*, *key weakening*, *key escrow*, or any other (custom) exemption mechanism. Applications or applets that use an exemption mechanism may be granted stronger encryption capabilities than those which don't. However, please note that cryptographic restrictions are no longer required for most countries, and thus exemption mechanisms may only be useful in those few countries whose governments mandate restrictions.

> **Note:**
>
> A *generator* creates objects with brand-new contents, whereas a *factory* creates objects from existing material (for example, an encoding).

An *engine* class provides the interface to the functionality of a specific type of cryptographic service (independent of a particular cryptographic algorithm). It defines *Application Programming Interface* (API) methods that allow applications to access the specific type of cryptographic service it provides. The actual implementations (from one or more providers) are those for specific algorithms. For example, the Signature engine class provides access to the functionality of a digital signature algorithm. The actual implementation supplied in a `SignatureSpi` subclass (see next paragraph) would be that for a specific kind of signature algorithm, such as SHA256withDSA or SHA512withRSA.

The application interfaces supplied by an engine class are implemented in terms of a **Service Provider Interface (SPI)**. That is, for each engine class, there is a corresponding abstract SPI class, which defines the Service Provider Interface methods that cryptographic service providers must implement.

Figure 3-1 Engine Classes

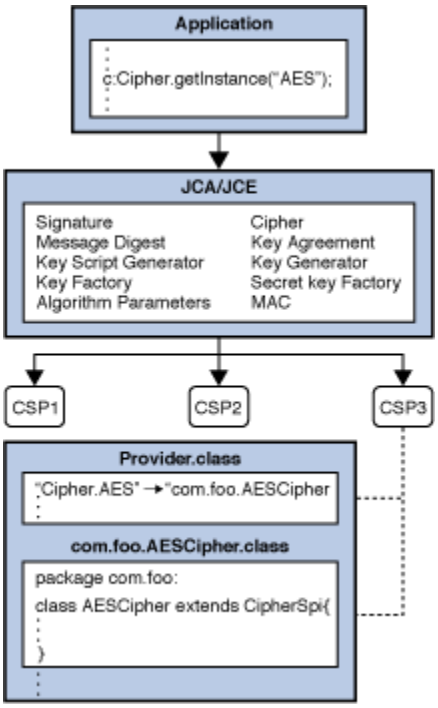

An instance of an engine class, the "API object", encapsulates (as a private field) an instance of the corresponding SPI class, the "SPI object". All API methods of an API object are declared "final", and their implementations invoke the corresponding SPI methods of the encapsulated SPI object. An instance of an engine class (and of its corresponding SPI class) is created by a call to the `getInstance` factory method of the engine class.

The name of each SPI class is the same as that of the corresponding engine class, followed by "Spi". For example, the SPI class corresponding to the Signature engine class is the `SignatureSpi` class.

Each SPI class is abstract. To supply the implementation of a particular type of service and for a specific algorithm, a provider must subclass the corresponding SPI class and provide implementations for all the abstract methods.

Another example of an engine class is the `MessageDigest` class, which provides access to a message digest algorithm. Its implementations, in `MessageDigestSpi` subclasses, may be those of various message digest algorithms such as SHA256 or SHA384.

As a final example, the `KeyFactory` engine class supports the conversion from opaque keys to transparent key specifications, and vice versa. See Key Specification Interfaces and Classes Required by Key Factories. The actual implementation supplied in a `KeyFactorySpi` subclass would be that for a specific type of keys, e.g., DSA public and private keys.

Steps to Implement and Integrate a Provider

Follow these steps to implement a provider and integrate it into the JCA framework:

- Step 1: Write your Service Implementation Code
- Step 2: Give your Provider a Name
- Step 3: Write Your Master Class, a Subclass of Provider
- Step 4: Create a Module Declaration for Your Provider
- Step 5: Compile Your Code
- Step 6: Place Your Provider in a JAR File
- Step 7: Sign Your JAR File, If Necessary
- Step 8: Prepare for Testing
- Step 9: Write and Compile Your Test Programs
- Step 10: Run Your Test Programs
- Step 11: Apply for U.S. Government Export Approval If Required
- Step 12: Document Your Provider and Its Supported Services
- Step 13: Make Your Class Files and Documentation Available to Clients

Step 1: Write your Service Implementation Code

The first thing you need to do is to write the code that provides algorithm-specific implementations of the cryptographic services you want to support. Your provider may supply implementations of cryptographic services already available in one or more of the security components of the JDK.

For cryptographic services not defined in JCA (for example, signatures and message digests), see Engine Classes and Algorithms.

For each cryptographic service you wish to implement, create a subclass of the appropriate SPI class. JCA defines the following engine classes:

- `SignatureSpi`
- `MessageDigestSpi`
- `KeyPairGeneratorSpi`
- `SecureRandomSpi`
- `AlgorithmParameterGeneratorSpi`
- `AlgorithmParametersSpi`
- `KeyFactorySpi`
- `CertificateFactorySpi`
- `KeyStoreSpi`
- `CipherSpi`
- `KeyAgreementSpi`
- `KeyGeneratorSpi`
- `MacSpi`
- `SecretKeyFactorySpi`
- `ExemptionMechanismSpi`

To know more about the JCA and other cryptographic classes, see Engine Classes and Corresponding Service Provider Interface Classes.

In the subclass, you need to:

1. Supply implementations for the abstract methods, whose names usually begin with `engine`. See Further Implementation Details and Requirements.
2. Depending on how you write your provider and register its algorithms (using either `String` objects or the `Provider.Service` class), the provider either:
   - Ensure that there is a public constructor without any arguments. Here's why: When one of your services is requested, Java Security looks up the subclass implementing that service, as specified by a property in your "master class" (see Step 3: Write Your Master Class, a Subclass of Provider). Java Security then creates the `Class` object associated with your subclass, and creates an instance of your subclass by calling the `newInstance` method on that `Class` object. `newInstance` requires your subclass to have a public constructor without any parameters. (A default constructor without arguments will automatically be generated if your subclass doesn't have any constructors. But if your subclass defines any constructors, you must explicitly define a public constructor without arguments.)
   - Override the `newInstance()` method in the registered `Provider.Service`. This is the preferred mechanism in JDK 9 and later.

Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations

When instantiating a provider's implementation (class) of a `Cipher`, `KeyAgreement`, `KeyGenerator`, `MAC`, or `SecretKey` factory, the framework will determine the provider's codebase (JAR file) and verify its signature. In this way, JCA authenticates the provider and ensures that only providers signed by a trusted entity can be plugged into

the JCA. Thus, one requirement for encryption providers is that they must be signed, as described in later steps.

In order for provider classes to become unusable if instantiated by an application directly, bypassing JCA, providers should implement the following:

- All SPI implementation classes in a provider package should be declared `final` (so that they cannot be subclassed), and their (SPI) implementation methods should be declared `protected`.
- All crypto-related helper classes in a provider package should have package-private scope, so that they cannot be accessed from outside the provider package.

For providers that may be exported outside the U.S., `CipherSpi` implementations must include an implementation of the `engineGetKeySize` method which, given a `Key`, returns the key size. If there are restrictions on available cryptographic strength specified in jurisdiction policy files, each `Cipher` initialization method calls `engineGetKeySize` and then compares the result with the maximum allowable key size for the particular location and circumstances of the applet or application being run. If the key size is too large, the initialization method throws an exception.

Additional *optional* features that providers may implement are:

- Optional: The `engineWrap` and `engineUnwrap` methods of `CipherSpi`. Wrapping a key enables secure transfer of the key from one place to another. Information about wrapping and unwrapping keys is provided in the `wrap`.
- Optional: One or more *exemption mechanisms*. An exemption mechanism is something such as key recovery, key escrow, or key weakening which, if implemented and enforced, may enable reduced cryptographic restrictions for an application (or applet) that uses it. To know more about the requirements for apps that utilize exemption mechanisms, see How to Make Applications Exempt from Cryptographic Restrictions.

Step 2: Give your Provider a Name

Decide on a unique name for your provider. This is the name to be used by client applications to refer to your provider, and it must not conflict with any other provider names.

Step 3: Write Your Master Class, a Subclass of Provider

Create a subclass of the `java.security.Provider` class. This is essentially a lookup table that advertises the algorithms that your provider implements.

You can use the following coding styles to subclass the `Provider` class:

- Create a provider that registers its services with `String` objects to store algorithm names and their associated implementation class name. These are stored in the `Hashtable<Object,Object>` superclass of `java.security.Provider`.
- Create a provider that uses the `Provider.Service` class, which uses a different method to store algorithm names and create new objects. The `Provider.Service` class enables you customize how the JCE framework requests services from your provider, such as how the framework creates new instances of your provider's services. This coding style is recommended, especially when using modules.

A provider can use either style, or even use both styles at the same time. Regardless of which style you choose, your subclass should be `final`.

Step 3.1: Create a Provider That Uses String Objects to Register Its Services

The following is an example of a provider that uses `String` objects to store implemented algorithm names:

```
package p;
public final class MyProvider extends Provider {
    public MyProvider() {
        super("MyProvider", "1.0",
            "Some info about my provider and which algorithms it
supports");
        // com.my.crypto.provider.MyCipher extends CipherSPI
        put("Cipher.MyCipher", "com.my.crypto.provider.MyCipher");
    }
}
```

To create a provider with this coding style, do the following:

- Call `super`, specifying the provider name (see Step 2: Give your Provider a Name) version number, and a string of information about the provider and algorithms it supports.

  ```
  super("MyProvider", "1.0",
       "Some info about my provider and which algorithms it supports");
  ```

- Set the values of various properties that are required for the Java Security API to look up the cryptographic services implemented by the provider.

  For each service implemented by the provider, there must be a property whose name is the type of service followed by a period and the name of the algorithm to which the service applies. The property value must specify the fully qualified name of the class implementing the service.

  For example, this following statement sets a property named `Cipher.MyCypher` whose value is `com.my.crypto.provider.MyCipher`, a class that extends `CipherSPI`:

  ```
  put("Cipher.MyCipher", "com.my.crypto.provider.MyCipher");
  ```

  The following list shows the various types of JCA services, where the actual algorithm name is substituted for *`algName`*:

  - `Signature.`*`algName`*
  - `MessageDigest.`*`algName`*
  - `KeyPairGenerator.`*`algName`*
  - `SecureRandom.`*`algName`*
  - `AlgorithmParameterGenerator.`*`algName`*
  - `AlgorithmParameters.`*`algName`*

- `KeyFactory.algName`
- `CertificateFactory.algName`
- `KeyStore.algName`
- `Cipher.algName`: *`algName`* may actually represent a *transformation*, and may be composed of an algorithm name, a particular mode, and a padding scheme. See Java Security Standard Algorithm Names
- `KeyAgreement.algName`
- `KeyGenerator.algName`
- `Mac.algName`
- `SecretKeyFactory.algName`
- `ExemptionMechanism.algName`: *`algName`* refers to the name of the exemption mechanism, which can be one of the following: `KeyRecovery`, `KeyEscrow`, or `KeyWeakening`. Case does *not* matter.

In all of these except `ExemptionMechanism` and `Cipher`, `algName` is the "standard" name of the algorithm, certificate type, or keystore type. See Java Security Standard Algorithm Names for the standard names that should be used.

The value of each property must be the fully qualified name of the class implementing the specified algorithm, certificate type, or keystore type. That is, it must be the package name followed by the class name, where the two are separated by a period.

As an example, the default provider named *SUN* implements the Digital Signature Algorithm (whose standard name is `SHA256withDSA`) in a class named `DSA` in the `sun.security.provider` package. Its subclass of `Provider` (which is the Sun class in the `sun.security.provider package`) sets the `Signature.SHA256withDSA` property to have the value `sun.security.provider.DSA` via the following:

```
put("Signature.SHA256withDSA", "sun.security.provider.DSA")
```

The list below shows more properties that can be defined for the various types of services, where the actual algorithm name is substituted for *algName*, certificate type for *certType,* keystore type for *storeType*, and attribute name for *attrName*:

- `Signature.algName` *`[one or more spaces]`* `attrName`
- `MessageDigest.algName` *`[one or more spaces]`* `attrName`
- `KeyPairGenerator.algName` *`[one or more spaces]`* `attrName`
- `SecureRandom.algName` *`[one or more spaces]`* `attrName`
- `KeyFactory.algName` *`[one or more spaces]`* `attrName`
- `CertificateFactory.certType` *`[one or more spaces]`* `attrName`
- `KeyStore.storeType` *`[one or more spaces]`* `attrName`
- `AlgorithmParameterGenerator.algName` *`[one or more spaces]`* `attrName`
- `AlgorithmParameters.algName` *`[one or more spaces]`* `attrName`
- `Cipher.algName` *`[one or more spaces]`* `attrName`
- `KeyAgreement.algName` *`[one or more spaces]`* `attrName`

- `KeyGenerator.algName` *[one or more spaces]* `attrName`
- `Mac.algName` *[one or more spaces]* `attrName`
- `SecretKeyFactory.algName` *[one or more spaces]* `attrName`
- `ExemptionMechanism.algName` *[one or more spaces]* `attrName`

In each of these, `attrName` is the "standard" name of the algorithm, certificate type, keystore type, or attribute. (See Java Security Standard Algorithm Names for the standard names that should be used.)

For a property in the above format, the value of the property must be the value for the corresponding attribute. (See Java Security Standard Algorithm Names for the definition of each standard attribute.)

For further master class property setting examples, see the JDK source code for the `sun.security.provider.Sun` and `com.sun.crypto.provider.SunJCE` classes. They show how the Sun and SunJCE providers set properties.

As an example, the default provider named SUN implements the `SHA256withDSA` Digital Signature Algorithm in software. The class `sun.security.provider.Sun` calls the method `SunEntries.putEntries`, which sets the properties for the SUN provider, including setting the property `Signature.SHA256withDSA ImplementedIn` to have the value `Software`:

```
put("Signature.SHA256withDSA ImplementedIn", "Software");
```

> **Note:**
>
> For examples of this coding style, see the source code for `sun.security.provider.Sun` and `sun.security.provider.SunEntries` classes.

Step 3.2: Create a Provider That Uses Provider.Service

The following is an example of a provider that uses a `Provider.Service` class:

```
package p;

public final class MyProvider extends Provider {

    public MyProvider() {
        super("MyProvider", "1.0",
            "Some info about my provider and which algorithms it
supports");
        putService(new ProviderService(this, "Cipher", "MyCipher",
"p.MyCipher"));
    }

    private static final class ProviderService extends Provider.Service {
        ProviderService(Provider p, String type, String algo, String cn) {
            super(p, type, algo, cn, null, null);
        }
```

```
        @Override
        public Object newInstance(Object ctrParamObj)
            throws NoSuchAlgorithmException {
            String type = getType();
            String algo = getAlgorithm();
            try {
                if (type.equals("Cipher")) {
                    if (algo.equals("MyCipher")) {
                        return new MyCipher();
                    }
                }
            } catch (Exception ex) {
                throw new NoSuchAlgorithmException(
                    "Error constructing " + type + " for "
                    + algo + " using SunMSCAPI", ex);
            }
            throw new ProviderException("No impl for " + algo + " " +
type);
        }
    }
}
```

To create a provider with this coding style, do the following:

- For each algorithm your provider supports, call `putService` with an instance of `Provider.Service`; the arguments of the `Provider.Service` constructor represent a supported algorithm.

  The following statement adds a service named `MyCipher` of type `Cipher`; the name of the class implementing this service is `p.MyCipher`. The argument of `putService` is a subclass of `Provider.Service`:

  ```
  putService(new ProviderService(this, "Cipher", "MyCipher",
  "p.MyCipher"));
  ```

  This example uses a subclass of `Provider.Service` named `ProviderService` (rather than `Provider.Service` itself) as it customizes how the JCE framework instantiates services. If you don't need to customize the behavior of `Provider.Service`, then you can call the `Provider.Service` constructor directly:

  ```
  public final class MyProvider extends Provider {
      public MyProvider() {
          super("MyProvider", "1.0",
              "Some info about my provider and which algorithms it
  supports");
          putService(new Provider.Service(
              this, "Cipher", "MyCipher", "p.MyCipher", null, null));
      }
  }
  ```

  Note that this example is essentially the same as the example described in Step 3.1: Create a Provider That Uses String Objects to Register Its Services.

- Override any method in `Provider.Service`, such as `newInstance`, to customize how the JCE framework handles the services in your provider.

  The example at the beginning of this section overrides the method `Provider.Service.newInstance`. The method returns an instance of `MyCipher` only if the requested service is `MyCipher`. If not, it throws a `NoSuchAlgorithmException` and a `ProviderException`.

  For more information about other methods you can override, see The Provider.Service Class.

> **Note:**
>
> For examples of this coding style, see the JDK source code contained in the `sun.security.mscapi` package.

Step 3.3: Specify Additional Information for Cipher Implementations

As mentioned above, in the case of a `Cipher` property, *algName* may actually represent a *transformation*. A *transformation* is a string that describes the operation (or set of operations) to be performed by a `Cipher` object on some given input. A transformation always includes the name of a cryptographic algorithm (e.g., *AES*), and may be followed by a mode and a padding scheme.

A transformation is of the form:

- *algorithm/mode/padding*, or
- *algorithm*

(In the latter case, provider-specific default values for the mode and padding scheme are used). For example, the following is a valid transformation:

```
Cipher c = Cipher.getInstance("AES/CBC/PKCS5Padding");
```

When requesting a block cipher in stream cipher mode (for example; `AES` in `CFB` or `OFB` mode), a client may optionally specify the number of bits to be processed at a time, by appending this number to the mode name as shown in the following sample transformations:

```
Cipher c1 = Cipher.getInstance("AES/CFB8/NoPadding");
Cipher c2 = Cipher.getInstance("AES/OFB32/PKCS5Padding");
```

If a number does not follow a stream cipher mode, a provider-specific default is used. (For example, the *SunJCE* provider uses a default of 128 bits.)

A provider may supply a separate class for each combination of *algorithm/mode/padding*. Alternatively, a provider may decide to provide more generic classes representing sub-transformations corresponding to *algorithm* or *algorithm/mode* or *algorithm//padding* (note the double slashes); in this case the requested mode and/or padding are set automatically by the `getInstance` methods of `Cipher`, which invoke the `engineSetMode` and `engineSetPadding` methods of the provider's subclass of `CipherSpi`.

That is, a `Cipher` property in a provider master class may have one of the formats shown in the table below.

Table 3-1 Cipher Property Format

`Cipher` Property Format	Description
`Cipher.`*algName*	A provider's subclass of `CipherSpi` implements *algName* with pluggable mode and padding
`Cipher.`*algName/mode*	A provider's subclass of `CipherSpi` implements *algName* in the specified *mode*, with pluggable padding
`Cipher.`*algName//padding*	A provider's subclass of `CipherSpi` implements *algName* with the specified *padding*, with pluggable mode
`Cipher.`*algName/mode/padding*	A provider's subclass of `CipherSpi` implements *algName* with the specified *mode* and *padding*

(See Java Security Standard Algorithm Names for the standard algorithm names, modes, and padding schemes that should be used.)

For example, a provider may supply a subclass of `CipherSpi` that implements *AES/ECB/PKCS5Padding*, one that implements *AES/CBC/PKCS5Padding*, one that implements *AES/CFB/PKCS5Padding*, and yet another one that implements *AES/OFB/PKCS5Padding*. That provider would have the following `Cipher` properties in its master class:

- `Cipher.`*`AES/ECB/PKCS5Padding`*
- `Cipher.`*`AES/CBC/PKCS5Padding`*
- `Cipher.`*`AES/CFB/PKCS5Padding`*
- `Cipher.`*`AES/OFB/PKCS5Padding`*

Another provider may implement a class for each of the above modes (i.e., one class for *ECB*, one for *CBC*, one for *CFB*, and one for *OFB*), one class for *PKCS5Padding*, and a generic *AES* class that subclasses from `CipherSpi`. That provider would have the following `Cipher` properties in its master class:

- `Cipher.`*`AES`*
- `Cipher.`*`AES`* `SupportedModes`
  - Example: `"ECB|CBC|CFB|OFB"`
- `Cipher.`*`AES`* `SupportedPaddings`
  - Example: `"NOPADDING|PKCS5Padding"`

The `getInstance` factory method of the `Cipher` engine class follows these rules in order to instantiate a provider's implementation of `CipherSpi` for a transformation of the form "*algorithm*":

1. Check if the provider has registered a subclass of `CipherSpi` for the specified "*algorithm*".
   - If the answer is YES, instantiate this class, for whose mode and padding scheme default values (as supplied by the provider) are used.
   - If the answer is NO, throw a `NoSuchAlgorithmException` exception.

2. The `getInstance` factory method of the `Cipher` engine class follows these rules in order to instantiate a provider's implementation of `CipherSpi` for a transformation of the form "*algorithm/mode/padding*":
   a. Check if the provider has registered a subclass of `CipherSpi` for the specified "*algorithm/mode/padding*" transformation.
      - If the answer is YES, instantiate it.
      - If the answer is NO, go to the next step.
   b. Check if the provider has registered a subclass of `CipherSpi` for the sub-transformation "*algorithm/mode*".
      - If the answer is YES, instantiate it, and call `engineSetPadding(`*`padding`*`)` on the new instance.
      - If the answer is NO, go to the next step.
   c. Check if the provider has registered a subclass of `CipherSpi` for the sub-transformation "*algorithm//padding*" (note the double slashes)
      - If the answer is YES, instantiate it, and call `engineSetMode(`*`mode`*`)` on the new instance.
      - If the answer is NO, go to the next step.
   d. Check if the provider has registered a subclass of `CipherSpi` for the sub-transformation "*algorithm*".
      - If the answer is YES, instantiate it, and call `engineSetMode(`*`mode`*`)` and `engineSetPadding(`*`padding`*`)` on the new instance.
      - If the answer is NO, throw a `NoSuchAlgorithmException` exception.

Step 4: Create a Module Declaration for Your Provider

This step is optional but recommended; it enables you to package your provider in a named module. A modular JDK can then locate your provider in the module path as opposed to the class path. The module system can more thoroughly check for dependencies in modules in the module path. Note that you can use named modules in a non-modular JDK; the module declaration will be ignored. Also, you can still package your providers in unnamed or automatic modules.

Create a module declaration for your provider and save it in a file named `module-info.java`. This module declaration includes the following:

- The name of your module.
- Any module upon which your provider depends.
- A `provides` directive if your module provides a service implementation.

The following example module declaration defines a module named `com.foo.MyProvider`. `p.MyProvider` is the fully qualified class name of a service implementation. Suppose that, in this example, `p.MyProvider` uses API in the package `javax.security.auth.kerberos`, which is in the module `java.security.jgss`. Thus, the directive `requires java.security.jgss` appears in the module declaration.

```
module com.foo.MyProvider {
    provides java.security.Provider with p.MyProvider;
```

```
requires java.security.jgss;
}
```

You can package a provider in three different kinds of modules:

- Named or explicit module: A module that appears on the module path and contains module configuration information in the `module-info.class` file.

  The JCE framework can use the `ServiceLoader` class (which simplifies provider configuration) to search for providers in explicit modules without any additional changes to the module. See Step 8.1: Configure the Provider and Step 10: Run Your Test Programs.
- Automatic module: A module that appears on the module path, but does not contain module configuration information in a `module-info.class` file (essentially a "regular" JAR file).
- Unnamed module: A module that appears on the class path. It may or may not have a `module-info.class` file; this file is ignored.

It is recommended that you package your providers in named modules as they provide better performance, stronger encapsulation, simpler configuration and greater flexibility.

You have a lot of flexibility when it comes to packaging and configuring your providers. However, this impacts how you start applications that use them. For example, you might have to specify additional `--add-exports` or `--add-modules` options. Named modules, in general, require fewer of these additional options. In addition named modules offer more flexibility. You can use them with non-modular JDKs or even as unnamed modules by specifying them in a modular JDK's class path. For more information about modules, see The State of the Module System and JEP 261: Module System.

Step 5: Compile Your Code

After you have created your implementation code (Step 1: Write your Service Implementation Code), given your provider a name (Step 2: Give your Provider a Name), created the master class (Step 3: Write Your Master Class, a Subclass of Provider), and created a module declaration (Step 4: Create a Module Declaration for Your Provider), use the Java compiler to compile your files.

Step 6: Place Your Provider in a JAR File

Add the File java.security.Provider to Use the ServiceLoader Class to Search for Providers

If your provider is packaged in an automatic or unnamed module (you did not create a module declaration as described in Step 4: Create a Module Declaration for Your Provider) and you want the use the `java.util.ServiceLoader` to search for your providers, then add the file `META-INF/services/java.security.Provider` to the JAR file and ensure that the file contains the fully qualified class name of your provider implementation.

The security provider loading mechanism uses the `ServiceLoader` class to search for providers before consulting the class path.

For example, if the fully qualified class name of your provider is `p.Provider` and all the compiled code of your provider is in the directory `classes`, then create a file named `classes/META-INF/services/java.security.Provider` that contains the following line:

```
p.MyProvider
```

Run the jar Command to Create a JAR File

The following command creates a JAR file named `MyProvider.jar`. All the compiled code for the module JAR file is in the directory `classes`. In addition, the module descriptor, `module-info.class`, is in the directory `classes`:

```
jar --create --file MyProvider.jar --module-version 1.0 -C classes
```

> **Note:**
>
> The `module-info.class` file and the `--module-version` option are optional. However, the `module-info.class` file is required if you want to create a modular JAR file. (A modular JAR file is a regular JAR file that has a `module-info.class` file in its top-level directory.)

See jar in *Java Platform, Standard Edition Tools Reference*.

Step 7: Sign Your JAR File, If Necessary

If your provider is supplying encryption algorithms through the `Cipher`, `KeyAgreement`, `KeyGenerator`, `Mac`, or `SecretKeyFactory` classes, you must sign your JAR file so that the JCA can authenticate the code at run time; see Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations. If you are not providing an implementation of this type, then you can skip this step.

Step 7.1: Get a Code-Signing Certificate

The next step is to request a code-signing certificate so that you can use it to sign your provider prior to testing. The certificate will be good for both testing and production. It will be valid for 5 years.

Below are the steps you should use to get a code-signing certificate. See keytool in the *Java Platform, Standard Edition Tools Reference*.

1. Use **keytool** to generate a RSA keypair, using RSA algorithm as an example:

```
keytool -genkeypair -alias <alias> \
        -keyalg RSA -keysize 2048 \
        -dname "cn=<Company Name>, \
        ou=Java Software Code Signing, \
        o=Oracle Corporation" \
```

```
    -keystore <keystore file name> \
    -storepass <keystore password>
```

This will generate a DSA keypair (a public key and an associated private key) and store it in an entry in the specified keystore. The public key is stored in a self-signed certificate. The keystore entry can subsequently be accessed using the specified alias.

The option values in angle brackets ("<" and ">") represent the actual values that must be supplied. For example, `<alias>` must be replaced with whatever alias name you wish to be used to refer to the newly-generated keystore entry in the future, and `<keystore file name>` must be replaced with the name of the keystore to be used.

> **Tip:**
>
> Do not surround actual values with angle brackets. For example, if you want your alias to be `myTestAlias`, specify the `-alias` option as follows:
>
> ```
> -alias myTestAlias
> ```

If you specify a keystore that doesn't yet exist, it will be created.

> **Note:**
>
> If command lines you type are not allowed to be as long as the `keytool -genkeypair` command you want to execute (for example, if you are typing to a Microsoft Windows DOS prompt), you can create and execute a plain-text batch file containing the command. That is, create a new text file that contains nothing but the full `keytool -genkeypair` command. (Remember to type it all on one line.) Save the file with a .bat extension. Then in your DOS window, type the file name (with its path, if necessary). This will cause the command in the batch file to be executed.

2. Use **keytool** to generate a certificate signing request.

```
keytool -certreq -alias <alias> \
    -file <csr file name> \
    -keystore <keystore file name> \
    -storepass <keystore password>
```

Here, *<alias>* is the alias for the DSA keypair entry created in the previous step. This command generates a Certificate Signing Request (CSR), using the PKCS#10 format. It stores the CSR in the file whose name is specified in *<csr file name>*.

3. Send the CSR, contact information, and other required documentation to the JCA Code Signing Certification Authority. See JCA Code Signing Certification Authority for contact information.
4. After the JCA Code Signing Certification Authority has received your email message, they will send you a request number via email. Once you receive this request number, you should print, fill out and send the Certification Form for CSPs. See Sending Certification Form for CSPs for contact information.
5. Use **keytool** to import the certificates received from the CA.

   Once you have received the two certificates from the JCA Code Signing Certification Authority, you can use **keytool** to import them into your keystore. First import the CA's certificate as a "trusted certificate":

   ```
   keytool -import -alias <alias for the CA cert> \
       -file <CA cert file name> \
       -keystore <keystore file name> \
       -storepass <keystore password>
   ```

   Then import the code-signing certificate:

   ```
   keytool -import -alias <alias> \
       -file <code-signing cert file name> \
       -keystore <keystore file name> \
       -storepass <keystore password>
   ```

   <alias> is the same alias as that which you created in Step 1 where you generated a DSA keypair. This command replaces the self-signed certificate in the keystore entry specified by *<alias>* with the one signed by the JCA Code Signing Certification Authority.

Now that you have in your keystore a certificate from an entity trusted by JCA (the JCA Code Signing Certification Authority), you can place your provider code in a JAR file (Step 6: Place Your Provider in a JAR File) and then use that certificate to sign the JAR file (Step 7.2: Sign Your Provider).

Step 7.2: Sign Your Provider

Sign the JAR file created in Step 6: Place Your Provider in a JAR File with the code-signing certificate obtained in Step 7.1: Get a Code-Signing Certificate. See jarsigner in *Java Platform, Standard Edition Tools Reference*.

```
jarsigner -keystore <keystore file name> \
    -storepass <keystore password> \
    <JAR file name> <alias>
```

Here, *<alias>* is the alias into the keystore for the entry containing the code-signing certificate received from the JCA Code Signing Certification Authority (the same alias as that specified in the commands in Step 7.1: Get a Code-Signing Certificate).

You can test verification of the signature via the following:

```
jarsigner -verify <JAR file name>
```

The text "jar verified" will be displayed if the verification was successful.

> **Note:**
>
> - You can also use the `jdk.security.jarsigner` API to sign JAR files.
> - If you include a signed JCE provider with your application and also want the JAR file signed for implementing other code-signing policies, you need to apply multiple signatures to the JCE provider JAR using the appropriate certificates/keys. The JCE signature is for acceptance of the provider JAR by the JCE framework, the other signature(s) can be used for making policy decisions. See jarsigner in *Java Platform, Standard Edition Tools Reference* for applying multiple signatures to a JAR file.
> - You cannot package signed providers in JMOD files.
> - Only providers that supply instances of Cipher, KeyAgreement, KeyGenerator, Mac, or SecretKFactory must be signed. If your provider only supplies instances of SecureRandom, MessageDigest, Signature, KeyStore, etc., the provider does not need to be signed.
> - You can link a provider in a custom runtime image with the `jlink` command as long as it doesn't have a Cipher, KeyAgreement, KeyGenerator, or MAC implementation.

Step 8: Prepare for Testing

The next steps describe how to install and configure your new provider so that it is available via the JCA.

Step 8.1: Configure the Provider

Register your provider so that the JCE framework can find your provider, either with the `ServiceLoader` class or in the class path or module path.

1. Open the `java.security` file in an editor:
   - Solaris, Linux, or macOS: `<java-home>/conf/security/java.security`
   - Windows: `<java-home>\conf\security\java.security`
2. In the `java.security` file, find the section where standard providers such as SUN, SunRsaSign, and SunJCE are configured as static providers; it looks like the following:

```
security.provider.1=SUN
security.provider.2=SunRsaSign
security.provider.3=SunEC
security.provider.4=SunJSSE
security.provider.5=SunJCE
security.provider.6=SunJGSS
security.provider.7=SunSASL
security.provider.8=XMLDSig
security.provider.9=SunPCSC
```

```
security.provider.10=JdkLDAP
security.provider.11=JdkSASL
security.provider.12=SunMSCAPI
security.provider.13=SunPKCS11
```

Each line in this section has the following form:

```
security.provider.n=provName|className
```

This declares a provider, and specifies its preference order *n*. The preference order is the order in which providers are searched for requested algorithms when no specific provider is requested. The order is 1-based; 1 is the most preferred, followed by 2, and so on.

provName is the provider's name and *className* is the fully qualified class name of the provider. You can use either of these two names.

3. Register your provider by adding to the `java.security` file a line with the form `security.provider.n=provName|className`.

   If you configured your provider so that the `ServiceLoader` class can search for it (because you packaged the provider in a named module as described in Step 4: Create a Module Declaration for Your Provider or added a `java.security.Provider` file as described in Add the File java.security.Provider to Use the ServiceLoader Class to Search for Providers), then specify just the provider's name.

   If you have not configured your provider so that `ServiceLoader` class can search for it, which means that the JCE framework will search for it in the class path or module path, then specify the fully qualified class name of your provider.

   For example, the highlighted line registers the provider `MyProvider` (whose fully qualified class name is `p.MyProvider` and has been configured so that the `ServiceLoader` class can search for it) as the 14th preferred provider:

   ```
   # ...
   security.provider.11=JdkSASL
   security.provider.12=SunMSCAPI
   security.provider.13=SunPKCS11
   security.provider.14=MyProvider
   ```

   If you are not sure if the `ServiceLoader` mechanism will be used, or if you'll be deploying on a non-modular system, then you can also register the provider again, this time using the full class name:

   ```
   security.provider.15=p.MyProvider
   ```

Alternatively, you can register providers dynamically. To do so, a program (such as your test program, to be written in Step 9: Write and Compile Your Test Programs) call either the `addProvider` or `insertProviderAt` method in the `Security` class:

```
ServiceLoader<Provider> sl =
ServiceLoader.load(java.security.Provider.class);
for (Provider p : sl) {
    System.out.println(p);
```

```
    if (p.getName().equals("MyProvider")) {
        Security.addProvider(p);
    }
}
```

This type of registration is not persistent and can only be done by code which is granted the following permission:

```
java.security.SecurityPermission "insertProvider.<provider name>"
```

For example, if the provider name is MyJCE, and if the provider's code is in the `myjce_provider.jar` file in the `/localWork` directory, then the following is a sample policy file that contains a `grant` statement that grants that permission:

```
grant codeBase "file:/localWork/myjce_provider.jar" {
    permission java.security.SecurityPermission
        "insertProvider.MyJCE";
};
```

Step 8.2: Set Provider Permissions

Permissions must be granted for when applications are run while a security manager is installed. A security manager may be installed for an application either through code in the application itself or through a command-line argument.

1. Your provider may need the following permissions granted to it in the client environment:
   - `java.lang.RuntimePermission` to get class protection domains. The provider may need to get its own protection domain in the process of doing self-integrity checking.
   - `java.security.SecurityPermission` to set provider properties.
2. To ensure your provider works when a security manager is installed, you need to test such an installation and execution environment. In addition, prior to testing your need to grant appropriate permissions to your provider and to any other providers it uses.

   For example, a sample statement granting permissions to a provider whose name is MyJCE and whose code is in `myjce_provider.jar` appears below. Such a statement could appear in a policy file. In this example, the `myjce_provider.jar` file is assumed to be in the `/localWork` directory.

```
grant codeBase "file:/localWork/myjce_provider.jar" {
    permission java.lang.RuntimePermission "getProtectionDomain";
    permission java.security.SecurityPermission
        "putProviderProperty.MyJCE";
};
```

Step 9: Write and Compile Your Test Programs

Write and compile one or more test programs that test your provider's incorporation into the Security API as well as the correctness of its algorithm(s). Create any supporting files needed, such as those for test data to be encrypted.

1. The first tests your program should perform are ones to ensure that your provider is found, and that its name, version number, and additional information is as expected.

   To do so, you could write code like the following, substituting your provider name for `MyPro`:

   ```
   import java.security.*;

   Provider p = Security.getProvider("MyPro");

   System.out.println("MyPro provider name is " + p.getName());
   System.out.println("MyPro provider version # is " + p.getVersion());
   System.out.println("MyPro provider info is " + p.getInfo());
   ```

2. You should ensure that your services are found.

   For instance, if you implemented the AES encryption algorithm, you could check to ensure it's found when requested by using the following code (again substituting your provider name for "MyPro"):

   ```
   Cipher c = Cipher.getInstance("AES", "MyPro");

   System.out.println("My Cipher algorithm name is " +
c.getAlgorithm());
   ```

3. Optional: If you don't specify a provider name in the call to `getInstance`, all registered providers will be searched, in preference order (see Step 8.1: Configure the Provider), until one implementing the algorithm is found.

4. Optional: If your provider implements an exemption mechanism, you should write a test applet or application that uses the exemption mechanism. Such an applet/application also needs to be signed, and needs to have a "permission policy file" bundled with it.

   See How to Make Applications Exempt from Cryptographic Restrictions for complete information on creating and testing such an application.

Step 10: Run Your Test Programs

When you run your test applications, the required `java` command options will vary depending on factors such as whether you packaged your provider as a named, automatic, or unnamed module and if you configured it so that the `ServiceLoader` class can search for it.

If you packaged your provider as a named module and have configured it so that the `ServiceLoader` class can search for it (by registering it with its name in the

`java.security` as described in Step 8.1: Configure the Provider), then run your test program with the following command:

```
java --module-path "jars" <other java options>
```

The directory `jars` contains your provider.

You may require more options depending on your provider code style (see Step 3.1: Create a Provider That Uses String Objects to Register Its Services and Step 3.2: Create a Provider That Uses Provider.Service), if you packaged your provider in a different kind of module, or if you have not configured it for the `ServiceLoader` class. The following table describes these options.

For the `java` commands, the name of the provider is `MyProvider`, its fully qualified class name is `p.MyProvider`, and it is packaged in the file `com.foo.MyProvider.jar`, which is in the directory `jars`.

Table 3-2 Expected Java Runtime Options for Various Provider Implementation Styles

Module Type	Provider Code Style	Configured for ServiceLoader Class?	Provider Name Used in java.security File	java Command
Unnamed	`String` objects or `Provider.Service`	No	Fully qualified class name	`java -cp "jars/com.foo.MyProvider.jar" <other java options>`
Unnamed	`String` objects or `Provider.Service`	Yes	Fully qualified class name or provider name	`java -cp "jars/com.foo.MyProvider.jar" <other java options>`
Automatic	`String` objects or `Provider.Service`	No	Fully qualified class name	`java --module-path "jars/com.foo.MyProvider.jar" --add-modules=com.foo.MyProvider <other java options>`
Automatic	`String` objects or `Provider.Service`	Yes	Fully qualified class name or provider name	`java --module-path "jars/com.foo.MyProvider.jar" <other java options>`
Named	`String` objects or `Provider.Service`	No	Fully qualified class name	`java --module-path "jars" --add-modules=com.foo.MyProvider --add-exports=com.foo.MyProvider/p=java.base <other java options>` You can remove the `--add-exports` option if you add `exports p` in the module declaration.
Named	`String` objects	Yes	Fully qualified class name	`java --module-path "jars" --add-exports=com.foo.MyProvider/p=java.base <other java options>` You can remove the `--add-exports` option if you add `exports p` in the module declaration.

Table 3-2 (Cont.) Expected Java Runtime Options for Various Provider Implementation Styles

Module Type	Provider Code Style	Configured for ServiceLoader Class?	Provider Name Used in java.security File	java Command
Named	`String` objects	Yes	Provider name	`java --module-path "jars" --add-exports=com.foo.MyProvider/p=java.base` *`<other java options>`* You can remove the `--add-exports` option if you add `exports p` in the module declaration.
Named	`Provider.Service`	Yes	Fully qualified class name	`java --module-path "jars" --add-exports=com.foo.MyProvider/p=java.base`*`<other java options>`* You can remove the `--add-exports` option if you add `exports p` in the module declaration.
Named	`Provider.Service`	Yes	Provider name	`java --module-path "jars"` *`<other java options>`*

Once you have determined the proper `java` options for your test programs, run them. Debug your code and continue testing as needed. If the Java runtime cannot seem to find one of your algorithms, review the previous steps and ensure that they are all completed.

Be sure to include testing of your programs using different installation options (for example, configured to use the `ServiceLoader` class or to be found in the class path or module path) and execution environments (with or without a security manager running).

1. Optional: If you find during testing that your code needs modification, make the changes and recompile Step 5: Compile Your Code.
2. Place the updated provider code in a JAR file (Step 6: Place Your Provider in a JAR File).
3. Sign the JAR file (Step 7: Sign Your JAR File, If Necessary).
4. Re-configure the provider (Step 8.1: Configure the Provider).
5. Optional: If needed, fix or add to the permissions (Step 8.2: Set Provider Permissions).
6. Run your programs.
7. Optional: If required, repeat steps 1 to 6.

Step 11: Apply for U.S. Government Export Approval If Required

All U.S. vendors whose providers may be exported outside the U.S. should apply to the Bureau of Industry and Security in the U.S. Department of Commerce for export approval.

Please consult your export counsel for more information.

> **Note:**
>
> If your provider calls `Cipher.getInstance()` and the returned `Cipher` object needs to perform strong cryptography regardless of what cryptographic strength is allowed by the user's downloaded jurisdiction policy files, you should include a copy of the `cryptoPerms` permission policy file which you intend to bundle in the JAR file for your provider and which specifies an appropriate permission for the required cryptographic strength. The necessity for this file is just like the requirement that applets and applications "exempt" from cryptographic restrictions must include a `cryptoPerms` permission policy file in their JAR file. See How to Make Applications Exempt from Cryptographic Restrictions.

Here are two URLs that may be useful:

- US Department of Commerce
- Bureau of Industry and Security

Step 12: Document Your Provider and Its Supported Services

The next step is to write documentation for your clients. At the minimum, you need to specify:

- The name programs should use to refer to your provider.

  > **Note:**
  >
  > As of this writing, provider name searches are **case-sensitive**. That is, if your master class specifies your provider name as "CryptoX" but a user requests "CRYPTOx", your provider will not be found. This behavior may change in the future, but for now be sure to warn your clients to use the exact case you specify.

- The types of algorithms and other services implemented by your provider.
- Instructions for installing the provider, similar to those provided in Step 8.1: Configure the Provider, except that the information and examples should be specific to your provider.
- The permissions your provider will require if a security manager is run, as described in Step 8.2: Set Provider Permissions.

In addition, your documentation should specify anything else of interest to clients, such as any default algorithm parameters.

Step 12.1: Indicate Whether Your Implementation is Cloneable for Message Digests and MACs

For each Message Digest and MAC algorithm, indicate whether or not your implementation is cloneable. This is not technically necessary, but it may save clients some time and coding by telling them whether or not intermediate Message Digests or MACs may be possible through cloning.

Clients who do not know whether or not a `MessageDigest` or `Mac` implementation is cloneable can find out by attempting to clone the object and catching the potential exception, as illustrated by the following example:

```
try {
    // try and clone it
    /* compute the MAC for i1 */
    mac.update(i1);
    byte[] i1Mac = mac.clone().doFinal();

    /* compute the MAC for i1 and i2 */
    mac.update(i2);
    byte[] i12Mac = mac.clone().doFinal();

    /* compute the MAC for i1, i2 and i3 */
    mac.update(i3);
    byte[] i123Mac = mac.doFinal();
} catch (CloneNotSupportedException cnse) {
    // have to use an approach not involving cloning
}
```

Where,

`mac`
Indicates the MAC object they received when they requested one via a call to `Mac.getInstance`

`i1`, `i2` and `i3`
Indicates input byte arrays, and they want to calculate separate hashes for:

- `i1`
- `i1 and i2`
- `i1, i2, and i3`

Key Pair Generators

For a key pair generator algorithm, in case the client does not explicitly initialize the key pair generator (via a call to an `initialize` method), each provider must supply and document a default initialization.

For example, the Diffie-Hellman key pair generator supplied by the *SunJCE* provider uses a default prime modulus size (`keysize`) of 2048 bits.

Key Factories

A provider should document all the key specifications supported by its (secret-)key factory.

Algorithm Parameter Generators

In case the client does not explicitly initialize the algorithm parameter generator (via a call to an `init` method in the `AlgorithmParameterGenerator` engine class), each provider must supply and document a default initialization.

For example, the *SunJCE* provider uses a default prime modulus size (`keysize`) of 2048 bits for the generation of Diffie-Hellman parameters, the *Sun* provider a default modulus prime size of 2048 bits for the generation of DSA parameters.

Signature Algorithms

If you implement a signature algorithm, you should document the format in which the signature (generated by one of the `sign` methods) is encoded.

For example, the SHA256withDSA signature algorithm supplied by the "SUN" provider encodes the signature as a standard `ASN.1 SEQUENCE` of two integers, `r` and `s`.

Random Number Generation (SecureRandom) Algorithms

For a random number generation algorithm, provide information regarding how "random" the numbers generated are, and the quality of the seed when the random number generator is self-seeding. Also note what happens when a `SecureRandom` object (and its encapsulated `SecureRandomSpi` implementation object) is deserialized: If subsequent calls to the `nextBytes` method (which invokes the `engineNextBytes` method of the encapsulated `SecureRandomSpi` object) of the restored object yield the exact same (random) bytes as the original object would, then let users know that if this behavior is undesirable, they should seed the restored random object by calling its `setSeed` method.

Certificate Factories

A provider should document what types of certificates (and their version numbers, if relevant), can be created by the factory.

Keystores

A provider should document any relevant information regarding the keystore implementation, such as its underlying data format.

Step 13: Make Your Class Files and Documentation Available to Clients

After writing, configuring, testing, installing and documenting your provider software, make documentation available to your customers.

Further Implementation Details and Requirements

This section provides additional information about alias names, service interdependencies, algorithm parameter generators and algorithm parameters.

Alias Names

In the JDK, the aliasing scheme enables clients to use aliases when referring to algorithms or types, rather than the standard names.

For many cryptographic algorithms and types, there is a single official "standard name" defined in the Java Security Standard Algorithm Names.

For example, "SHA-256" is the standard name for the SHA-256 Message Digest algorithm defined in RFC 1321. `DiffieHellman` is the standard for the Diffie-Hellman key agreement algorithm defined in PKCS3.

In the JDK, there is an aliasing scheme that enables clients to use aliases when referring to algorithms or types, rather than their standard names.

For example, the "SUN" provider's master class (`Sun.java`) defines the alias `"SHA1/DSA"` for the algorithm whose standard name is `"SHA1withDSA"`. Thus, the following statements are equivalent:

```
Signature sig = Signature.getInstance("SHA1withDSA", "SUN");

Signature sig = Signature.getInstance("SHA1/DSA", "SUN");
```

Aliases can be defined in your "master class" (see Step 3: Write Your Master Class, a Subclass of Provider). To define an alias, create a property named

```
Alg.Alias.engineClassName.aliasName
```

where *engineClassName* is the name of an engine class (e.g., `Signature`), and *aliasName* is your alias name. The *value* of the property must be the standard algorithm (or type) name for the algorithm (or type) being aliased.

As an example, the "SUN" provider defines the alias `"SHA1/DSA"` for the signature algorithm whose standard name is `"SHA1withDSA"` by setting a property named `Alg.Alias.Signature.SHA1/DSA` to have the value `SHA1withDSA` via the following:

```
put("Alg.Alias.Signature.SHA1/DSA", "SHA1withDSA");
```

> **Note:**
>
> The aliases defined by one provider are available only to that provider and not to any other providers. Thus, aliases defined by the *SunJCE* provider are available only to the *SunJCE* provider.

Service Interdependencies

Some algorithms require the use of other types of algorithms. For example, a PBE algorithm usually needs to use a message digest algorithm in order to transform a password into a key.

If you are implementing one type of algorithm that requires another, you can do one of the following:

- Provide your own implementations for both.
- Let your implementation of one algorithm use an instance of the other type of algorithm, as supplied by the default *Sun* provider that is included with every Java SE Platform installation. For example, if you are implementing a PBE algorithm

that requires a message digest algorithm, you can obtain an instance of a class implementing the SHA256 message digest algorithm by calling:

```
MessageDigest.getInstance("SHA256", "SUN")
```

- Let your implementation of one algorithm use an instance of the other type of algorithm, as supplied by another specific provider. This is only appropriate if you are sure that all clients who will use your provider will also have the other provider installed.
- Let your implementation of one algorithm use an instance of the other type of algorithm, as supplied by another (unspecified) provider. That is, you can request an algorithm by name, but without specifying any particular provider, as in:

  ```
  MessageDigest.getInstance("SHA256")
  ```

  This is only appropriate if you are sure that there will be at least one implementation of the requested algorithm (in this case, SHA256) installed on each Java platform where your provider will be used.

Here are some common types of algorithm interdependencies:

Signature and Message Digest Algorithms

A signature algorithm often requires use of a message digest algorithm. For example, the SHA256withDSA signature algorithm requires the SHA256 message digest algorithm.

Signature and (Pseudo-)Random Number Generation Algorithms

A signature algorithm often requires use of a (pseudo-)random number generation algorithm. For example, such an algorithm is required in order to generate a DSA signature.

Key Pair Generation and Message Digest Algorithms

A key pair generation algorithm often requires use of a message digest algorithm. For example, DSA keys are generated using the SHA-256 message digest algorithm.

Algorithm Parameter Generation and Message Digest Algorithms

An algorithm parameter generator often requires use of a message digest algorithm. For example, DSA parameters are generated using the SHA-256 message digest algorithm.

Keystores and Message Digest Algorithms

A keystore implementation will often utilize a message digest algorithm to compute keyed hashes (where the key is a user-provided password) to check the integrity of a keystore and make sure that the keystore has not been tampered with.

Key Pair Generation Algorithms and Algorithm Parameter Generators

A key pair generation algorithm sometimes needs to generate a new set of algorithm parameters. It can either generate the parameters directly, or use an algorithm parameter generator.

Key Pair Generation, Algorithm Parameter Generation, and (Pseudo-)Random Number Generation Algorithms

A key pair generation algorithm may require a source of randomness in order to generate a new key pair and possibly a new set of parameters associated with the keys. That source of randomness is represented by a `SecureRandom` object. The implementation of the key pair generation algorithm may generate the key parameters itself, or may use an algorithm parameter generator to generate them, in which case it may or may not initialize the algorithm parameter generator with a source of randomness.

Algorithm Parameter Generators and Algorithm Parameters

An algorithm parameter generator's `engineGenerateParameters` method must return an `AlgorithmParameters` instance.

Signature and Key Pair Generation Algorithms or Key Factories

If you are implementing a signature algorithm, your implementation's `engineInitSign` and `engineInitVerify` methods will require passed-in keys that are valid for the underlying algorithm (e.g., DSA keys for the DSS algorithm). You can do one of the following:

- Also create your own classes implementing appropriate interfaces (e.g. classes implementing the `DSAPrivateKey` and `DSAPublicKey` interfaces from the package `java.security.interfaces`), and create your own key pair generator and/or key factory returning keys of those types. Require the keys passed to `engineInitSign` and `engineInitVerify` to be the types of keys you have implemented, that is, keys generated from your key pair generator or key factory. Or you can,
- Accept keys from other key pair generators or other key factories, as long as they are instances of appropriate interfaces that enable your signature implementation to obtain the information it needs (such as the private and public keys and the key parameters). For example, the `engineInitSign` method for a DSS `Signature` class could accept any private keys that are instances of `java.security.interfaces.DSAPrivateKey`.

Keystores and Key and Certificate Factories

A keystore implementation will often utilize a key factory to parse the keys stored in the keystore, and a certificate factory to parse the certificates stored in the keystore.

Default Initialization

In case the client does not explicitly initialize a key pair generator or an algorithm parameter generator, each provider of such a service must supply (and document) a default initialization.

Sun

Default Key Pair Generator Parameter Requirements

If you implement a key pair generator, your implementation should supply default parameters that are used when clients don't specify parameters.

The documentation you supply (Step 12: Document Your Provider and Its Supported Services) should state what the default parameters are.

For example, the DSA key pair generator in the *Sun* provider supplies a set of pre-computed `p`, `q`, and `g` default values for the generation of 512, 768, 1024, and 2048-bit key pairs. The following `p`, `q`, and `g` values are used as the default values for the generation of 1024-bit DSA key pairs:

```
p = fd7f5381 1d751229 52df4a9c 2eece4e7 f611b752 3cef4400 c31e3f80
    b6512669 455d4022 51fb593d 8d58fabf c5f5ba30 f6cb9b55 6cd7813b
    801d346f f26660b7 6b9950a5 a49f9fe8 047b1022 c24fbba9 d7feb7c6
    1bf83b57 e7c6a8a6 150f04fb 83f6d3c5 1ec30235 54135a16 9132f675
    f3ae2b61 d72aeff2 2203199d d14801c7

q = 9760508f 15230bcc b292b982 a2eb840b f0581cf5

g = f7e1a085 d69b3dde cbbcab5c 36b857b9 7994afbb fa3aea82 f9574c0b
    3d078267 5159578e bad4594f e6710710 8180b449 167123e8 4c281613
    b7cf0932 8cc8a6e1 3c167a8b 547c8d28 e0a3ae1e 2bb3a675 916ea37f
    0bfa2135 62f1fb62 7a01243b cca4f1be a8519089 a883dfe1 5ae59f06
    928b665e 807b5525 64014c3b fecf492a
```

(The `p` and `q` values given here were generated by the prime generation standard, using the 160-bit

```
SEED:  8d515589 4229d5e6 89ee01e6 018a237e 2cae64cd
```

With this seed, the algorithm found `p` and `q` when the counter was at 92.)

The Provider.Service Class

`Provider.Service` class offers an alternative way for providers to advertise their services and supports additional features.

Since its introduction, security providers have published their service information via appropriately formatted key-value String pairs they put in their Hashtable entries. While this mechanism is simple and convenient, it limits the amount customization possible. As a result, JDK 5.0 introduced a second option, the `Provider.Service` class. It offers an alternative way for providers to advertise their services and supports additional features as described below. Note that this addition is fully compatible with the older method of using String valued Hashtable entries. A provider on JDK 5.0 can choose either method as it prefers, or even use both at the same time.

A `Provider.Service` object encapsulates all information about a service. This is the provider that offers the service, its type (e.g. `MessageDigest` or `Signature`), the algorithm name, and the name of the class that implements the service. Optionally, it also includes a list of alternate algorithm names for this service (aliases) and attributes, which are a map of (name, value) String pairs. In addition, it defines the methods `newInstance()` and `supportsParameter()`. They have default implementations, but can be overridden by providers if needed, as may be the case with providers that interface with hardware security tokens.

The `newInstance()` method is used by the security framework when it needs to construct new implementation instances. The default implementation uses reflection to invoke the standard constructor for the respective type of service. For all standard services except `CertStore`, this is the no-args constructor. The `constructorParameter` to `newInstance()` must be null in theses cases. For services of type `CertStore`, the constructor that takes a `CertStoreParameters` object is invoked, and `constructorParameter` must be a non-null instance of `CertStoreParameters.` A security provider can override the `newInstance()` method to implement instantiation as appropriate for that implementation. It could use direct invocation or call a constructor that passes additional information specific to the Provider instance or token. For example, if multiple Smartcard readers are present on the system, it might pass information about which reader the newly created service is to be associated with. However, despite customization all implementations must follow the conventions about `constructorParameter` described above.

The `supportsParameter()` tests whether the Service can use the specified parameter. It returns false if this service cannot use the parameter. It returns true if this service can use the parameter, if a fast test is infeasible, or if the status is unknown. It is used by the security framework with some types of services to quickly exclude non-matching implementations from consideration. It is currently only defined for the following standard services: `Signature`, `Cipher`, `Mac`, and `KeyAgreement`. The `parameter` must be an instance of `Key` in these cases. For example, for `Signature` services, the framework tests whether the service can use the supplied Key before instantiating the service. The default implementation examines the attributes `SupportedKeyFormats` and `SupportedKeyClasses` as described below. Again, a provider may override this methods to implement additional tests.

The `SupportedKeyFormats` attribute is a list of the supported formats for encoded keys (as returned by `key.getFormat()`) separated by the "|" (pipe) character. For example, `X.509|PKCS#8`. The `SupportedKeyClasses` attribute is a list of the names of classes of interfaces separated by the "|" character. A key object is considered to be acceptable if it is assignable to at least one of those classes or interfaces named. In other words, if the class of the key object is a subclass of one of the listed classes (or the class itself) or if it implements the listed interface. An example value is `"java.security.interfaces.RSAPrivateKey|java.security.interfaces.RSAPublicKey"` .

Four methods have been added to the Provider class for adding and looking up Services. As mentioned earlier, the implementation of those methods and also of the existing Properties methods have been specifically designed to ensure compatibility with existing Provider subclasses. This is achieved as follows:

If legacy Properties methods are used to add entries, the Provider class makes sure that the property strings are parsed into equivalent Service objects prior to lookup via `getService()`. Similarly, if the `putService()` method is used, equivalent property strings are placed into the provider's hashtable at the same time. If a provider implementation overrides any of the methods in the Provider class, it has to ensure that its implementation does not interfere with this conversion. To avoid problems, we recommend that implementations do not override any of the methods in the `Provider` class.

Signature Formats

The signature algorithm should specify the format in which the signature is encoded.

If you implement a signature algorithm, the documentation you supply (Step 12: Document Your Provider and Its Supported Services) should specify the format in which the signature (generated by one of the `sign` methods) is encoded.

For example, the *SHA1withDSA* signature algorithm supplied by the *Sun* provider encodes the signature as a standard ASN.1 sequence of two `ASN.1 INTEGER` values: `r` and `s`, in that order:

```
SEQUENCE ::= {
        r INTEGER,
        s INTEGER }
```

DSA Interfaces and their Required Implementations

The Java Security API contains interfaces (in the `java.security.interfaces` package) for the convenience of programmers implementing DSA services.

The Java Security API contains the following interfaces:

- `Interface DSAKey`
- `Interface DSAKeyPairGenerator`
- `Interface DSAParams`
- `Interface DSAPrivateKey`
- `Interface DSAPublicKey`

The following sections discuss requirements for implementations of these interfaces.

DSAKeyPairGenerator

The interface `Interface DSAKeyPairGenerator` is obsolete. It used to be needed to enable clients to provide DSA-specific parameters to be used rather than the default parameters your implementation supplies. However, in Java it is no longer necessary; a new `KeyPairGenerator initialize` method that takes an `AlgorithmParameterSpec` parameter enables clients to indicate algorithm-specific parameters.

DSAParams Implementation

If you are implementing a DSA key pair generator, you need a class implementing `Interface DSAParams` for holding and returning the `p`, `q`, and `g` parameters.

A `DSAParams` implementation is also required if you implement the `DSAPrivateKey` and `DSAPublicKey` interfaces. `DSAPublicKey` and `DSAPrivateKey` both extend the DSAKey interface, which contains a `getParams` method that must return a `DSAParams` object.

> **Note:**
>
> There is a `DSAParams` implementation built into the JDK: the `java.security.spec.DSAParameterSpec` class.

DSAPrivateKey and DSAPublicKey Implementations

If you implement a DSA key pair generator or key factory, you need to create classes implementing the `Interface DSAPrivateKey` and `Interface DSAPublicKey` interfaces.

If you implement a DSA key pair generator, your `generateKeyPair` method (in your `KeyPairGeneratorSpi` subclass) will return instances of your implementations of those interfaces.

If you implement a DSA key factory, your `engineGeneratePrivate` method (in your `KeyFactorySpi` subclass) will return an instance of your `DSAPrivateKey` implementation, and your `engineGeneratePublic` method will return an instance of your `DSAPublicKey` implementation.

Also, your `engineGetKeySpec` and `engineTranslateKey` methods will expect the passed-in key to be an instance of a `DSAPrivateKey` or `DSAPublicKey` implementation. The `getParams` method provided by the interface implementations is useful for obtaining and extracting the parameters from the keys and then using the parameters, for example as parameters to the `DSAParameterSpec` constructor called to create a parameter specification from parameter values that could be used to initialize a `KeyPairGenerator` object for DSA.

If you implement a DSA signature algorithm, your `engineInitSign` method (in your `SignatureSpi` subclass) will expect to be passed a `DSAPrivateKey` and your `engineInitVerify` method will expect to be passed a `DSAPublicKey`.

Please note: The `DSAPublicKey` and `DSAPrivateKey` interfaces define a very generic, provider-independent interface to DSA public and private keys, respectively. The `engineGetKeySpec` and `engineTranslateKey` methods (in your `KeyFactorySpi` subclass) could additionally check if the passed-in key is actually an instance of their provider's own implementation of `DSAPrivateKey` or `DSAPublicKey`, e.g., to take advantage of provider-specific implementation details. The same is true for the DSA signature algorithm `engineInitSign` and `engineInitVerify` methods (in your `SignatureSpi` subclass).

To see what methods need to be implemented by classes that implement the `DSAPublicKey` and `DSAPrivateKey` interfaces, first note the following interface signatures:

In the `java.security.interfaces` package:

```
public interface DSAPrivateKey extends DSAKey,
              java.security.PrivateKey

public interface DSAPublicKey extends DSAKey,
              java.security.PublicKey

public interface DSAKey
```

In the `java.security` package:

```
public interface PrivateKey extends Key
```

```
public interface PublicKey extends Key

public interface Key extends java.io.Serializable
```

In order to implement the `DSAPrivateKey` and `DSAPublicKey` interfaces, you must implement the methods they define as well as those defined by interfaces they extend, directly or indirectly.

Thus, for private keys, you need to supply a class that implements

- The `getX` method from the `Interface DSAPrivateKey` interface.
- The `getParams` method from the `Interface DSAKey` interface, since `DSAPrivateKey` extends `DSAKey`. Note: The `getParams` method returns a `DSAParams` object, so you must also have a `DSAParams` implementation.
- The `getAlgorithm`, `getEncoded`, and `getFormat` methods from the `Interface Key` interface, since `DSAPrivateKey` extends `java.security.PrivateKey`, and `PrivateKey` extends `Key`.
  Similarly, for public DSA keys, you need to supply a class that implements:
  - The `getY` method from the `Interface DSAPublicKey` interface.
  - The `getParams` method from the `Interface DSAKey` interface, since `DSAPublicKey` extends DSAKey.

    > **Note:**
    >
    > The `getParams` method returns a `DSAParams` object, so you must also have a DSAParams Implementation.

  - The `getAlgorithm`, `getEncoded`, and `getFormat` methods from the `Interface Key`, since `DSAPublicKey` extends `java.security.PublicKey`, and `PublicKey` extends `Key`.

RSA Interfaces and their Required Implementations

The Java Security API contains the interfaces (in the `java.security.interfaces` package) for the convenience of programmers implementing RSA services.

- `Interface RSAPrivateKey`
- `Interface RSAPrivateCrtKey`
- `Interface RSAPublicKey`

The following sections discuss requirements for implementations of these interfaces.

RSAPrivateKey, RSAPrivateCrtKey, and RSAPublicKey Implementations

If you implement an RSA key pair generator or key factory, you need to create classes implementing the `Interface RSAPublicKey` (and/or `Interface RSAPrivateCrtKey`) and `Interface RSAPublicKey` interfaces. (`RSAPrivateCrtKey` is the interface to an RSA private key, using the *Chinese Remainder Theorem* (CRT) representation.)

If you implement an RSA key pair generator, your `generateKeyPair` method (in your `KeyPairGeneratorSpi` subclass) will return instances of your implementations of those interfaces.

If you implement an RSA key factory, your `engineGeneratePrivate` method (in your `KeyFactorySpi` subclass) will return an instance of your `RSAPrivateKey` (or `RSAPrivateCrtKey`) implementation, and your `engineGeneratePublic` method will return an instance of your `RSAPublicKey` implementation.

Also, your `engineGetKeySpec` and `engineTranslateKey` methods will expect the passed-in key to be an instance of an `RSAPrivateKey`, `RSAPrivateCrtKey`, or `RSAPublicKey` implementation.

If you implement an RSA signature algorithm, your `engineInitSign` method (in your `SignatureSpi` subclass) will expect to be passed either an `RSAPrivateKey` or an `RSAPrivateCrtKey`, and your `engineInitVerify` method will expect to be passed an `RSAPublicKey`.

Please note: The `RSAPublicKey`, `RSAPrivateKey`, and `RSAPrivateCrtKey` interfaces define a very generic, provider-independent interface to RSA public and private keys. The `engineGetKeySpec` and `engineTranslateKey` methods (in your `KeyFactorySpi` subclass) could additionally check if the passed-in key is actually an instance of their provider's own implementation of `RSAPrivateKey`, `RSAPrivateCrtKey`, or `RSAPublicKey`, e.g., to take advantage of provider-specific implementation details. The same is true for the RSA signature algorithm `engineInitSign` and `engineInitVerify` methods (in your `SignatureSpi` subclass).

To see what methods need to be implemented by classes that implement the `RSAPublicKey`, `RSAPrivateKey`, and `RSAPrivateCrtKey` interfaces, first note the following interface signatures:

In the `java.security.interfaces` package:

```
public interface RSAPrivateKey extends java.security.PrivateKey

public interface RSAPrivateCrtKey extends RSAPrivateKey

public interface RSAPublicKey extends java.security.PublicKey
```

In the `java.security` package:

```
public interface PrivateKey extends Key

public interface PublicKey extends Key

public interface Key extends java.io.Serializable
```

In order to implement the `RSAPrivateKey`, `RSAPrivateCrtKey`, and `RSAPublicKey` interfaces, you must implement the methods they define as well as those defined by interfaces they extend, directly or indirectly.

Thus, for RSA private keys, you need to supply a class that implements:

- The `getModulus` and `getPrivateExponent` methods from the `Interface RSAPrivateKey` interface.
- The `getAlgorithm`, `getEncoded`, and `getFormat` methods from the `Interface Key` interface, since `RSAPrivateKey` extends `java.security.PrivateKey`, and `PrivateKey` extends `Key`.

Similarly, for RSA private keys using the *Chinese Remainder Theorem* (CRT) representation, you need to supply a class that implements:

- All the methods listed above for RSA private keys, since `RSAPrivateCrtKey` extends `java.security.interfaces.RSAPrivateKey`.
- The `getPublicExponent`, `getPrimeP`, `getPrimeQ`, `getPrimeExponentP`, `getPrimeExponentQ`, and `getCrtCoefficient` methods from the `Interface RSAPrivateKey` interface.

For public RSA keys, you need to supply a class that implements:

- The `getModulus` and `getPublicExponent` methods from the `Interface RSAPublicKey` interface.
- The `getAlgorithm`, `getEncoded`, and `getFormat` methods from the `Interface Key` interface, since `RSAPublicKey` extends `java.security.PublicKey`, and `PublicKey` extends `Key`.

JCA contains a number of `AlgorithmParameterSpec` implementations for the most frequently used cipher and key agreement algorithm parameters. If you are operating on algorithm parameters that should be for a different type of algorithm not provided by JCA, you will need to supply your own `AlgorithmParameterSpec` implementation appropriate for that type of algorithm.

Diffie-Hellman Interfaces and their Required Implementations

JCA contains interfaces (in the `javax.crypto.interfaces` package) for the convenience of programmers implementing Diffie-Hellman services.

- `Interface DHPublicKey`
- `Interface DHKey`
- `Interface DHPrivateKey`

The following sections discuss requirements for implementations of these interfaces.

DHPrivateKey and DHPublicKey Implementations

If you implement a Diffie-Hellman key pair generator or key factory, you need to create classes implementing the `Interface DHPrivateKey` and `Interface DHPublicKey` interfaces.

If you implement a Diffie-Hellman key pair generator, your `generateKeyPair` method (in your `KeyPairGeneratorSpi` subclass) will return instances of your implementations of those interfaces.

If you implement a Diffie-Hellman key factory, your `engineGeneratePrivate` method (in your `KeyFactorySpi` subclass) will return an instance of your `DHPrivateKey` implementation, and your `engineGeneratePublic` method will return an instance of your `DHPublicKey` implementation.

Also, your `engineGetKeySpec` and `engineTranslateKey` methods will expect the passed-in key to be an instance of a `DHPrivateKey` or `DHPublicKey` implementation. The `getParams` method provided by the interface implementations is useful for obtaining and extracting the parameters from the keys. You can then use the parameters, for example, as parameters to the `DHParameterSpec` constructor called to create a parameter specification from parameter values used to initialize a `KeyPairGenerator` object for Diffie-Hellman.

If you implement the Diffie-Hellman key agreement algorithm, your `engineInit` method (in your `KeyAgreementSpi` subclass) will expect to be passed a `DHPrivateKey` and your `engineDoPhase` method will expect to be passed a `DHPublicKey`.

> **Note:**
>
> The `DHPublicKey` and `DHPrivateKey` interfaces define a very generic, provider-independent interface to Diffie-Hellman public and private keys, respectively. The `engineGetKeySpec` and `engineTranslateKey` methods (in your `KeyFactorySpi` subclass) could additionally check if the passed-in key is actually an instance of their provider's own implementation of `DHPrivateKey` or `DHPublicKey`, e.g., to take advantage of provider-specific implementation details. The same is true for the Diffie-Hellman algorithm `engineInit` and `engineDoPhase` methods (in your `KeyAgreementSpi` subclass).

To see what methods need to be implemented by classes that implement the `DHPublicKey` and `DHPrivateKey` interfaces, first note the following interface signatures:

In the `javax.crypto.interfaces` package:

```
public interface DHPrivateKey extends DHKey, java.security.PrivateKey

public interface DHPublicKey extends DHKey, java.security.PublicKey

public interface DHKey
```

In the `java.security` package:

```
public interface PrivateKey extends Key

public interface PublicKey extends Key

public interface Key extends java.io.Serializable
```

To implement the `DHPrivateKey` and `DHPublicKey` interfaces, you must implement the methods they define as well as those defined by interfaces they extend, directly or indirectly.

Thus, for private keys, you need to supply a class that implements:

- The `getX` method from the `Interface DHPrivateKey` interface.

- The `getParams` method from the `Interface DHKey` interface, since `DHPrivateKey` extends `DHKey`.
- The `getAlgorithm`, `getEncoded`, and `getFormat` methods from the `Interface Key` interface, since `DHPrivateKey` extends `java.security.PrivateKey`, and `PrivateKey` extends `Key`.

Similarly, for public Diffie-Hellman keys, you need to supply a class that implements:

- The `getY` method from the `Interface DHPublicKey` interface.
- The `getParams` method from the `Interface DHKey` interface, since `DHPublicKey` extends `DHKey`.
- The `getAlgorithm`, `getEncoded`, and `getFormat` methods from the `Interface Key` interface, since `DHPublicKey` extends `java.security.PublicKey`, and `PublicKey` extends `Key`.

Interfaces for Other Algorithm Types

As noted above, the Java Security API contains interfaces for the convenience of programmers implementing services like DSA, RSA and ECC. If there are services without API support, you need to define your own APIs.

If you are implementing a key pair generator for a different algorithm, you should create an interface with one or more `initialize` methods that clients can call when they want to provide algorithm-specific parameters to be used rather than the default parameters your implementation supplies. Your subclass of `KeyPairGeneratorSpi` should implement this interface.

For algorithms without direct API support, it is recommended that you create similar interfaces and provide implementation classes. Your public key interface should extend the `Interface PublicKey` interface. Similarly, your private key interface should extend the `Interface PrivateKey` interface.

Algorithm Parameter Specification Interfaces and Classes

An algorithm parameter specification is a transparent representation of the sets of parameters used with an algorithm.

A *transparent* representation of parameters means that you can access each value individually, through one of the *get* methods defined in the corresponding specification class (e.g., `DSAParameterSpec` defines `getP`, `getQ`, and `getG` methods, to access the p, q, and g parameters, respectively).

This is contrasted with an *opaque* representation, as supplied by the `AlgorithmParameters` engine class, in which you have no direct access to the key material values; you can only get the name of the algorithm associated with the parameter set (via `getAlgorithm`) and some kind of encoding for the parameter set (via `getEncoded`).

If you supply an `AlgorithmParametersSpi`, `AlgorithmParameterGeneratorSpi`, or `KeyPairGeneratorSpi` implementation, you must utilize the `AlgorithmParameterSpec` interface, since each of those classes contain methods that take an `AlgorithmParameterSpec` parameter. Such methods need to determine which actual implementation of that interface has been passed in, and act accordingly.

JCA contains a number of `AlgorithmParameterSpec` implementations for the most frequently used signature, cipher and key agreement algorithm parameters. If you are operating on algorithm parameters that should be for a different type of algorithm not provided by JCA, you will need to supply your own `AlgorithmParameterSpec` implementation appropriate for that type of algorithm.

Java defines the following algorithm parameter specification interfaces and classes in the `java.security.spec` and `javax.crypto.spec` packages:

The AlgorithmParameterSpec Interface

`AlgorithmParameterSpec` is an interface to a transparent specification of cryptographic parameters.

This interface contains no methods or constants. Its only purpose is to group (and provide type safety for) all parameter specifications. All parameter specifications must implement this interface.

The DSAParameterSpec Class

This class (which implements the `AlgorithmParameterSpec` and `DSAParams` interfaces) specifies the set of parameters used with the DSA algorithm. It has the following methods:

```
public BigInteger getP()

public BigInteger getQ()

public BigInteger getG()
```

These methods return the DSA algorithm parameters: the prime `p`, the sub-prime `q`, and the base `g`.

Many types of DSA services will find this class useful - for example, it is utilized by the DSA signature, key pair generator, algorithm parameter generator, and algorithm parameters classes implemented by the *Sun* provider. As a specific example, an algorithm parameters implementation must include an implementation for the `getParameterSpec` method, which returns an `AlgorithmParameterSpec`. The DSA algorithm parameters implementation supplied by *Sun* returns an instance of the `DSAParameterSpec` class.

The IvParameterSpec Class

This class (which implements the `AlgorithmParameterSpec` interface) specifies the initialization vector (IV) used with a cipher in feedback mode.

Table 3-3 Method in `IvParameterSpec`

Method	Description
`byte[] getIV()`	Returns the initialization vector (IV).

The OAEPParameterSpec Class

This class specifies the set of parameters used with OAEP Padding, as defined in the PKCS #1 standard.

Table 3-4 Methods in `OAEPParameterSpec`

Method	Description
`String getDigestAlgorithm()`	Returns the message digest algorithm name.
`String getMGFAlgorithm()`	Returns the mask generation function algorithm name.
`AlgorithmParameterSpec getMGFParameters()`	Returns the parameters for the mask generation function.
`PSource getPSource()`	Returns the source of encoding input P.

The PBEParameterSpec Class

This class (which implements the `AlgorithmParameterSpec` interface) specifies the set of parameters used with a password-based encryption (PBE) algorithm.

Table 3-5 Methods in `PBEParameterSpec`

Method	Description
`int getIterationCount()`	Returns the iteration count.
`byte[] getSalt()`	Returns the salt.

The RC2ParameterSpec Class

This class (which implements the `AlgorithmParameterSpec` interface) specifies the set of parameters used with the RC2 algorithm.

Table 3-6 Methods in `RC2ParameterSpec`

Method	Description
`boolean equals(Object obj)`	Tests for equality between the specified object and this object.
`int getEffectiveKeyBits()`	Returns the effective key size in bits.
`byte[] getIV()`	Returns the IV or null if this parameter set does not contain an IV.
`int hashCode()`	Calculates a hash code value for the object.

The RC5ParameterSpec Class

This class (which implements the `AlgorithmParameterSpec` interface) specifies the set of parameters used with the RC5 algorithm.

Table 3-7 Methods in `RC5ParameterSpec`

Method	Description
`boolean equals(Object obj)`	Tests for equality between the specified object and this object.
`byte[] getIV()`	Returns the IV or null if this parameter set does not contain an IV.
`int getRounds()`	Returns the number of rounds.
`int getVersion()`	Returns the version.
`int getWordSize()`	Returns the word size in bits.
`int hashCode()`	Calculates a hash code value for the object.

The DHParameterSpec Class

This class (which implements the `AlgorithmParameterSpec` interface) specifies the set of parameters used with the Diffie-Hellman algorithm.

Table 3-8 Methods in DHParameterSpec

Method	Description
`BigInteger getG()`	Returns the base generator `g`.
`int getL()`	Returns the size in bits, `l`, of the random exponent (private value).
`BigInteger getP()`	Returns the prime modulus `p`.

Many types of Diffie-Hellman services will find this class useful; for example, it is used by the Diffie-Hellman key agreement, key pair generator, algorithm parameter generator, and algorithm parameters classes implemented by the "SunJCE" provider. As a specific example, an algorithm parameters implementation must include an implementation for the `getParameterSpec` method, which returns an `AlgorithmParameterSpec`. The Diffie-Hellman algorithm parameters implementation supplied by "SunJCE" returns an instance of the `DHParameterSpec` class.

Key Specification Interfaces and Classes Required by Key Factories

A key factory provides bi-directional conversions between opaque keys (of type `Key`) and key specifications. If you implement a key factory, you thus need to understand and utilize key specifications. In some cases, you also need to implement your own key specifications.

Key specifications are transparent representations of the key material that constitutes a key. If the key is stored on a hardware device, its specification may contain information that helps identify the key on the device.

A *transparent* representation of keys means that you can access each key material value individually, through one of the *get* methods defined in the corresponding specification class. For example, `java.security.spec.DSAPrivateKeySpec` defines `getX`, `getP`, `getQ`, and `getG` methods, to access the private key `x`, and the DSA algorithm parameters used to calculate the key: the prime `p`, the sub-prime `q`, and the base `g`.

This is contrasted with an *opaque* representation, as defined by the Key interface, in which you have no direct access to the parameter fields. In other words, an "opaque" representation gives you limited access to the key - just the three methods defined by the Key interface: `getAlgorithm`, `getFormat`, and `getEncoded`.

A key may be specified in an algorithm-specific way, or in an algorithm-independent encoding format (such as ASN.1). For example, a DSA private key may be specified by its components *x*, *p*, *q*, and *g* (see `DSAPrivateKeySpec`), or it may be specified using its DER encoding (see `PKCS8EncodedKeySpec`).

Java defines the following key specification interfaces and classes in the `java.security.spec` and `javax.crypto.spec` packages:

The `KeySpec` Interface

This interface contains no methods or constants. Its only purpose is to group (and provide type safety for) all key specifications. All key specifications must implement this interface.

Java supplies several classes implementing the `KeySpec` interface:

- `DSAPrivateKeySpec`
- `DSAPublicKeySpec`
- `RSAPrivateKeySpec`
- `RSAPublicKeySpec`
- `EncodedKeySpec`
- `PKCS8EncodedKeySpec`
- `X509EncodedKeySpec`

If your provider uses key types (e.g., `Your_PublicKey_type` and `Your_PrivateKey_type`) for which the JDK does not already provide corresponding `KeySpec` classes, there are two possible scenarios, one of which requires that you implement your own key specifications:

1. If your users will never have to access specific key material values of your key type, you will not have to provide any `KeySpec` classes for your key type.

   In this scenario, your users will always create `Your_PublicKey_type` and `Your_PrivateKey_type` keys through the appropriate `KeyPairGenerator` supplied by your provider for that key type. If they want to store the generated keys for later usage, they retrieve the keys' encodings (using the `getEncoded` method of the `Key` interface). When they want to create an `Your_PublicKey_type` or `Your_PrivateKey_type` key from the encoding (e.g., in order to initialize a Signature object for signing or verification), they create an instance of `X509EncodedKeySpec` or `PKCS8EncodedKeySpec` from the encoding, and feed it to the appropriate `KeyFactory` supplied by your provider for that algorithm, whose `generatePublic` and `generatePrivate` methods will return the requested `PublicKey` (an instance of `Your_PublicKey_type`) or `PrivateKey` (an instance of `Your_PrivateKey_type`) object, respectively.

2. If you anticipate a need for users to access specific key material values of your key type, or to construct a key of your key type from key material and associated parameter values, rather than from its encoding (as in the above case), you have to specify new `KeySpec` classes (classes that implement the `KeySpec` interface) with the appropriate constructor methods and *get* methods for returning key

material fields and associated parameter values for your key type. You will specify those classes in a similar manner as is done by the `DSAPrivateKeySpec` and `DSAPublicKeySpec` classes. You need to ship those classes along with your provider classes, for example, as part of your provider JAR file.

The DSAPrivateKeySpec Class

This class (which implements the `KeySpec` Interface) specifies a DSA private key with its associated parameters. It has the following methods:

Table 3-9 Methods in DSAPrivateKeySpec

Method in DSAPrivateKeySpec	Description
`public BigInteger getX()`	Returns the private key x.
`public BigInteger getP()`	Returns the prime p.
`public BigInteger getQ()`	Returns the sub-prime q.
`public BigInteger getG()`	Returns the base g.

These methods return the private key `x`, and the DSA algorithm parameters used to calculate the key: the prime `p`, the sub-prime `q`, and the base `g`.

The DSAPublicKeySpec Class

This class (which implements the `KeySpec` Interface) specifies a DSA public key with its associated parameters. It has the following methods:

Table 3-10 Methods in DSAPublicKeySpec

Method in DSAPublicKeySpec	Description
`public BigInteger getY()`	returns the public key y.
`public BigInteger getP()`	Returns the prime p.
`public BigInteger getQ()`	Returns the sub-prime q.
`public BigInteger getG()`	Returns the base g.

The RSAPrivateKeySpec Class

This class (which implements the `KeySpec` Interface) specifies an RSA private key. It has the following methods:

Table 3-11 Methods in RSAPrivateKeySpec

Method in RSAPrivateKeySpec	Description
`public BigInteger getModulus()`	Returns the modulus.
`public BigInteger getPrivateExponent()`	Returns the private exponent.

These methods return the RSA modulus `n` and private exponent `d` values that constitute the RSA private key.

The RSAPrivateCrtKeySpec Class

This class (which extends the `RSAPrivateKeySpec` class) specifies an RSA private key, as defined in the PKCS#1 standard, using the *Chinese Remainder Theorem* (CRT) information values. It has the following methods (in addition to the methods inherited from its superclass `RSAPrivateKeySpec`):

Table 3-12 Methods in RSAPrivateCrtKeySpec

Method in RSAPrivateCrtKeySpec	Description
`public BigInteger getPublicExponent()`	Returns the public exponent.
`public BigInteger getPrimeP()`	Returns the prime P.
`public BigInteger getPrimeQ()`	Returns the prime Q.
`public BigInteger getPrimeExponentP()`	Returns the primeExponentP.
`public BigInteger getPrimeExponentQ()`	Returns the primeExponentQ.
`public BigInteger getCrtCoefficient()`	Returns the crtCoefficient.

These methods return the public exponent `e` and the CRT information integers: the prime factor `p` of the modulus `n`, the prime factor `q` of `n`, the exponent `d mod (p-1)`, the exponent `d mod (q-1)`, and the Chinese Remainder Theorem coefficient `(inverse of q) mod p`.

An RSA private key logically consists of only the modulus and the private exponent. The presence of the CRT values is intended for efficiency.

The RSAPublicKeySpec Class

This class (which implements the `KeySpec` Interface) specifies an RSA public key. It has the following methods:

Table 3-13 Methods in RSAPublicKeySpec

Method in RSAPublicKeySpec	Description
`public BigInteger getModulus()`	Returns the modulus.
`public BigInteger getPublicExponent()`	Returns the public exponent.

The EncodedKeySpec Class

This abstract class (which implements the `KeySpec` Interface) represents a public or private key in encoded format.

Table 3-14 Methods in EncodedKeySpec

Method in EncodedKeySpec	Description
`public abstract byte[] getEncoded()`	Returns the encoded key.

Table 3-14 (Cont.) Methods in EncodedKeySpec

Method in EncodedKeySpec	Description
`public abstract String getFormat()`	Returns the name of the encoding format.

The JDK supplies two classes implementing the `EncodedKeySpec` interface: `PKCS8EncodedKeySpec` and `X509EncodedKeySpec`. If desired, you can supply your own `EncodedKeySpec` implementations for those or other types of key encodings.

The PKCS8EncodedKeySpec Class

This class, which is a subclass of `EncodedKeySpec`, represents the DER encoding of a private key, according to the format specified in the PKCS #8 standard.

Its `getEncoded` method returns the key bytes, encoded according to the PKCS #8 standard. Its `getFormat` method returns the string "PKCS#8".

The X509EncodedKeySpec Class

This class, which is a subclass of `EncodedKeySpec`, represents the DER encoding of a public or private key, according to the format specified in the X.509 standard.

Its `getEncoded` method returns the key bytes, encoded according to the X.509 standard. Its `getFormat` method returns the string "X.509".`DHPrivateKeySpec`, `DHPublicKeySpec`, `DESKeySpec`, `DESedeKeySpec`, `PBEKeySpec`, and `SecretKeySpec`.

The DHPrivateKeySpec Class

This class (which implements the `KeySpec` interface) specifies a Diffie-Hellman private key with its associated parameters.

Table 3-15 Methods in DHPrviateKeySpec

Method in DHPrivateKeySpec	Description
`BigInteger getG()`	Returns the base generator `g`.
`BigInteger getP()`	Returns the prime modulus `p`.
`BigInteger getX()`	Returns the private value `x`.

The DHPublicKeySpec Class

Table 3-16 Methods in DHPublicKeySpec

Method in DHPublicKeySpec	Description
`BigInteger getG()`	Returns the base generator `g`.
`BigInteger getP()`	Returns the prime modulus `p`.
`BigInteger getY()`	Returns the public value `y`.

The DESKeySpec Class

This class (which implements the `KeySpec` interface) specifies a DES key.

Table 3-17 Methods in DESKeySpec

Method in DESKeySpec	Description
`byte[] getKey()`	Returns the DES key bytes.
`static boolean isParityAdjusted(byte[] key, int offset)`	Checks if the given DES key material is parity-adjusted.
`static boolean isWeak(byte[] key, int offset)`	Checks if the given DES key material is weak or semi-weak.

The DESedeKeySpec Class

This class (which implements the `KeySpec` interface) specifies a DES-EDE (Triple DES) key.

Table 3-18 Methods in DESedeKeySpec

Method in DESedeKeySpec	Description
`byte[] getKey()`	Returns the DES-EDE key.
`static boolean isParityAdjusted(byte[] key, int offset)`	Checks if the given DES-EDE key is parity-adjusted.

The PBEKeySpec Class

This class implements the `KeySpec` interface. A user-chosen password can be used with password-based encryption (PBE); the password can be viewed as a type of raw key material. An encryption mechanism that uses this class can derive a cryptographic key from the raw key material.

Table 3-19 Methods in PBEKeySpec

Method in PBEKeySpec	Description
`void clearPassword`	Clears the internal copy of the password.
`int getIterationCount`	Returns the iteration count or 0 if not specified.
`int getKeyLength`	Returns the to-be-derived key length or 0 if not specified.
`char[] getPassword`	Returns a copy of the password.
`byte[] getSalt`	Returns a copy of the salt or null if not specified.

The SecretKeySpec Class

This class implements the `KeySpec` interface. Since it also implements the `SecretKey` interface, it can be used to construct a `SecretKey` object in a provider-independent fashion, i.e., without having to go through a provider-based `SecretKeyFactory`.

Table 3-20 Methods in SecretKeySpec

Method in SecretKeySpec	Description
`boolean equals (Object obj)`	Indicates whether some other object is "equal to" this one.
`String getAlgorithm()`	Returns the name of the algorithm associated with this secret key.
`byte[] getEncoded()`	Returns the key material of this secret key.
`String getFormat()`	Returns the name of the encoding format for this secret key.
`int hashCode()`	Calculates a hash code value for the object.

Secret-Key Generation

If you provide a secret-key generator (subclass of `javax.crypto.KeyGeneratorSpi`) for a particular secret-key algorithm, you may return the generated secret-key object.

The generated secret-key object (which must be an instance of `javax.crypto.SecretKey`, see engineGenerateKey) can be returned in one of the following ways:

- You implement a class whose instances represent secret-keys of the algorithm associated with your key generator. Your key generator implementation returns instances of that class. This approach is useful if the keys generated by your key generator have provider-specific properties.
- Your key generator returns an instance of `SecretKeySpec`, which already implements the `javax.crypto.SecretKey` interface. You pass the (raw) key bytes and the name of the secret-key algorithm associated with your key generator to the `SecretKeySpec` constructor. This approach is useful if the underlying (raw) key bytes can be represented as a byte array and have no key-parameters associated with them.

Adding New Object Identifiers

The following information applies to providers who supply an algorithm that is not listed as one of the standard algorithms in Java Security Standard Algorithm Names.

Mapping from OID to Name

Sometimes the JCA needs to instantiate a cryptographic algorithm implementation from an algorithm identifier (for example, as encoded in a certificate), which by definition includes the object identifier (OID) of the algorithm. For example, in order to verify the signature on an X.509 certificate, the JCA determines the signature algorithm from the signature algorithm identifier that is encoded in the certificate, instantiates a Signature object for that algorithm, and initializes the Signature object for verification.

For the JCA to find your algorithm, you must provide the object identifier of your algorithm as an alias entry for your algorithm in the provider master file.

```
put("Alg.Alias.<engine_type>.1.2.3.4.5.6.7.8",
    "<algorithm_alias_name>");
```

Note that if your algorithm is known under more than one object identifier, you need to create an alias entry for each object identifier under which it is known.

An example of where the JCA needs to perform this type of mapping is when your algorithm ("`Foo`") is a signature algorithm and users run the `keytool` command and specify your (signature) algorithm alias.

```
% keytool -genkeypair -sigalg 1.2.3.4.5.6.7.8
```

In this case, your provider master file should contain the following entries:

```
put("Signature.Foo", "com.xyz.MyFooSignatureImpl");
put("Alg.Alias.Signature.1.2.3.4.5.6.7.8", "Foo");
```

Other examples of where this type of mapping is performed are (1) when your algorithm is a keytype algorithm and your program parses a certificate (using the X.509 implementation of the SUN provider) and extracts the public key from the certificate in order to initialize a Signature object for verification, and (2) when `keytool` users try to access a private key of your keytype (for example, to perform a digital signature) after having generated the corresponding keypair. In these cases, your provider master file should contain the following entries:

```
put("KeyFactory.Foo", "com.xyz.MyFooKeyFactoryImpl");
put("Alg.Alias.KeyFactory.1.2.3.4.5.6.7.8", "Foo");
```

Mapping from Name to OID

If the JCA needs to perform the inverse mapping (that is, from your algorithm name to its associated OID), you need to provide an alias entry of the following form for one of the OIDs under which your algorithm should be known:

```
put("Alg.Alias.Signature.OID.1.2.3.4.5.6.7.8", "MySigAlg");
```

If your algorithm is known under more than one object identifier, prefix the preferred one with "OID."

An example of where the JCA needs to perform this kind of mapping is when users run `keytool` in any mode that takes a `-sigalg` option. For example, when the `-genkeypair` and `-certreq` commands are invoked, the user can specify your (signature) algorithm with the `-sigalg` option.

Ensuring Exportability

A key feature of JCA is the exportability of the JCA framework and of the provider cryptography implementations if certain conditions are met.

By default, an application can use cryptographic algorithms of any strength. However, due to import regulations in some countries, you may have to limit those algorithms' strength. You do this with jurisdiction policy files; see Cryptographic Strength Configuration. The JCA framework will enforce the restrictions specified in the installed jurisdiction policy files.

As noted elsewhere, you can write just one version of your provider software, implementing cryptography of maximum strength. It is up to JCA, not your provider, to enforce any jurisdiction policy file-mandated restrictions regarding the cryptographic algorithms and maximum cryptographic strengths available to applets/applications in different locations.

The conditions that must be met by your provider in order to enable it to be plugged into JCA are the following:

- The provider code should be written in such a way that provider classes become unusable if instantiated by an application directly, bypassing JCA. See Step 1: Write your Service Implementation Code in Steps to Implement and Integrate a Provider.
- The provider package must be signed by an entity trusted by the JCA framework. (See Step 7.1: Get a Code-Signing Certificate through Step 7.2: Sign Your Provider.) U.S. vendors whose providers may be exported outside the U.S. first need to apply for U.S. government export approval. (See Step 11: Apply for U.S. Government Export Approval If Required.)

Sample Code for MyProvider

The following is the complete source code for an example provider, MyProvider. It's a portable provider; you can specify it in a class or module path. It consists of two modules:

- `com.example.MyProvider`: Contains an example provider that demonstrate how to write a provider with the `Provider.Service` mechanism. You must compile, package, and sign the provider, then specify it in your class or module path as described in Steps to Implement and Integrate a Provider.
- `com.example.MyApp`: Contains a sample application that uses the MyProvider provider. It finds and loads this provider with the `ServiceLoader` mechanism, and then registers it dynamically with the `Security.addProvider()` method.

This example consists of the following files:

- src/com.example.MyProvider/module-info.java
- src/com.example.MyProvider/com/example/MyProvider/MyProvider.java
- src/com.example.MyProvider/com/example/MyProvider/MyCipher.java
- src/com.example.MyProvider/META-INF/services/java.security.Provider
- src/com.example.MyApp/module-info.java

- src/com.example.MyApp/com/example/MyApp/MyApp.java
- RunTest.sh

src/com.example.MyProvider/module-info.java

See Step 4: Create a Module Declaration for Your Provider for information about the module declaration, which is specified in `module-info.java`.

```
module com.example.MyProvider {
    provides java.security.Provider with com.example.MyProvider.MyProvider;
}
```

src/com.example.MyProvider/com/example/MyProvider/MyProvider.java

The `MyProvider` class is an example of a provider that uses the `Provider.Service` class. See Step 3.2: Create a Provider That Uses Provider.Service.

```
package com.example.MyProvider;

import java.security.*;
import java.util.*;

/**
 * Test JCE provider.
 *
 * Registers services using Provider.Service and overrides newInstance().
 */
public final class MyProvider extends Provider {

    public MyProvider() {
        super("MyProvider", "1.0", "My JCE provider");

        final Provider p = this;

        AccessController.doPrivileged((PrivilegedAction<Void>) () -> {
            putService(new ProviderService(p, "Cipher",
                    "MyCipher", "com.example.MyProvider.MyCipher"));
            return null;
        });
    }

    private static final class ProviderService extends Provider.Service {

        ProviderService(Provider p, String type, String algo, String cn) {
            super(p, type, algo, cn, null, null);
        }

        ProviderService(Provider p, String type, String algo, String cn,
                String[] aliases, HashMap<String, String> attrs) {
            super(p, type, algo, cn,
                    (aliases == null ? null : Arrays.asList(aliases)),
attrs);
        }
```

```
        @Override
        public Object newInstance(Object ctrParamObj)
                throws NoSuchAlgorithmException {

            String type = getType();
            if (ctrParamObj != null) {
                throw new InvalidParameterException(
                        "constructorParameter not used with " + type
                        + " engines");
            }
            String algo = getAlgorithm();
            try {
                if (type.equals("Cipher")) {
                    if (algo.equals("MyCipher")) {
                        return new MyCipher();
                    }
                }
            } catch (Exception ex) {
                throw new NoSuchAlgorithmException(
                        "Error constructing " + type + " for "
                        + algo + " using SunMSCAPI", ex);
            }
            throw new ProviderException("No impl for " + algo
                    + " " + type);
        }
    }

    @Override
    public String toString() {
        return "MyProvider [getName()=" + getName()
                + ", getVersionStr()=" + getVersionStr() + ", getInfo()="
                + getInfo() + "]";
    }
}
```

src/com.example.MyProvider/com/example/MyProvider/MyCipher.java

The `MyCipher` class extends the `CipherSPI`, which is a Server Provider Interface (SPI). Each cryptographic service that a provider implements has a subclass of the appropriate SPI. See Step 1: Write your Service Implementation Code.

> **Note:**
>
> This code is only a stub provider that demonstrates how to write a provider; it's missing the actual cryptographic algorithm implementation. The `MyCipher` class would contain an actual cryptographic algorithm implementation if MyProvider were a real security provider.

```
package com.example.MyProvider;

import java.security.*;
```

```
import java.security.spec.*;
import javax.crypto.*;

/**
 * Implementation represents a test Cipher.
 *
 * All are stubs.
 */
public class MyCipher extends CipherSpi {

    @Override
    protected byte[] engineDoFinal(byte[] input, int inputOffset, int
inputLen)
            throws IllegalBlockSizeException, BadPaddingException {
        return null;
    }

    @Override
    protected int engineDoFinal(byte[] input, int inputOffset, int
inputLen,
            byte[] output, int outputOffset) throws ShortBufferException,
            IllegalBlockSizeException, BadPaddingException {
        return 0;
    }

    @Override
    protected int engineGetBlockSize() {
        return 0;
    }

    @Override
    protected byte[] engineGetIV() {
        return null;
    }

    @Override
    protected int engineGetOutputSize(int inputLen) {
        return 0;
    }

    @Override
    protected AlgorithmParameters engineGetParameters() {
        return null;
    }

    @Override
    protected void engineInit(int opmode, Key key, SecureRandom random)
            throws InvalidKeyException {
    }

    @Override
    protected void engineInit(int opmode, Key key,
            AlgorithmParameterSpec params, SecureRandom random)
            throws InvalidKeyException, InvalidAlgorithmParameterException
{
```

```
    }

    @Override
    protected void engineInit(int opmode, Key key, AlgorithmParameters
params,
            SecureRandom random) throws InvalidKeyException,
            InvalidAlgorithmParameterException {
    }

    @Override
    protected void engineSetMode(String mode) throws
NoSuchAlgorithmException {
    }

    @Override
    protected void engineSetPadding(String padding)
            throws NoSuchPaddingException {
    }

    @Override
    protected int engineGetKeySize(Key key)
            throws InvalidKeyException {
        return 0;
    }

    @Override
    protected byte[] engineUpdate(byte[] input, int inputOffset, int
inputLen) {
        return null;
    }

    @Override
    protected int engineUpdate(byte[] input, int inputOffset, int inputLen,
            byte[] output, int outputOffset) throws ShortBufferException {
        return 0;
    }
}
```

src/com.example.MyProvider/META-INF/services/java.security.Provider

The `java.security.Provider` file enables automatic or unnamed modules to use the `ServiceLoader` class to search for your providers. See Step 6: Place Your Provider in a JAR File.

```
com.example.MyProvider.MyProvider
```

src/com.example.MyApp/module-info.java

This file contains a `uses` directive, which specifies a service that the module requires. This directive helps the module system locate providers and ensure that they run

reliably. This is the complement to the `provides` directive in the `MyProvider` module definition.

```
module com.example.MyApp {
    uses java.security.Provider;
}
```

src/com.example.MyApp/com/example/MyApp/MyApp.java

```
package com.example.MyApp;

import java.util.*;
import java.security.*;
import javax.crypto.*;

/**
 * A simple JCE test client to access a simple test Provider/Cipher
 * implementation in a signed modular jar.
 */
public class MyApp {

    private static final String PROVIDER = "MyProvider";
    private static final String CIPHER = "MyCipher";

    public static void main(String[] args) throws Exception {

        /*
         * Registers MyProvider dynamically.
         *
         * Could do statically by editing the java.security file.
         * Use the first form if using ServiceLoader ("uses" or
         * META-INF/service), the second if using the traditional class
         * lookup method.  Both if provider could be deployed to either.
         *
         * security.provider.14=MyProvider
         * security.provider.15=com.example.MyProvider.MyProvider
         */
        ServiceLoader<Provider> sl =
            ServiceLoader.load(java.security.Provider.class);
        for (Provider p : sl) {
            if (p.getName().equals(PROVIDER)) {
                System.out.println("Registering the Provider");
                Security.addProvider(p);
            }
        }

        /*
         * Get a MyCipher from MyProvider and initialize it.
         */
        Cipher cipher = Cipher.getInstance(CIPHER, PROVIDER);
        cipher.init(Cipher.ENCRYPT_MODE, (Key) null);

        /*
         * What Provider did we get?
```

```
         */
        Provider p = cipher.getProvider();
        Class c = p.getClass();
        Module m = c.getModule();
        System.out.println(p.getName() + ": version "
            + p.getVersionStr() + "\n"
            + p.getInfo() + "\n    "
            + ((m.getName() == null) ? "<UNNAMED>" : m.getName())
            + "/" + c.getName());
    }
}
```

RunTest.sh

```
#!/bin/sh

#
# A simple example to show how a JCE provider could be developed in a
# modular JDK, for deployment as either Named/Unnamed modules.
#

#
# Edit as appropriate
#
JDK_DIR=d:/java/jdk9
KEYSTORE=YourKeyStore
STOREPASS=YourStorePass
SIGNER=YourAlias

echo "-----------"
echo "Clean/Init"
echo "-----------"
rm -rf mods jars
mkdir mods jars

echo "--------------------"
echo "Compiling MyProvider"
echo "--------------------"
${JDK_DIR}/bin/javac.exe \
    --module-source-path src \
    -d mods \
    $(find src/com.example.MyProvider -name '*.java' -print)

echo "-----------------------------------"
echo "Packaging com.example.MyProvider.jar"
echo "-----------------------------------"
${JDK_DIR}/bin/jar.exe --create \
    --file jars/com.example.MyProvider.jar \
    --verbose \
    --module-version=1.0 \
    -C mods/com.example.MyProvider . \
    -C src/com.example.MyProvider META-INF/services

echo "----------------------------------"
```

```
echo "Signing com.example.MyProvider.jar"
echo "----------------------------------"
${JDK_DIR}/bin/jarsigner.exe \
    -keystore ${KEYSTORE} \
    -storepass ${STOREPASS} \
    jars/com.example.MyProvider.jar ${SIGNER}

echo "---------------"
echo "Compiling MyApp"
echo "---------------"
${JDK_DIR}/bin/javac.exe \
    --module-source-path src \
    -d mods \
    $(find src/com.example.MyApp -name '*.java' -print)

echo "------------------------------"
echo "Packaging com.example.MyApp.jar"
echo "------------------------------"
${JDK_DIR}/bin/jar.exe --create \
    --file jars/com.example.MyApp.jar \
    --verbose \
    --module-version=1.0 \
    -C mods/com.example.MyApp .

echo "------------------------"
echo "Test1                   "
echo "Named Provider/Named App"
echo "------------------------"
${JDK_DIR}/bin/java.exe \
    --module-path 'jars' \
    -m com.example.MyApp/com.example.MyApp.MyApp

echo "--------------------------"
echo "Test2                     "
echo "Named Provider/Unnamed App"
echo "--------------------------"
${JDK_DIR}/bin/java.exe \
    --module-path 'jars/com.example.MyProvider.jar' \
    --class-path 'jars/com.example.MyApp.jar' \
    com.example.MyApp.MyApp

echo "--------------------------"
echo "Test3                     "
echo "Unnamed Provider/Named App"
echo "--------------------------"
${JDK_DIR}/bin/java.exe \
    --module-path 'jars/com.example.MyApp.jar' \
    --class-path 'jars/com.example.MyProvider.jar' \
    -m com.example.MyApp/com.example.MyApp.MyApp

echo "----------------------------"
echo "Test4                       "
echo "Unnamed Provider/Unnamed App"
echo "----------------------------"
${JDK_DIR}/bin/java.exe \
```

```
--class-path \
    'jars/com.example.MyProvider.jar;jars/com.example.MyApp.jar' \
com.example.MyApp.MyApp
```

4
JDK Providers Documentation

This document contains the technical details of the providers that are included in the JDK. It is assumed that readers have a strong understanding of the Java Cryptography Architecture and Provider Architecture.

> **Note:**
>
> The Java Security Standard Algorithm Names contains more information about the standard names used in this document.

Topics

Introduction to JDK Providers

Import Limits on Cryptographic Algorithms

Cipher Transformations

SecureRandom Implementations

The SunPKCS11 Provider

The SUN Provider

The SunRsaSign Provider

The SunJSSE Provider

The SunJCE Provider

The SunJGSS Provider

The SunSASL Provider

The XMLDSig Provider

The SunPCSC Provider

The SunMSCAPI Provider

The SunEC Provider

The OracleUcrypto Provider

The Apple Provider

The JdkLDAP Provider

The JdkSASL Provider

Introduction to JDK Providers

The Java platform defines a set of APIs spanning major security areas, including cryptography, public key infrastructure, authentication, secure communication, and access control. These APIs enable developers to easily integrate security mechanisms into their application code.

The Java Cryptography Architecture (JCA) and its Provider Architecture are core concepts of the Java Development Kit (JDK). It is assumed that readers have a solid understanding of this architecture.

Reminder: Cryptographic implementations in the JDK are distributed through several different providers ("SUN", "SunJSSE", "SunJCE", "SunRsaSign") for both historical reasons and by the types of services provided. General purpose applications **SHOULD NOT** request cryptographic services from specific providers. That is:

```
getInstance("...", "SunJCE");  // not recommended
```

versus

```
getInstance("...");            // recommended
```

Otherwise, applications are tied to specific providers that may not be available on other Java implementations. They also might not be able to take advantage of available optimized providers (for example, hardware accelerators via PKCS11 or native OS implementations such as Microsoft's MSCAPI) that have a higher preference order than the specific requested provider.

The following table lists the modules and the supported Java Cryptographic Service Providers:

Table 4-1 Modules and the Java Cryptographic Service Providers

Module	Provider(s)
java.base	SUN, SunRsaSign, SunJSSE, SunJCE[1], Apple
java.naming	JdkLDAP
java.security.jgss	SunJGSS
java.security.sasl	SunSASL
java.smartcardio	SunPCSC
java.xml.crypto	XMLDSig
jdk.crypto.cryptoki	SunPKCS11[1]
jdk.crypto.ec	SunEC[1]
jdk.crypto.mscapi	SunMSCAPI[1]
jdk.crypto.ucrypto	OracleUcrypto[1]
jdk.security.jgss	JdkSASL

[1] Indicates JCE crypto providers previously distributed as signed JAR files (JCE providers contain `Cipher/KeyAgreement/KeyGenerator/Mac/SecretKeyFactory` implementations).

Import Limits on Cryptographic Algorithms

By default, an application can use cryptographic algorithms of any strength. However, due to import regulations in some locations, you may have to limit the strength of those algorithms. The JDK provides two different sets of jurisdiction policy files in the directory `<java-home>/conf/security/policy` that determine the strength of cryptographic algorithms. Information about jurisdiction policy files and how to activate them is available in Cryptographic Strength Configuration.

Consult your export/import control counsel or attorney to determine the exact requirements for your location.

For the "limited" configuration, the following table lists the maximum key sizes allowed by the "limited" set of jurisdiction policy files:

Table 4-2 Maximum Keysize of Cryptographic Algorithms

Algorithm	Maximum Keysize
DES	64
DESede	*
RC2	128
RC4	128
RC5	128
RSA	*
all others	128

Cipher Transformations

The `javax.crypto.Cipher.getInstance(String transformation)` factory method generates `Cipher` objects using transformations of the form **algorithm/mode/padding**. If the mode/padding are omitted, the SunJCE and SunPKCS11 providers use ECB as the default mode and PKCS5Padding as the default padding for many symmetric ciphers.

It is recommended to use transformations that fully specify the algorithm, mode, and padding instead of relying on the defaults. The defaults are provider specific and can vary among providers.

> **Note:**
>
> ECB works well for single blocks of data and can be parallelized, but absolutely should not be used for multiple blocks of data.

SecureRandom Implementations

The following table lists the default preference order of the available `SecureRandom` implementations.

Table 4-3 Default SecureRandom Implementations

OS	Algorithm Name	Provider Name
Solaris	1. PKCS11[1,2]	SunPKCS11
	2. NativePRNG[3]	SUN
	3. DRBG	SUN
	4. SHA1PRNG[3]	SUN
	5. NativePRNGBlocking	SUN
	6. NativePRNGNonBlocking	SUN
Linux	1. NativePRNG[3]	SUN
	2. DRBG	SUN
	3. SHA1PRNG[3]	SUN
	4. NativePRNGBlocking	SUN
	5. NativePRNGNonBlocking	SUN
macOS	1. NativePRNG[3]	SUN
	2. DRBG	SUN
	3. SHA1PRNG[3]	SUN
	4. NativePRNGBlocking	SUN
	5. NativePRNGNonBlocking	SUN
Windows	1. DRBG	SUN
	2. SHA1PRNG	SUN
	3. Windows-PRNG[4]	SunMSCAPI

1 The SunPKCS11 provider is available on all platforms, but is only enabled by default on Solaris as it is the only OS with a native PKCS11 implementation automatically installed and configured. On other platforms, applications or deployers must specifically install and configure a native PKCS11 library, and then configure and enable the SunPKCS11 provider to use it.

2 The PKCS11 `SecureRandom` implementation for Solaris has been disabled due to the performance overhead of small-sized requests (see JDK-8098581: SecureRandom.nextBytes() hurts performance with small size requests). Edit `sunpkcs11-solaris.cfg` to reenable.

3 On Solaris, Linux, and OS X, if the *entropy gathering device* in `java.security` is set to `file:/dev/urandom` or `file:/dev/random`, then NativePRNG is preferred to SHA1PRNG. Otherwise, SHA1PRNG is preferred.

4 There is currently no NativePRNG on Windows. Access to the equivalent functionality is via the SunMSCAPI provider.

The SunPKCS11 Provider

The Cryptographic Token Interface Standard (PKCS#11) provides native programming interfaces to cryptographic mechanisms, such as hardware cryptographic accelerators and Smart Cards. When properly configured, the `SunPKCS11` provider enables applications to use the standard JCA/JCE APIs to access native PKCS#11 libraries. **The `SunPKCS11` provider itself does not contain cryptographic functionality**, it is simply a conduit between the Java environment and the native PKCS11 providers. The PKCS#11 Reference Guide has a much more detailed treatment of this provider.

The SUN Provider

Algorithms

The following algorithms are available in the `SUN` provider:

Table 4-4 Algorithms in SUN provider

Engine	Algorithm Names
`AlgorithmParameterGenerator`	DSA
`AlgorithmParameters`	DSA
`CertificateFactory`	X.509
`CertPathBuilder`	PKIX
`CertPathValidator`	PKIX
`CertStore`	Collection
`Configuration`	JavaLoginConfig
`KeyFactory`	DSA
`KeyPairGenerator`	DSA
`KeyStore`	PKCS12[1] JKS DKS CaseExactJKS
`MessageDigest`	MD2 MD5 SHA-1 SHA-224 SHA-256 SHA-384 SHA-512 SHA-512/224 SHA-512/256 SHA3-224 SHA3-256 SHA3-384 SHA3-512
`Policy`	JavaPolicy

Table 4-4 (Cont.) Algorithms in SUN provider

Engine	Algorithm Names
`SecureRandom`	DRBG (The following mechanisms and algorithms are supported: Hash_DRBG and HMAC_DRBG with SHA-224, SHA-512/224, SHA-256, SHA-512/256, SHA-384 and SHA-512. CTR_DRBG (both use derivation function and not use) with AES-128, AES-192 and AES-256. Prediction resistance and reseeding supported for each combination, and security strength can be requested from 112 up to the highest strength one supports.) SHA1PRNG (Initial seeding is currently done via a combination of system attributes and the `java.security` entropy gathering device.) NativePRNG (`nextBytes()` uses `/dev/urandom`, `generateSeed()` uses `/dev/random`) NativePRNGBlocking (`nextBytes()` and `generateSeed()` use `/dev/random`) NativePRNGNonBlocking (`nextBytes()` and `generateSeed()` use `/dev/urandom`)
`Signature`	NONEwithDSA SHA1withDSA SHA224withDSA SHA256withDSA NONEwithDSAinP1363Format SHA1withDSAinP1363Format SHA224withDSAinP1363Format SHA256withDSAinP1363Format

[1] The PKCS12 `KeyStore` implementation does not support the `KeyBag` type.

OIDs Associated with SHA Message Digests and DSA Signatures

The following table lists OIDs associated with SHA Message Digests:

Table 4-5 OIDs associated with SHA Message Digests

SHA Message Digest	OID
SHA-224	2.16.840.1.101.3.4.2.4
SHA-256	2.16.840.1.101.3.4.2.1
SHA-384	2.16.840.1.101.3.4.2.2
SHA-512	2.16.840.1.101.3.4.2.3

Table 4-5 (Cont.) OIDs associated with SHA Message Digests

SHA Message Digest	OID
SHA-512/224	2.16.840.1.101.3.4.2.5
SHA-512/256	2.16.840.1.101.3.4.2.6
SHA3-224	2.16.840.1.101.3.4.2.7
SHA3-256	2.16.840.1.101.3.4.2.8
SHA3-384	2.16.840.1.101.3.4.2.9
SHA3-512	2.16.840.1.101.3.4.2.10

The following table lists OIDs associated with DSA Signatures:

Table 4-6 OIDs associated with DSA Signatures

DSA Signature	OID
SHA1withDSA	1.2.840.10040.4.3 1.3.14.3.2.13 1.3.14.3.2.27
SHA224withDSA	2.16.840.1.101.3.4.3.1
SHA256withDSA	2.16.840.1.101.3.4.3.2

Keysize Restrictions

The `SUN` provider uses the following default keysizes (in bits) and enforces the following restrictions:

Table 4-7 KeyPairGenerator Algorithm Keysize Restrictions

Algorithm Name	Default Keysize	Restrictions/Comments
DSA	2048	Keysize must be a multiple of 64, ranging from 512 to 1024, plus 2048 and 3072.

Table 4-8 AlgorithmParameterGenerator Algorithm Keysize Restrictions

Algorithm Name	Default Keysize	Restrictions/Comments
DSA	2048	Keysize must be a multiple of 64, ranging from 512 to 1024, plus 2048 and 3072.

CertificateFactory/CertPathBuilder/CertPathValidator/CertStore Implementations

See Appendix B: CertPath Implementation in SUN Provider in the Java PKI Programmer's Guide Additional details on the `SUN` provider implementations for `CertificateFactory`, `CertPathBuilder`, `CertPathValidator` and `CertStore` are documented in of the .

The SunRsaSign Provider

Algorithms

The following algorithms are available in the `SunRsaSign` provider:

Table 4-9 SunRsaSign Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`AlgorithmParameters`	RSASSA-PSS
`KeyFactory`	RSA RSASSA-PSS
`KeyPairGenerator`	RSA RSASSA-PSS
`Signature`	MD2withRSA MD5withRSA SHA1withRSA SHA224withRSA SHA256withRSA SHA384withRSA SHA512withRSA SHA512/224withRSA SHA512/256withRSA RSASSA-PSS SHA1withRSAandMGF1 SHA224withRSAandMGF1 SHA256withRSAandMGF1 SHA384withRSAandMGF1 SHA512withRSAandMGF1 SHA512/224withRSAandMGF1 SHA512/256withRSAandMGF1

Keysize Restrictions

The `SunRsaSign` provider uses the following default keysize (in bits) and enforces the following restriction:

KeyPairGenerator

Table 4-10 SunRsaSign Provider Keysize Restrictions

Alg. Name	Default Keysize	Restrictions/Comments
RSA and RSASSA-PSS	2048	Keysize must range between 512 and 16384 bits. If the key size exceeds 3072, then the public exponent length cannot exceed 64 bits.

The SunJSSE Provider

Algorithms

The following algorithms are available in the `SunJSSE` provider:

Table 4-11 Algorithms in SunJSSE Provider

Engine	Algorithm Name(s)
`KeyFactory`	RSA **Note:** The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for `KeyFactory`.
`KeyManagerFactory`	PKIX: A factory for `X509ExtendedKeyManager` instances that manage X.509 certificate-based key pairs for local side authentication according to the rules defined by the IETF PKIX working group in RFC 5280. This `KeyManagerFactory` currently supports initialization using a `KeyStore` object or `javax.net.ssl.KeyStoreBuilderParameters`. SunX509: A factory for `X509ExtendedKeyManager` instances that manage X.509 certificate-based key pairs for local side authentication, but with less strict checking of certificate usage/validity and chain verification. This `KeyManagerFactory` supports initialization using a `Keystore` object, but does not currently support initialization using the class `javax.net.ssl.ManagerFactoryParameters`. **Note:** The SunX509 factory is for backwards compatibility with older releases, and should no longer be used.
`KeyPairGenerator`	RSA **Note:** The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for `KeyPairGenerator`.
`KeyStore`	PKCS12 **Note:** The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for `KeyStore`.

Table 4-11 (Cont.) Algorithms in SunJSSE Provider

Engine	Algorithm Name(s)
`Signature`	MD2withRSA MD5withRSA SHA1withRSA **Note:** The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for `Signature`. These `Signature` algorithms are planned to be removed from the `SunJSSE` provider in a future release.
`SSLContext`	SSL SSLv3 TLS TLSv1 TLSv1.1 TLSv1.2 TLSv1.3 DTLS DTLSv1.0 DTLSv1.2
`TrustManagerFactory`	PKIX: A factory for `X509ExtendedTrustManager` instances that validate certificate chains according to the rules defined by the IETF PKIX working group in RFC 5280. This `TrustManagerFactory` currently supports initialization using a `KeyStore` object or `javax.net.ssl.CertPathTrustManagerParameters`. SunX509: A factory for `X509ExtendedTrustManager` instances that validate certificate chains, but with less strict checking of certificate usage/validity and chain verification. This `TrustManagerFactory` supports initialization using a `Keystore` object, but does not currently support initialization using the class `javax.net.ssl.ManagerFactoryParameters`. **Note:** The SunX509 factory is for backwards compatibility with older releases, and should no longer be used.

SunJSSE Provider Protocol Parameters

The `SunJSSE` provider supports the following `protocol` parameters:

Table 4-12 SunJSSE Provider Protocol Parameters

Protocol	Enabled by Default for Client	Enabled by Default for Server
SSL	Yes	Yes
SSLv3	No	No
TLS	Yes	Yes
TLSv1	Yes	Yes
TLSv1.1	Yes	Yes
TLSv1.2	Yes	Yes
TLSv1.3	Yes	Yes
SSLv2Hello	No	Yes
DTLS	Yes	Yes
DTLSv1.0	Yes	Yes
DTLSv1.2	Yes	Yes

> **Note:**
>
> Starting with JDK 8u31, the SSLv3 protocol (Secure Socket Layer) has been disabled by the `jdk.tls.disabledAlgorithms` Security Property. See Disabled and Restricted Cryptographic Algorithms and RFC 7568: Deprecating Secure Sockets Layer Version 3.0.

The following table illustrates which connection combinations are possible when using SSLv2Hellos:

Table 4-13 Connections Possible Using SSLv2Hellos

Client	Server	Connection Possible?
Enabled	Enabled	Yes
Not enabled	Enabled	Yes (most interoperable: SunJSSE default)
Enabled	Not enabled	No
Not enabled	Not enabled	Yes

SunJSSE Cipher Suites

The following are the currently implemented SunJSSE cipher suites for this JDK release, sorted by order of preference. Not all of these cipher suites are available for use by default. See JSSE Cipher Suite Names in *Java Security Standard Algorithm Names* to determine which protocols that each cipher suite supports.

- TLS_AES_128_GCM_SHA256
- TLS_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256

- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_DSS_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_DSS_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_DHE_DSS_WITH_AES_256_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA
- TLS_ECDH_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_DSS_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_DSS_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA

- TLS_ECDH_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_DSS_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA
- TLS_EMPTY_RENEGOTIATION_INFO_SCSV
- TLS_DH_anon_WITH_AES_256_GCM_SHA384
- TLS_DH_anon_WITH_AES_128_GCM_SHA256
- TLS_DH_anon_WITH_AES_256_CBC_SHA256
- TLS_ECDH_anon_WITH_AES_256_CBC_SHA
- TLS_DH_anon_WITH_AES_256_CBC_SHA
- TLS_DH_anon_WITH_AES_128_CBC_SHA256
- TLS_ECDH_anon_WITH_AES_128_CBC_SHA
- TLS_DH_anon_WITH_AES_128_CBC_SHA
- TLS_ECDH_anon_WITH_3DES_EDE_CBC_SHA
- SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_RC4_128_SHA
- TLS_ECDHE_RSA_WITH_RC4_128_SHA
- SSL_RSA_WITH_RC4_128_SHA
- TLS_ECDH_ECDSA_WITH_RC4_128_SHA
- TLS_ECDH_RSA_WITH_RC4_128_SHA
- SSL_RSA_WITH_RC4_128_MD5
- TLS_ECDH_anon_WITH_RC4_128_SHA
- SSL_DH_anon_WITH_RC4_128_MD5
- SSL_RSA_WITH_DES_CBC_SHA
- SSL_DHE_RSA_WITH_DES_CBC_SHA
- SSL_DHE_DSS_WITH_DES_CBC_SHA
- SSL_DH_anon_WITH_DES_CBC_SHA
- SSL_RSA_EXPORT_WITH_DES40_CBC_SHA
- SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA
- SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
- SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA

- SSL_RSA_EXPORT_WITH_RC4_40_MD5
- SSL_DH_anon_EXPORT_WITH_RC4_40_MD5
- TLS_RSA_WITH_NULL_SHA256
- TLS_ECDHE_ECDSA_WITH_NULL_SHA
- TLS_ECDHE_RSA_WITH_NULL_SHA
- SSL_RSA_WITH_NULL_SHA
- TLS_ECDH_ECDSA_WITH_NULL_SHA
- TLS_ECDH_RSA_WITH_NULL_SHA
- TLS_ECDH_anon_WITH_NULL_SHA
- SSL_RSA_WITH_NULL_MD5

> **Note:**
>
> - The cipher suite order of preference may change in future releases.
> - TLS_EMPTY_RENEGOTIATION_INFO_SCSV is a pseudo-cipher suite that supports RFC 5746.

The cipher suites available by default in a JDK release change as new algorithms are developed and old algorithms are found to be less effective than previously thought. Oracle JDK uses two mechanisms to restrict the availability of these algorithms:

- The `jdk.tls.disabledAlgorithms` Security Property, which disables categories of cipher suites. For example, if this Security Property contains `RC4`, then all RC4-based cipher suites would be disabled.
- Moving the cipher suite to the list of suites not enabled by default.

See Disabled and Restricted Cryptographic Algorithms for information about the `jdk.tls.disabledAlgorithms` Security Property.

Determining Current List of Cipher Suites Available by Default

To determine the current list of cipher suites available by default, run the following application, `AvailableCipherSuites.java`:

```
import java.util.*;
import java.security.*;
import javax.net.ssl.*;

public class AvailableCipherSuites {

    public static void main(String[] args) throws Exception {

        // If an argument is present, then remove the
        // jdk.tls.disabledAlgorithms restrictions and
        // print all implemented cipher suites.
        if (args.length != 0) {
            Security.setProperty("jdk.tls.disabledAlgorithms", "");
```

```
        }

        SSLContext sslc = SSLContext.getDefault();
        SSLSocketFactory sslf = sslc.getSocketFactory();
        SSLSocket ssls = (SSLSocket) sslf.createSocket();

        ArrayList<String> enabled = new ArrayList(
                Arrays.asList(ssls.getEnabledCipherSuites()));

        ArrayList<String> supported = new ArrayList(
                Arrays.asList(ssls.getSupportedCipherSuites()));
        supported.removeAll(enabled);
        System.out.println("Enabled by Default Cipher Suites");
        System.out.println("-------------------------------");
        enabled.stream().forEach(System.out::println);

        System.out.println();

        System.out.println("Not Enabled by Default Cipher Suites");
        System.out.println("-----------------------------------");
        supported.stream().forEach(System.out::println);
    }
}
```

Without arguments, this application calls the methods `SSLSocket.getSupportedCipherSuites()` and `SSLSocket.getEnabledCipherSuites()` to print the available enabled and supported cipher suites that are not disabled by the `jdk.tls.disabledAlgorithms` Security Property. Running this application with an argument removes the "Disabled by Security Property" restriction; as a result, the application prints all possible cipher suites.

Tighter Checking of EncryptedPreMasterSecret Version Numbers

Prior to the JDK 7 release, the SSL/TLS implementation did not check the version number in PreMasterSecret, and the SSL/TLS client did not send the correct version number by default. Unless the system property `com.sun.net.ssl.rsaPreMasterSecretFix` is set to `true`, the TLS client sends the active negotiated version, but not the expected maximum version supported by the client.

For compatibility, this behavior is preserved for SSL version 3.0 and TLS version 1.0. However, for TLS version 1.1 or later, the implementation tightens checking the PreMasterSecret version numbers as required by RFC 5246. Clients always send the correct version number, and servers check the version number strictly. The system property, `com.sun.net.ssl.rsaPreMasterSecretFix`, is not used in TLS 1.1 or later.

The SunJCE Provider

Algorithms

The following algorithms are available in the SunJCE provider:

Table 4-14 SunJCE Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`AlgorithmParameterGenerator`	DiffieHellman
`AlgorithmParameters`	AES Blowfish DES DESede DiffieHellman GCM OAEP PBE PBES2 PBEWithHmacSHA1AndAES_128 PBEWithHmacSHA224AndAES_128 PBEWithHmacSHA256AndAES_128 PBEWithHmacSHA384AndAES_128 PBEWithHmacSHA512AndAES_128 PBEWithHmacSHA1AndAES_256 PBEWithHmacSHA224AndAES_256 PBEWithHmacSHA256AndAES_256 PBEWithHmacSHA384AndAES_256 PBEWithHmacSHA512AndAES_256 PBEWithMD5AndDES PBEWithMD5AndTripleDES PBEWithSHA1AndDESede PBEWithSHA1AndRC2_40 PBEWithSHA1AndRC2_128 PBEWithSHA1AndRC4_40 PBEWithSHA1AndRC4_128 RC2
`Cipher`	See Table 4-15
`KeyAgreement`	DiffieHellman
`KeyFactory`	DiffieHellman
`KeyGenerator`	AES ARCFOUR Blowfish DES DESede HmacMD5 HmacSHA1 HmacSHA224 HmacSHA256 HmacSHA384 HmacSHA512 RC2

Table 4-14 (Cont.) SunJCE Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`KeyPairGenerator`	DiffieHellman
`KeyStore`	JCEKS
`Mac`	HmacMD5 HmacSHA1 HmacSHA224 HmacSHA256 HmacSHA384 HmacSHA512 HmacSHA512/224 HmacSHA512/256 HmacPBESHA1 PBEWithHmacSHA1 PBEWithHmacSHA224 PBEWithHmacSHA256 PBEWithHmacSHA384 PBEWithHmacSHA512
`SecretKeyFactory`	DES DESede PBEWithMD5AndDES PBEWithMD5AndTripleDES PBEWithSHA1AndDESede PBEWithSHA1AndRC2_40 PBEWithSHA1AndRC2_128 PBEWithSHA1AndRC4_40 PBEWithSHA1AndRC4_128 PBKDF2WithHmacSHA1 PBKDF2WithHmacSHA224 PBKDF2WithHmacSHA256 PBKDF2WithHmacSHA384 PBKDF2WithHmacSHA512 PBEWithHmacSHA1AndAES_128 PBEWithHmacSHA224AndAES_128 PBEWithHmacSHA256AndAES_128 PBEWithHmacSHA384AndAES_128 PBEWithHmacSHA512AndAES_128 PBEWithHmacSHA1AndAES_256 PBEWithHmacSHA224AndAES_256 PBEWithHmacSHA256AndAES_256 PBEWithHmacSHA384AndAES_256 PBEWithHmacSHA512AndAES_256

Cipher Transformations

The following table lists cipher transformations available in the SunJCE provider.

Table 4-15 SunJCE Provider Cipher Transformations

Algorithm Names	Modes	Paddings
AES	ECB, CBC, PCBC, CFB[1], CFB8..CFB128, OFB[1], OFB8..OFB128	NoPadding, PKCS5Padding, ISO10126Padding
AES	CTR, CTS, GCM	NoPadding
AES_128, AES_192, AES_256	ECB, CBC, OFB, CFB, GCM	NoPadding
AESWrap	ECB	NoPadding
AESWrap_128	ECB	NoPadding
AESWrap_192	ECB	NoPadding
AESWrap_256	ECB	NoPadding
ARCFOUR	ECB	NoPadding
Blowfish, DES, DESede, RC2	ECB, CBC, PCBC, CTR, CTS, CFB[1], CFB8..CFB64, OFB[1], OFB8..OFB64	NoPadding, PKCS5Padding, ISO10126Padding
DESedeWrap	CBC	NoPadding
PBEWithMD5AndDES, PBEWithMD5AndTripleDES[2], PBEWithSHA1AndDESede, PBEWithSHA1AndRC2_40, PBEWithSHA1AndRC2_128, PBEWithSHA1AndRC4_40, PBEWithSHA1AndRC4_128, PBEWithHmacSHA1AndAES_128, PBEWithHmacSHA224AndAES_128, PBEWithHmacSHA256AndAES_128, PBEWithHmacSHA384AndAES_128, PBEWithHmacSHA512AndAES_128, PBEWithHmacSHA1AndAES_256, PBEWithHmacSHA224AndAES_256, PBEWithHmacSHA256AndAES_256, PBEWithHmacSHA384AndAES_256, PBEWithHmacSHA512AndAES_256	CBC	PKCS5Padding

Table 4-15 (Cont.) SunJCE Provider Cipher Transformations

Algorithm Names	Modes	Paddings
RSA	ECB	NoPadding, PKCS1Padding, OAEPPadding, OAEPWithMD5AndMGF1Padding, OAEPWithSHA-1AndMGF1Padding, OAEPWithSHA-1AndMGF1Padding, OAEPWithSHA-224AndMGF1Padding, OAEPWithSHA-256AndMGF1Padding, OAEPWithSHA-384AndMGF1Padding, OAEPWithSHA-512AndMGF1Padding, OAEPWithSHA-512/224AndMGF1Padding, OAEPWithSHA-512/2256ndMGF1Padding

[1] CFB/OFB with no specified value defaults to the block size of the algorithm. (for example, AES is 128; Blowfish, DES, DESede, and RC2 are 64.)
[2] PBEWithMD5AndTripleDES is a proprietary algorithm that has not been standardized.

Keysize Restrictions

The SunJCE provider uses the following default keysizes (in bits) and enforces the following restrictions:

Table 4-16 KeyGenerator Keysize Restrictions

Algorithm Name	Default Keysize	Restrictions/Comments
AES	128	Keysize must be equal to 128, 192, or 256.
ARCFOUR (RC4)	128	Keysize must range between 40 and 1024 (inclusive).
Blowfish	128	Keysize must be a multiple of 8, ranging from 32 to 448 (inclusive).
DES	56	Keysize must be equal to 56.

Table 4-16 (Cont.) KeyGenerator Keysize Restrictions

Algorithm Name	Default Keysize	Restrictions/Comments
DESede (Triple DES)	168	Keysize must be equal to 112 or 168. A keysize of 112 will generate a Triple DES key with 2 intermediate keys, and a keysize of 168 will generate a Triple DES key with 3 intermediate keys. Due to the "Meet-In-The-Middle" problem, even though 112 or 168 bits of key material are used, the **effective** keysize is 80 or 112 bits respectively.
HmacMD5	512	No keysize restriction.
HmacSHA1	512	No keysize restriction.
HmacSHA224	224	No keysize restriction.
HmacSHA256	256	No keysize restriction.
HmacSHA384	384	No keysize restriction.
HmacSHA512	512	No keysize restriction.
RC2	128	Keysize must range between 40 and 1024 (inclusive).

Note:

The various Password-Based Encryption (PBE) algorithms use various algorithms to generate key data, and ultimately depends on the targeted Cipher algorithm. For example,
”PBEWithMD5AndDES” will always generate 56–bit keys.

Table 4-17 KeyPairGenerator Algorithm Keysize Restrictions

Algorithm Name	Default Keysize	Restrictions/Comments
Diffie-Hellman (DH)	2048	Keysize must be a multiple of 64, ranging from 512 to 1024, plus 1536, 2048, 3072, 4096, 6144, 8192.

Table 4-18 AlgorithmParameterGenerator Algorithm Keysize Restrictions

Algorithm Name	Default Keysize	Restrictions/Comments
Diffie-Hellman (DH)	2048	Keysize must be a multiple of 64, ranging from 512 to 1024, plus 2048 and 3072.

The SunJGSS Provider

Algorithms

The following algorithms are available in the SunJGSS provider:

Table 4-19 SunJGSS Provider Algorithm Names

OID	Name
`1.2.840.113554.1.2.2`	Kerberos v5
`1.3.6.1.5.5.2`	SPNEGO

The SunSASL Provider

Algorithms

The following algorithms are available in the `SunSASL` provider:

Table 4-20 SunSASL Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`SaslClient`	CRAM-MD5 DIGEST-MD5 EXTERNAL NTLM PLAIN
`SaslServer`	CRAM-MD5 DIGEST-MD5 NTLM

The XMLDSig Provider

Algorithms

The following algorithms are available in the `XMLDSig` provider:

Table 4-21 XMLDSig Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`KeyInfoFactory`	DOM

Table 4-21 (Cont.) XMLDSig Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`TransformService`	• http://www.w3.org/TR/2001/REC-xml-c14n-20010315 – `CanonicalizationMethod.INCLUSIVE` • http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments – `CanonicalizationMethod.INCLUSIVE_WITH_COMMENTS` • http://www.w3.org/2001/10/xml-exc-c14n# – `CanonicalizationMethod.EXCLUSIVE` • http://www.w3.org/2001/10/xml-exc-c14n#WithComments – `CanonicalizationMethod.EXCLUSIVE_WITH_COMMENTS` • http://www.w3.org/2000/09/xmldsig#base64 – `Transform.BASE64` • http://www.w3.org/2000/09/xmldsig#enveloped-signature – `Transform.ENVELOPED` • http://www.w3.org/TR/1999/REC-xpath-19991116 – `Transform.XPATH` • http://www.w3.org/2002/06/xmldsig-filter2 – `Transform.XPATH2` • http://www.w3.org/TR/1999/REC-xslt-19991116 – `Transform.XSLT`
`XMLSignatureFactory`	DOM

The SunPCSC Provider

The SunPCSC provider enables applications to use the Java Smart Card I/O API to interact with the PC/SC Smart Card stack of the underlying operating system. Consult your operating system documentation for details.

On Solaris and Linux, SunPCSC accesses the PC/SC stack via the `libpcsclite.so` library. It looks for this library in the directories `/usr/$LIBISA` and `/usr/local/$LIBISA`, where `$LIBISA` is expanded to `lib/64` on 64-bit Solaris and `lib64` on 64-bit Linux. The system property `sun.security.smartcardio.library` may also be set to the full filename of an alternate `libpcsclite.so` implementation. On Windows, SunPCSC always calls into `winscard.dll` and no Java-level configuration is necessary or possible.

If PC/SC is available on the host platform, the SunPCSC implementation can be obtained via `TerminalFactory.getDefault()` and `TerminalFactory.getInstance("PC/SC")`. If PC/SC is not available or not correctly configured, a `getInstance()` call will fail with a `NoSuchAlgorithmException` and `getDefault()` will return a JRE built-in implementation that does not support any terminals.

Algorithms

The following algorithms are available in the `SunPCSC` provider:

Table 4-22 The SunPCSC Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`TerminalFactory`	PC/SC

The SunMSCAPI Provider

The `SunMSCAPI` provider enables applications to use the standard JCA/JCE APIs to access the native cryptographic libraries, certificates stores and key containers on Windows. The `SunMSCAPI` provider itself does not contain cryptographic functionality, it is simply a conduit between the Java environment and the native cryptographic services on Windows.

Algorithms

The following algorithms are available in the `SunMSCAPI` provider:

Table 4-23 The SunMSCAPI Algorithm Names for Engine Classes

Engine	Algorithm Names
`Cipher`	RSA *RSA/ECB/PKCS1Padding only*
`KeyPairGenerator`	RSA
`KeyStore`	**Windows-MY** : The keystore type that identifies the native Microsoft Windows MY keystore. It contains the user's personal certificates and associated private keys. **Windows-ROOT**: The keystore type that identifies the native Microsoft Windows ROOT keystore. It contains the certificates of Root certificate authorities and other self-signed trusted certificates.
`SecureRandom`	**Windows-PRNG** : The name of the native pseudo-random number generation (PRNG) algorithm.
`Signature`	MD5withRSA MD2withRSA NONEwithRSA SHA1withRSA SHA256withRSA SHA384withRSA SHA512withRSA RSASSA-PSS

Keysize Restrictions

The SunMSCAPI provider uses the following default keysizes (in bits) and enforce the following restrictions:

KeyGenerator

Table 4-24 The SunMSCAPI Provider Keysize Restrictions

Alg. Name	Default Keysize	Restrictions/Comments
RSA	2048	Keysize ranges from 512 bits to 16,384 bits (depending on the underlying Microsoft Windows cryptographic service provider).

The SunEC Provider

The `SunEC` provider implements Elliptical Curve Cryptography (ECC). Compared to traditional cryptosystems such as RSA, ECC offers equivalent security with smaller key sizes, which results in faster computations, lower power consumption, and memory and bandwidth savings. Applications can now use the standard JCA/JCE APIs to access ECC functionality without the dependency on external ECC libraries (through `SunPKCS11`).

Algorithms

The following algorithms are available in the `SunEC` provider:

Table 4-25 The SunEC Provider Names for Engine Classes

Engine	Algorithm Name(s)
`AlgorithmParameters`	EC
`KeyAgreement`	ECDH, X25519, X448, XDH
`KeyFactory`	EC, X25519, X448, XDH
`KeyPairGenerator`	EC, X25519, X448, XDH
`Signature`	NONEwithECDSA SHA1withECDSA SHA224withECDSA SHA256withECDSA SHA384withECDSA SHA512withECDSA NONEwithECDSAinP1363Format SHA1withECDSAinP1363Format SHA224withECDSAinP1363Format SHA256withECDSAinP1363Format SHA384withECDSAinP1363Format SHA512withECDSAinP1363Format

Note:

- The XDH algorithm can be initialized with either X25519 or X448 parameters and keys.
- The X25519 algorithm supports X25519 parameters and keys only. Similarly, the X448 algorithm supports X448 parameters and keys only.

Keysize Restrictions

The SunEC provider uses the following default keysizes (in bits) and enforces the following restrictions:

Table 4-26 The SunEC Provider Keysize Restrictions

KeyPairGenerator Algorithm Name	Default Keysize	Restrictions/Comments
EC	256	Keysize must range from 112 to 571 (inclusive).
X25519	255	Keysize must be 255
X448	448	Keysize must be 448
XDH	255	Keysize must be 255 or 448

Supported Elliptic Curve Names

Table 4-27 lists the elliptic curve (EC) names that `SunEC` supports for the EC, Elliptic-Curve Diffie-Hellman (ECDH), and Elliptic Curve Digital Signature Algorithm (ECDSA) algorithms. The first column lists all the `SunEC` supported EC names. The second column specifies the EC name's object identifier. The third column specifies any additional `SunEC` supported EC names. (A value of N/A means that there are no additional names.) All strings that appear in one row refer to the same EC. For example, the strings `secp256r1`, `1.2.840.10045.3.1.7`, `NIST P-256`, and `X9.62 prime256v1` refer to the same EC. You can use EC names to create parameter specifications for EC parameter generation with the `ECGenParameterSpec` class.

Table 4-27 SunEC Supported EC Names

Supported Name	Object Identifier	Additional Supported Names
brainpoolP256r1	1.3.36.3.3.2.8.1.1.7	N/A
brainpoolP320r1	1.3.36.3.3.2.8.1.1.9	N/A
brainpoolP384r1	1.3.36.3.3.2.8.1.1.11	N/A
brainpoolP512r1	1.3.36.3.3.2.8.1.1.13	N/A
secp112r1	1.3.132.0.6	N/A
secp112r2	1.3.132.0.7	N/A
secp128r1	1.3.132.0.28	N/A
secp128r2	1.3.132.0.29	N/A
secp160k1	1.3.132.0.9	N/A
secp160r1	1.3.132.0.8	N/A

Table 4-27 (Cont.) SunEC Supported EC Names

Supported Name	Object Identifier	Additional Supported Names
secp160r2	1.3.132.0.30	N/A
secp192k1	1.3.132.0.31	N/A
secp192r1	1.2.840.10045.3.1.1	NIST P-192, X9.62 prime192v1
secp224k1	1.3.132.0.32	N/A
secp224r1	1.3.132.0.33	NIST P-224
secp256k1	1.3.132.0.10	N/A
secp256r1	1.2.840.10045.3.1.7	NIST P-256, X9.62 prime256v1
secp384r1	1.3.132.0.34	NIST P-384
secp521r1	1.3.132.0.35	NIST P-521
sect113r1	1.3.132.0.4	N/A
sect113r2	1.3.132.0.5	N/A
sect131r1	1.3.132.0.22	N/A
sect131r2	1.3.132.0.23	N/A
sect163k1	1.3.132.0.1	NIST K-163
sect163r1	1.3.132.0.2	N/A
sect163r2	1.3.132.0.15	NIST B-163
sect193r1	1.3.132.0.24	N/A
sect193r2	1.3.132.0.25	N/A
sect233k1	1.3.132.0.26	NIST K-233
sect233r1	1.3.132.0.27	NIST B-233
sect239k1	1.3.132.0.3	N/A
sect283k1	1.3.132.0.16	NIST K-283
sect283r1	1.3.132.0.17	NIST B-283
sect409k1	1.3.132.0.36	NIST K-409
sect409r1	1.3.132.0.37	NIST B-409
sect571k1	1.3.132.0.38	NIST K-571
sect571r1	1.3.132.0.39	NIST B-571
X9.62 c2tnb191v1	1.2.840.10045.3.0.5	N/A
X9.62 c2tnb191v2	1.2.840.10045.3.0.6	N/A
X9.62 c2tnb191v3	1.2.840.10045.3.0.7	N/A
X9.62 c2tnb239v1	1.2.840.10045.3.0.11	N/A
X9.62 c2tnb239v2	1.2.840.10045.3.0.12	N/A
X9.62 c2tnb239v3	1.2.840.10045.3.0.13	N/A
X9.62 c2tnb359v1	1.2.840.10045.3.0.18	N/A
X9.62 c2tnb431r1	1.2.840.10045.3.0.20	N/A
X9.62 prime192v2	1.2.840.10045.3.1.2	N/A
X9.62 prime192v3	1.2.840.10045.3.1.3	N/A
X9.62 prime239v1	1.2.840.10045.3.1.4	N/A
X9.62 prime239v2	1.2.840.10045.3.1.5	N/A
X9.62 prime239v3	1.2.840.10045.3.1.6	N/A

The OracleUcrypto Provider

The Solaris-only security provider `OracleUcrypto` leverages the Solaris Ucrypto library to offload and delegate cryptographic operations supported by the Oracle SPARC T4 based on-core cryptographic instructions. The `OracleUcrypto` provider itself does not contain cryptographic functionality; it is simply a conduit between the Java environment and the Solaris Ucrypto library.

If the underlying Solaris Ucrypto library does not support a particular algorithm, then the `OracleUcrypto` provider will not support it either. Consequently, at runtime, the supported algorithms consists of the intersection of those that the Solaris Ucrypto library supports and those that the `OracleUcrypto` provider recognizes.

Algorithms

The following algorithms are available in the `OracleUcrypto` provider:

Table 4-28 The OracleUcrypto Provider Algorithm Names for Engine Classes

Engine	Algorithm Name(s)
`Cipher`	AES RSA AES/ECB/NoPadding AES/ECB/PKCS5Padding AES/CBC/NoPadding AES/CBC/PKCS5Padding AES/CTR/NoPadding AES/GCM/NoPadding AES/CFB128/NoPadding AES/CFB128/PKCS5Padding AES_128/ECB/NoPadding AES_192/ECB/NoPadding AES_256/ECB/NoPadding AES_128/CBC/NoPadding AES_192/CBC/NoPadding AES_256/CBC/NoPadding AES_128/GCM/NoPadding AES_192/GCM/NoPadding AES_256/GCM/NoPadding RSA/ECB/PKCS1Padding RSA/ECB/NoPadding
`Signature`	MD5withRSA SHA1withRSA SHA256withRSA SHA384withRSA SHA512withRSA

Table 4-28 (Cont.) The OracleUcrypto Provider Algorithm Names for Engine Classes

Engine	Algorithm Name(s)
`MessageDigest`	MD5 SHA SHA-224 SHA-256 SHA-384 SHA-512 SHA3–224 SHA3–256 SHA3–384 SHA3–512

Keysize Restrictions

The `OracleUcrypto` provider does not specify any default keysizes or keysize restrictions; these are specified by the underlying Solaris Ucrypto library.

OracleUcrypto Provider Configuration File

The `OracleUcrypto` provider has a configuration file named `ucrypto-solaris.cfg` that resides in the `$JAVA_HOME/conf/security` directory. Modify this configuration file to specify which algorithms to disable by default. For example, the following configuration file disables AES with CFB128 mode by default:

```
#
# Configuration file for the OracleUcrypto provider
#
disabledServices = {
  Cipher.AES/CFB128/PKCS5Padding
  Cipher.AES/CFB128/NoPadding
}
```

The Apple Provider

The `Apple` provider implements a `java.security.KeyStore` that provides access to the macOS Keychain.

Algorithms

The following algorithms are available in the `Apple` provider:

Table 4-29 The Apple Provider Algorithm Name for Engine Classes

Engine	Algorithm Name(s)
`KeyStore`	KeychainStore

The JdkLDAP Provider

The `JdkLDAP` provider replaces the LDAP CertStore implementation in the `SUN` provider.

Algorithms

The following algorithms are available in the `JdkLDAP` provider:

Table 4-30 The JdkLDAP Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`CertStore`	LDAP

The JdkSASL Provider

Algorithms

The following algorithms are available in the `JdkSASL` provider:

Table 4-31 The JdkSASL Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
`SaslClient`	GSSAPI
`SaslServer`	GSSAPI

5
PKCS#11 Reference Guide

The Java platform defines a set of programming interfaces for performing cryptographic operations. These interfaces are collectively known as the Java Cryptography Architecture (JCA) and the Java Cryptography Extension (JCE). See Java Cryptography Architecture (JCA) Reference Guide.

The cryptographic interfaces are provider-based. Specifically, applications talk to Application Programming Interfaces (APIs), and the actual cryptographic operations are performed in configured providers which adhere to a set of Service Provider Interfaces (SPIs). This architecture supports different provider implementations. Some providers may perform cryptographic operations in software; others may perform the operations on a hardware token (for example, on a smartcard device or on a hardware cryptographic accelerator).

The Cryptographic Token Interface Standard, PKCS#11, is produced by RSA Security and defines native programming interfaces to cryptographic tokens, such as hardware cryptographic accelerators and smartcards. Existing applications that use the JCA and JCE APIs can access native PKCS#11 tokens with the PKCS#11 provider. No modifications to the application are required. The only requirement is to properly configure the provider.

Although an application can make use of most PKCS#11 features using existing APIs, some applications might need more flexibility and capabilities. For example, an application might want to deal with smartcards being removed and inserted dynamically more easily. Or, a PKCS#11 token might require authentication for some non-key-related operations and therefore, the application must be able to log into the token without using keystore. The JCA gives applications greater flexibility in dealing with different providers.

This document describes how native PKCS#11 tokens can be configured into the Java platform for use by Java applications. It also describes how the JCA makes it easier for applications to deal with different types of providers, including PKCS#11 providers.

SunPKCS11 Provider

The SunPKCS11 provider, in contrast to most other providers, does not implement cryptographic algorithms itself. Instead, it acts as a bridge between the Java JCA and JCE APIs and the native PKCS#11 cryptographic API, translating the calls and conventions between the two.

This means that Java applications calling standard JCA and JCE APIs can, without modification, take advantage of algorithms offered by the underlying PKCS#11 implementations, such as, for example,

- Cryptographic smartcards,
- Hardware cryptographic accelerators, and
- High performance software implementations.

> **Note:**
>
> Java SE only facilitates accessing native PKCS#11 implementations, it does not itself include a native PKCS#11 implementation. However, cryptographic devices such as Smartcards and hardware accelerators often come with software that includes a PKCS#11 implementation, which you need to install and configure according to manufacturer's instructions.

SunPKCS11 Requirements

The SunPKCS11 provider requires an implementation of PKCS#11 v2.20 or later to be installed on the system. This implementation must take the form of a shared-object library (`.so` on Solaris and Linux) or dynamic-link library (`.dll` on Windows or `.dylib` on macOS). Consult your vendor documentation to find out if your cryptographic device includes such a PKCS#11 implementation, how to configure it, and what the name of the library file is.

The SunPKCS11 provider supports a number of algorithms, provided that the underlying PKCS#11 implementation offers them. The algorithms and their corresponding PKCS#11 mechanisms are listed in the table in SunPKCS11 Provider Supported Algorithms.

SunPKCS11 Configuration

The SunPKCS11 provider is in the module `jdk.crypto.cryptoki`. To use the provider, you must first install it statically or programmatically.

To install the provider statically, add the provider to the Java security properties file (*`java-home`*`/conf/security/java.security`).

For example, here's a fragment of the `java.security` file that installs the SunPKCS11 provider with the configuration file `/opt/bar/cfg/pkcs11.cfg`.

```
# configuration for security providers 1-12 omitted
security.provider.13=SunPKCS11 /opt/bar/cfg/pkcs11.cfg
```

To install the provider dynamically, create an instance of the provider with the appropriate configuration filename and then install it. Here is an example.

```
String configName = "/opt/bar/cfg/pkcs11.cfg";
Provider p = Security.getProvider("SunPKCS11");
p = p.configure(configName);
Security.addProvider(p);
```

Note:

Save the returned `Provider` object from the `configure` method, then add that object, as demonstrated in this example:

```
p = p.configure(configName);
Security.addProvider(p);
```

Don't add the provider from which you called the `configure` method:

```
p.configure(configName);
Security.addProvider(p);
```

If this provider cannot be configured in-place, then a new provider is created and returned. Therefore, always use the provider returned from the `configure` method.

To use more than one slot per PKCS#11 implementation, or to use more than one PKCS#11 implementation, simply repeat the installation for each with the appropriate configuration file. This will result in a SunPKCS11 provider instance for each slot of each PKCS#11 implementation.

The configuration file is a text file that contains entries in the following format:

`attribute=value`
The valid values for *attribute* and *value* are described in the table in this section:

The two mandatory attributes are `name` and `library`.
Here is a sample configuration file:

```
name = FooAccelerator
library = /opt/foo/lib/libpkcs11.so
```

Comments are denoted by lines starting with the # (number) symbol.

Table 5-1 Attributes in the PKCS#11 Provider Configuration File

Attribute	Value	Description
library	Pathname of PKCS#11 implementation	This is the full pathname (including extension) of the PKCS#11 implementation; the format of the pathname is platform dependent. For example, `/opt/foo/lib/libpkcs11.so` might be the pathname of a PKCS#11 implementation on Solaris and Linux while `C:\foo\mypkcs11.dll` might be the pathname on Windows or `/opt/local/lib/libpkcs11.dylib` on macOS. **Note**: • To configure Mozilla Network Security Services (NSS) as the PKCS#11 implementation then set this attribute to the full pathname of the NSS softokn3 library. • Depending on your platform, you may have to set the environment variable `LD_LIBRARY_PATH` (Solaris or Linux), `PATH` (Windows), or `DYLD_LIBRARY_PATH` (macOS) to include the enclosing directory to enable the operating system to locate the dependent libraries.
name	Name suffix of this provider instance	This string is concatenated with the prefix `SunPKCS11-` to produce this provider instance's name (that is, the string returned by its `Provider.getName()` method). For example, if the `name` attribute is `"FooAccelerator"`, then the provider instance's name will be `"SunPKCS11-FooAccelerator"`.

Table 5-1 (Cont.) Attributes in the PKCS#11 Provider Configuration File

Attribute	Value	Description
description	Description of this provider instance	This string will be returned by the provider instance's `Provider.getInfo()` method. If none is specified, a default description will be returned.
slot	Slot id	This is the id of the slot that this provider instance is to be associated with. For example, you would use `1` for the slot with the id `1` under PKCS#11. At most one of `slot` or `slotListIndex` may be specified. If neither is specified, the default is a `slotListIndex` of `0`.
slotListIndex	Slot index	This is the slot index that this provider instance is to be associated with. It is the index into the list of all slots returned by the PKCS#11 function `C_GetSlotList`. For example, `0` indicates the first slot in the list. At most one of `slot` or `slotListIndex` may be specified. If neither is specified, the default is a `slotListIndex` of `0`.

Table 5-1 (Cont.) Attributes in the PKCS#11 Provider Configuration File

Attribute	Value	Description

enabledMechanisms	Brace enclosed, whitespace-separated list of PKCS#11 mechanisms to enable	This is the list PKCS#11 mechanisms that this provider instance should use, provided that they are supported by both the SunPKCS11 provider and PKCS#11 token. All other mechanisms will be ignored. Each entry in the list is the name of a PKCS#11 mechanism. Here is an example that lists two PKCS#11 mechanisms.
		
		`enabledMechanisms = {`
		`  CKM_RSA_PKCS`
		
		`CKM_RSA_PKCS_KEY_PAIR_G`
		`EN`
		`}`
		
		At most one of `enabledMechanisms` or `disabledMechanisms` may be specified. If neither is specified, the mechanisms enabled are those that are supported by both the SunPKCS11 provider (see SunPKCS11 Provider Supported Algorithms) and the PKCS#11 token.

disabledMechanisms	Brace enclosed, whitespace-separated list of PKCS#11 mechanisms to disable	This is the list of PKCS#11 mechanisms that this provider instance should ignore. Any mechanism listed will be ignored by the provider, even if they are supported by the token and the SunPKCS11 provider. The strings `SecureRandom` and `KeyStore` may be specified to disable those services. At most one of `enabledMechanisms` or `disabledMechanisms` may be specified. If neither is specified, the mechanisms enabled are those that are supported by both the SunPKCS11 provider (see SunPKCS11 Provider Supported Algorithms) and the PKCS#11 token.

Table 5-1 (Cont.) Attributes in the PKCS#11 Provider Configuration File

Attribute	Value	Description
attributes	See below	The `attributes` option can be used to specify additional PKCS#11 that should be set when creating PKCS#11 key objects. This makes it possible to accommodate tokens that require particular attributes. For details, see the section below.

Attributes Configuration

The attributes option allows you to specify additional PKCS#11 attributes that should be set when creating PKCS#11 key objects. By default, the SunPKCS11 provider only specifies mandatory PKCS#11 attributes when creating objects. For example, for RSA public keys it specifies the key type and algorithm (CKA_CLASS and CKA_KEY_TYPE) and the key values for RSA public keys (CKA_MODULUS and CKA_PUBLIC_EXPONENT). The PKCS#11 library you are using will assign implementation specific default values to the other attributes of an RSA public key, for example that the key can be used to encrypt and verify messages (CKA_ENCRYPT and CKA_VERIFY = true).

The `attributes` option can be used if you do not like the default values your PKCS#11 implementation assigns or if your PKCS#11 implementation does not support defaults and requires a value to be specified explicitly. Note that specifying attributes that your PKCS#11 implementation does not support or that are invalid for the type of key in question may cause the operation to fail at runtime.

The option can be specified zero or more times, the options are processed in the order specified in the configuration file as described below. The `attributes` option has the format:

```
attributes(operation, keytype, keyalgorithm) = {
  name1 = value1
  [...]
}
```

Valid values for `operation` are:

- `generate`, for keys generated via a `KeyPairGenerator` or `KeyGenerator`
- `import`, for keys created via a `KeyFactory` or `SecretKeyFactory`. This also applies to Java software keys automatically converted to PKCS#11 key objects when they are passed to the initialization method of a cryptographic operation, for example `Signature.initSign()`.
- `*`, for keys created using either a generate or a create operation.

Valid values for `keytype` are `CKO_PUBLIC_KEY`, `CKO_PRIVATE_KEY`, and `CKO_SECRET_KEY`, for public, private, and secret keys, respectively, and `*` to match any type of key.

Valid values for `keyalgorithm` are one of the `CKK_xxx` constants from the PKCS#11 specification, or `*` to match keys of any algorithm. The algorithms currently supported

by the SunPKCS11 provider include CKK_RSA, CKK_DSA, CKK_DH, CKK_AES, CKK_DES, CKK_DES3, CKK_RC4, CKK_BLOWFISH, CKK_GENERIC_SECRET, and CKK_EC.

The attribute names and values are specified as a list of one or more name-value pairs. `name` must be a `CKA_xxx` constant from the PKCS#11 specification, for example `CKA_SENSITIVE`. `value` can be one of the following:

- A boolean value, `true` or `false`
- An integer, in decimal form (default) or in hexadecimal form if it begins with `0x`.
- `null`, indicating that this attribute should *not* be specified when creating objects.

If the `attributes` option is specified multiple times, the entries are processed in the order specified with the attributes aggregated and later attributes overriding earlier ones. For example, consider the following configuration file excerpt:

```
attributes(*,CKO_PRIVATE_KEY,*) = {
  CKA_SIGN = true
}

attributes(*,CKO_PRIVATE_KEY,CKK_DH) = {
  CKA_SIGN = null
}

attributes(*,CKO_PRIVATE_KEY,CKK_RSA) = {
  CKA_DECRYPT = true
}
```

The first entry says to specify `CKA_SIGN = true` for all private keys. The second option overrides that with `null` for Diffie-Hellman private keys, so the `CKA_SIGN` attribute will not specified for them at all. Finally, the third option says to also specify `CKA_DECRYPT = true` for RSA private keys. That means RSA private keys will have both `CKA_SIGN = true` and `CKA_DECRYPT = true` set.

There is also a special form of the `attributes` option. You can write `attributes = compatibility` in the configuration file. That is a shortcut for a whole set of attribute statements. They are designed to provider maximum compatibility with existing Java applications, which may expect, for example, all key components to be accessible and secret keys to be usable for both encryption and decryption. The `compatibility` attributes line can be used together with other `attributes` lines, in which case the same aggregation and overriding rules apply as described earlier.

Accessing Network Security Services (NSS)

Network Security Services (NSS) is a set of open source security libraries whose crypto APIs are based on PKCS#11 but it includes special features that are outside of the PKCS#11 standard. The SunPKCS11 provider includes code to interact with these NSS specific features, including several NSS specific configuration directives, which are described below.

For best results, we recommend that you use the latest version of NSS available. It should be at least version 3.12.

The SunPKCS11 provider uses NSS specific code when any of the `nss` configuration directives described below are used. In that case, the regular configuration commands `library`, `slot`, and `slotListIndex` cannot be used.

Table 5-2 NSS Attributes and Values

Attribute	Value	Description
nssLibraryDirectory	directory containing the NSS and NSPR libraries	This is the full pathname of the directory containing the NSS and NSPR libraries. It must be specified unless NSS has already been loaded and initialized by another component running in the same process as the Java VM. **Note**: Depending on your platform, you may have to set the environment variable `LD_LIBRARY_PATH` (Solaris or Linux), `PATH` (Windows), or `DYLD_LIBRARY_PATH` (macOS) to include this directory to enable the operating system to locate the dependent libraries.
nssSecmodDirectory	directory containing the NSS DB files	The full pathname of the directory containing the NSS configuration and key information (`secmod.db`, `key3.db`, and `cert8.db`). This directive must be specified unless NSS has already been initialized by another component (see above) or NSS is used without database files as described below.
nssDbMode	one of `readWrite`, `readOnly`, and `noDb`	This directives determines how the NSS database is accessed. In read-write mode, full access is possible but only one process at a time should be accessing the databases. Read-only mode disallows modifications to the files. The noDb mode allows NSS to be used without database files purely as a cryptographic provider. It is not possible to create persistent keys using the PKCS11 KeyStore.

Table 5-2 (Cont.) NSS Attributes and Values

Attribute	Value	Description
nssModule	one of `keystore`, `crypto`, `fips`, and `trustanchors`	NSS makes its functionality available using several different libraries and slots. This directive determines which of these modules is accessed by this instance of SunPKCS11. The `crypto` module is the default in `noDb` mode. It supports crypto operations without login but no persistent keys. The `fips` module is the default if the NSS `secmod.db` has been set to FIPS-140 compliant mode. In this mode, NSS restricts the available algorithms and the PKCS#11 attributes with which keys can be created. The `keystore` module is the default in other configurations. It supports persistent keys using the PKCS11 KeyStore, which are stored in the NSS DB files. This module requires login. The `trustanchors` module enables access to NSS trust anchor certificates via the PKCS11 KeyStore, if `secmod.db` has been configured to include the trust anchor library.

Example 5-1 SunPKCS11 Configuration Files for NSS

NSS as a pure cryptography provider

```
name = NSScrypto
nssLibraryDirectory = /opt/tests/nss/lib
nssDbMode = noDb
attributes = compatibility
```

NSS as a FIPS 140 compliant crypto token

```
name = NSSfips
nssLibraryDirectory = /opt/tests/nss/lib
nssSecmodDirectory = /opt/tests/nss/fipsdb
nssModule = fips
```

Troubleshooting PKCS#11

There could be issues with PKCS#11 which requires debugging. To show debug info about Library, Slots, Token, and Mechanism, add `showInfo=true` in the SunPKCS11 provider configuration file, which is `<java-home>/conf/security/sunpkcs11-solaris.cfg` or the configuration file that you specified statically or dynamically as described in SunPKCS11 Configuration.

For additional debugging info, users can start or restart the Java processes with one of the following options:

- For general SunPKCS11 provider debugging info:

  ```
  -Djava.security.debug=sunpkcs11
  ```

- For PKCS#11 keystore specific debugging info:

  ```
  -Djava.security.debug=pkcs11keystore
  ```

Disabling PKCS#11 Providers and/or Individual PKCS#11 Mechanisms

As part of the troubleshooting process, it could be helpful to temporarily disable a PKCS#11 provider or the specific mechanism of a given provider.

Please note that disabling a PKCS#11 provider, is only a temporary measure. By disabling the PKCS#11 provider, the provider is no longer available which can cause applications to break or have a performance impact. Once the issue has been identified, only that specific mechanism should remain disabled.

Disabling PKCS#11 Providers

A PKCS#11 provider can be disabled by using one of the following methods:

1. Disable PKCS#11 for a single Java process. Start or restart the Java process with the following Java command line flag:

   ```
   -Dsun.security.pkcs11.enable-solaris=false
   ```

   > **Note:**
   >
   > This step is only applicable to the SunPKCS11 provider when backed by the default Solaris PKCS#11 provider file (`<java_home>/conf/security/sunpkcs11-solaris.cfg`).

2. Disable PKCS#11 for all Java processes run with a particular Java installation: This can be done dynamically by using the API (not shown in this section) or statically by editing the `<java_home>/conf/security/java.security` file and commenting out the SunPKCS11 security provider (do not forget to re-number the order of providers, if necessary) as shown below.

   ```
   #
   # List of providers and their preference orders (see above):
   ```

```
#
security.provider.1=SUN
security.provider.2=SunRsaSign
security.provider.3=SunEC
security.provider.4=SunJSSE
security.provider.5=SunJCE
security.provider.6=SunJGSS
security.provider.7=SunSASL
security.provider.8=XMLDSig
security.provider.9=SunPCSC
security.provider.10=JdkLDAP
security.provider.11=JdkSASL
security.provider.12=SunMSCAPI
#security.provider.13=SunPKCS11
```

Start or restart the Java processes being run on this installation of Java.

Disabling Specific Mechanisms

When an issue occurs in one of the mechanisms of PKCS#11, it can be resolved by disabling only that particular mechanism, rather than the entire PKCS#11 provider (do not forget to re-enable the PKCS#11 provider if it was disabled earlier).

Note:

To disable the PKCS#11 SecureRandom implementation only, you can add `SecureRandom` to the list of disabled mechanisms in the `<java-home>/conf/security/sunpkcs11-solaris.cfg` file:

```
name = Solaris

description = SunPKCS11 accessing Solaris Cryptographic Framework

library = /usr/lib/$ISA/libpkcs11.so

handleStartupErrors = ignoreAll

# Use the X9.63 encoding for EC points (do not wrap in an ASN.1
OctetString).
useEcX963Encoding = true

attributes = compatibility

disabledMechanisms = {
  CKM_DSA_KEY_PAIR_GEN
  SecureRandom
}
```

Application Developers

Java applications can use the existing JCA and JCE APIs to access PKCS#11 tokens through the SunPKCS11 provider.

Token Login

You can login to the keystore using a Personal Identification Number and perform PKCS#11 operations.

Certain PKCS#11 operations, such as accessing private keys, require a login using a Personal Identification Number, or PIN, before the operations can proceed. The most common type of operations that require login are those that deal with keys on the token. In a Java application, such operations often involve first loading the keystore. When accessing the PKCS#11 token as a keystore via the `java.security.KeyStore` class, you can supply the PIN in the password input parameter to the `load` method. The PIN will then be used by the SunPKCS11 provider for logging into the token. Here is an example.

```
char[] pin = ...;
KeyStore ks = KeyStore.getInstance("PKCS11");
ks.load(null, pin);
```

This is fine for an application that treats PKCS#11 tokens as static keystores. For an application that wants to accommodate PKCS#11 tokens more dynamically, such as smartcards being inserted and removed, you can use the new `KeyStore.Builder` class. Here is an example of how to initialize the builder for a PKCS#11 keystore with a callback handler.

```
KeyStore.CallbackHandlerProtection chp =
    new KeyStore.CallbackHandlerProtection(new MyGuiCallbackHandler());
KeyStore.Builder builder =
    KeyStore.Builder.newInstance("PKCS11", null, chp);
```

For the SunPKCS11 provider, the callback handler must be able to satisfy a `PasswordCallback`, which is used to prompt the user for the PIN. Whenever the application needs access to the keystore, it uses the builder as follows.

```
KeyStore ks = builder.getKeyStore();
Key key = ks.getKey(alias, null);
```

The builder will prompt for a password as needed using the previously configured callback handler. The builder will prompt for a password only for the initial access. If the user of the application continues using the same Smartcard, the user will not be prompted again. If the user removes and inserts a different smartcard, the builder will prompt for a password for the new card.

Depending on the PKCS#11 token, there may be non-key-related operations that also require token login. Applications that use such operations can use the `java.security.AuthProvider` class. The `AuthProvider` class extends from

`java.security.Provider` and defines methods to perform login and logout operations on a provider, as well as to set a callback handler for the provider to use.

For the SunPKCS11 provider, the callback handler must be able to satisfy a `PasswordCallback`, which is used to prompt the user for the PIN.

Here is an example of how an application might use an `AuthProvider` to log into the token. (Note that you must configure the SunPKCS11 provider before using it.)

```
    Provider p = Security.getProvider("SunPKCS11");
    AuthProvider aprov = (AuthProvider)p.configure(<provider configuration
file>);
    aprov.login(subject, new MyGuiCallbackHandler());
```

Token Keys

Java `Key` objects may or may not contain actual key material.

- A software `Key` object does contain the actual key material and allows access to that material.
- An unextractable key on a secure token (such as a smartcard) is represented by a Java `Key` object that does not contain the actual key material. The `Key` object only contains a reference to the actual key.

Applications and providers must use the correct interfaces to represent these different types of `Key` objects. Software `Key` objects (or any `Key` object that has access to the actual key material) should implement the interfaces in the `java.security.interfaces` and `javax.crypto.interfaces` packages (such as `DSAPrivateKey`). Key objects representing unextractable token keys should only implement the relevant generic interfaces in the `java.security` and `javax.crypto` packages (`PrivateKey`, `PublicKey`, or `SecretKey`). Identification of the algorithm of a key should be performed using the `Key.getAlgorithm()` method.

Note that a `Key` object for an unextractable token key can only be used by the provider associated with that token.

Delayed Provider Selection

Java cryptography `getInstance()` methods, such as `Cipher.getInstance("AES")`, return the implementation from the first provider that implemented the requested algorithm. However, the JDK delays the selection of the provider until the relevant initialization method is called. The initialization method accepts a `Key` object and can determine at that point which provider can accept the specified `Key` object. This ensures that the selected provider can use the specified `Key` object. (If an application attempts to use a `Key` object for an unextractable token key with a provider that only accepts software key objects, then the provider throws an `InvalidKeyException`. This is an issue for the `Cipher`, `KeyAgreement`, `Mac`, and `Signature` classes.) The following represents the affected initialization methods.

- `Cipher.init(..., Key key, ...)`
- `KeyAgreement.init(Key key, ...)`
- `Mac.init(Key key, ...)`
- `Signature.initSign(PrivateKey privateKey)`

Furthermore, if an application calls the initialization method multiple times (each time with a different key, for example), the proper provider for the given key is selected each time. In other words, a different provider may be selected for each initialization call.

Although this delayed provider selection is hidden from the application, it does affect the behavior of the `getProvider()` method for `Cipher`, `KeyAgreement`, `Mac`, and `Signature`. If `getProvider()` is called *before* the initialization operation has occurred (and therefore before provider selection has occurred), then the first provider that supports the requested algorithm is returned. This may not be the same provider as the one selected *after* the initialization method is called. If `getProvider()` is called *after* the initialization operation has occurred, then the actual selected provider is returned. It is recommended that applications only call `getProvider()` after they have called the relevant initialization method.

In addition to `getProvider()`, the following additional methods are similarly affected.

- `Cipher.getBlockSize`
- `Cipher.getExcemptionMechanism`
- `Cipher.getIV`
- `Cipher.getOutputSize`
- `Cipher.getParameters`
- `Mac.getMacLength`
- `Signature.getParameters`
- `Signature.setParameter`

JAAS KeyStoreLoginModule

The JDK comes with a JAAS keystore login module, `KeyStoreLoginModule`, that allows an application to authenticate using its identity in a specified keystore. After authentication, the application would acquire its principal and credentials information (certificate and private key) from the keystore. By using this login module and configuring it to use a PKCS#11 token as a keystore, the application can acquire this information from a PKCS#11 token.

Use the following options to configure the `KeyStoreLoginModule` to use a PKCS#11 token as the keystore.

- `keyStoreURL="NONE"`
- `keyStoreType="PKCS11"`
- `keyStorePasswordURL=`*`some_pin_url`*

where

some_pin_url
The location of the PIN. If the `keyStorePasswordURL` option is omitted, then the login module will get the PIN via the application's callback handler, supplying it with a

PasswordCallback . Here is an example of a configuration file that uses a PKCS#11 token as a keystore.

```
other {
    com.sun.security.auth.module.KeyStoreLoginModule required
        keyStoreURL="NONE"
  keyStoreType="PKCS11"
  keyStorePasswordURL="file:/home/joe/scpin";
};
```

If more than one SunPKCS11 provider has been configured dynamically or in the `java.security` security properties file, you can use the `keyStoreProvider` option to target a specific provider instance. The argument to this option is the name of the provider. For the SunPKCS11 provider, the provider name is of the form `SunPKCS11-`*`TokenName`*, where *`TokenName`* is the name suffix that the provider instance has been configured with, as detailed in Table 5-1. For example, the following configuration file names the PKCS#11 provider instance with name suffix `SmartCard`.

```
other {
    com.sun.security.auth.module.KeyStoreLoginModule required
        keyStoreURL="NONE"
  keyStoreType="PKCS11"
  keyStorePasswordURL="file:/home/joe/scpin"
  keyStoreProvider="SunPKCS11-SmartCard";
};
```

Some PKCS#11 tokens support login via a protected authentication path. For example, a smartcard may have a dedicated PIN-pad to enter the pin. Biometric devices will also have their own means to obtain authentication information. If the PKCS#11 token has a protected authentication path, then use the `protected=true` option and omit the `keyStorePasswordURL` option. Here is an example of a configuration file for such a token.

```
other {
    com.sun.security.auth.module.KeyStoreLoginModule required
        keyStoreURL="NONE"
  keyStoreType="PKCS11"
  protected=true;
};
```

Tokens as JSSE Keystore and Trust Stores

To use PKCS#11 tokens as JSSE keystores or trust stores, the JSSE application can use the APIs described in Token Login to instantiate a `KeyStore` that is backed by a PKCS#11 token and pass it to its key manager and trust manager. The JSSE application will then have access to the keys on the token.

JSSE also supports configuring the use of keystores and trust stores via system properties, as described in the Java Secure Socket Extension (JSSE) Reference Guide. To use a PKCS#11 token as a keystore or trust store, set the `javax.net.ssl.keyStoreType` and `javax.net.ssl.trustStoreType` system properties, respectively, to "PKCS11", and set the `javax.net.ssl.keyStore` and

`javax.net.ssl.trustStore` system properties, respectively, to `NONE`. To specify the use of a specific provider instance, use the `javax.net.ssl.keyStoreProvider` and `javax.net.ssl.trustStoreProvider` system properties (for example, "SunPKCS11-SmartCard").

Using keytool and jarsigner with PKCS#11 Tokens

If the SunPKCS11 provider has been configured in the `java.security` security properties file (located in the `$JAVA_HOME/conf/security` directory of the Java runtime), then keytool and jarsigner can be used to operate on the PKCS#11 token by specifying the following options.

- `-keystore NONE`
- `-storetype PKCS11`

Here an example of a command to list the contents of the configured PKCS#11 token.

```
keytool -keystore NONE -storetype PKCS11 -list
```

The PIN can be specified using the `-storepass` option. If none has been specified, then `keytool` and `jarsigner` will prompt for the token PIN. If the token has a protected authentication path (such as a dedicated PIN-pad or a biometric reader), then the `-protected` option must be specified, and no password options can be specified.

If more than one SunPKCS11 provider has been configured in the `java.security` security properties file, you can use the `-providerName` option to target a specific provider instance. The argument to this option is the name of the provider.

- `-providerName` *`providerName`*

For the SunPKCS11 provider, *`providerName`* is of the form `SunPKCS11-`*`TokenName`* where:

TokenName
The name suffix that the provider instance has been configured with, as detailed in Table 5-1. For example, the following command lists the contents of the PKCS#11 keystore provider instance with name suffix `SmartCard`.

```
keytool -keystore NONE -storetype PKCS11 \
        -providerName SunPKCS11-SmartCard \
        -list
```

If the SunPKCS11 provider has not been configured in the `java.security` security properties file, you can use the following options to instruct `keytool` and `jarsigner` to install the provider dynamically.

- `-providerClass sun.security.pkcs11.SunPKCS11`
- `-providerArg` *`ConfigFilePath`*

ConfigFilePath
The path to the token configuration file. Here is an example of a command to list a PKCS#11 keystore when the SunPKCS11 provider has not been configured in the `java.security` file.

```
keytool -keystore NONE -storetype PKCS11 \
        -providerClass sun.security.pkcs11.SunPKCS11 \
        -providerArg /foo/bar/token.config \
        -list
```

> **Note:**
>
> Sometimes the hardware token is too small to store the certificates. You can use the `jarsigner` tool's `-certchain` option to load them from an external file.

Signing JAR File with jdk.security.jarsigner API and PKCS#11 Token

To sign a JAR file with the `jdk.security.jarsigner` API and a PKCS#11 token as a keystore, follow these steps:

1. Access the PKCS#11 token's keystore as described in Token Login:

```
char[] pin = ...;
KeyStore ks = KeyStore.getInstance("PKCS11");
ks.load(null, pin);
KeyStore ks = builder.getKeyStore();
Key key = ks.getKey(alias, null);
```

2. Create a `JarSigner` object with the `JarSigner.Builder(Keystore.PrivateKeyEntry)` constructor:

```
JarSigner mySigner = JarSigner.Builder(key).build();
```

3. Sign the JAR file:

```
try (ZipFile in = new ZipFile(inputFile);
    FileOutputStream out = new FileOutputStream(outputFile)) {
    mySigner.sign(in, out);
}
```

Keystore Entry Syntax in Policy File

The `keystore` entry in the default policy implementation has the following syntax, which accommodates a PIN and multiple PKCS#11 provider instances:

```
keystore "some_keystore_url", "keystore_type", "keystore_provider";
keystorePasswordURL "some_password_url";
```

Where

keystore_provider
The keystore provider name (for example, `"SunPKCS11-SmartCard"`).

some_password_url
A URL pointing to the location of the token PIN. Both *keystore_provider* and the *keystorePasswordURL* line are optional. If *keystore_provider* has not been specified, then the first configured provider that supports the specified keystore type is used. If the `keystorePasswordURL` line has not been specified, then no password is used.

See Default Policy Implementation and Policy File Syntax.

Example 5-2 Keystore Policy Entry for a PKCS#11 Token

The following is an example keystore policy entry for a PKCS#11 token:

```
keystore "NONE", "PKCS11", "SunPKCS11-SmartCard";
keystorePasswordURL "file:/foo/bar/passwordFile";
```

Provider Developers

The `java.security.Provider` class enables provider developers to more easily support PKCS#11 tokens and cryptographic services through provider services and parameter support.

See Example Provider for an example of a simple provider designed to demonstrate provider services and parameter support.

Provider Services

For each service implemented by the provider, there must be a property whose name is the type of service (`Cipher`, `Signature`, etc), followed by a period and the name of the algorithm to which the service applies. The property value must specify the fully qualified name of the class implementing the service. Here is an example of a provider setting `KeyAgreement.DiffieHellman` property to have the value `com.sun.crypto.provider.DHKeyAgreement`.

```
put("KeyAgreement.DiffieHellman", "com.sun.crypto.provider.DHKeyAgreement")
```

The public static nested class `Provider.Service` encapsulates the properties of a provider service (including its type, attributes, algorithm name, and algorithm aliases). Providers can instantiate `Provider.Service` objects and register them by calling the `Provider.putService()` method. This is equivalent to creating a `Property` entry and calling the `Provider.put()` method. Note that legacy `Property` entries registered via `Provider.put` are still supported.

Here is an example of a provider creating a `Service` object with the `KeyAgreement` type, for the `DiffieHellman` algorithm, implemented by the class `com.sun.crypto.provider.DHKeyAgreement`.

```
Service s = new Service(this, "KeyAgreement", "DiffieHellman",
    "com.sun.crypto.provider.DHKeyAgreement", null, null);
putService(s);
```

Using `Provider.Service` objects instead of legacy `Property` entries has a couple of major benefits. One benefit is that it allows the provider to have greater flexibility when Instantiating Engine Classes. Another benefit is that it allows the provider to test Parameter Support. These features are discussed in detail next.

Instantiating Engine Classes

By default, the Java Cryptography framework looks up the provider property for a particular service and directly instantiates the engine class registered for that property. A provider can to override this behavior and instantiate the engine class for the requested service itself.

To override the default behavior, the provider overrides the `Provider.Service.newInstance()` method to add its custom behavior. For example, the provider might call a custom constructor, or might perform initialization using information not accessible outside the provider (or that are only known by the provider).

Parameter Support

The Java Cryptography framework may attempt a fast check to determine whether a provider's service implementation can use an application-specified parameter. To perform this fast check, the framework calls `Provider.Service.supportsParameter()`.

The framework relies on this fast test during delayed provider selection (see Delayed Provider Selection). When an application invokes an initialization method and passes it a `Key` object, the framework asks an underlying provider whether it supports the object by calling its `Service.supportsParameter()` method. If `supportsParameter()` returns `false`, the framework can immediately remove that provider from consideration. If `supportsParameter()` returns `true`, the framework passes the `Key` object to that provider's initialization engine class implementation. A provider that requires software `Key` objects should override this method to return `false` when it is passed non-software keys. Likewise, a provider for a PKCS#11 token that contains unextractable keys should only return `true` for `Key` objects that it created, and which therefore correspond to the keys on its respective token.

> **Note:**
>
> The default implementation of `supportsParameter()` returns `true`. This allows existing providers to work without modification. However, because of this lenient default implementation, the framework must be prepared to catch exceptions thrown by providers that reject the `Key` object inside their initialization engine class implementations. The framework treats these cases the same as when `supportsParameter()` returns `false`.

Parameter Support

The Java Cryptography framework may attempt a fast check to determine whether a provider's service implementation can use an application-specified parameter. To perform this fast check, the framework calls `Provider.Service.supportsParameter()`.

The framework relies on this fast test during delayed provider selection (see Delayed Provider Selection). When an application invokes an initialization method and passes it a `Key` object, the framework asks an underlying provider whether it supports the object

by calling its `Service.supportsParameter()` method. If `supportsParameter()` returns `false`, the framework can immediately remove that provider from consideration. If `supportsParameter()` returns `true`, the framework passes the `Key` object to that provider's initialization engine class implementation. A provider that requires software `Key` objects should override this method to return `false` when it is passed non-software keys. Likewise, a provider for a PKCS#11 token that contains unextractable keys should only return `true` for `Key` objects that it created, and which therefore correspond to the keys on its respective token.

> **Note:**
>
> The default implementation of `supportsParameter()` returns `true`. This allows existing providers to work without modification. However, because of this lenient default implementation, the framework must be prepared to catch exceptions thrown by providers that reject the `Key` object inside their initialization engine class implementations. The framework treats these cases the same as when `supportsParameter()` returns `false`.

SunPKCS11 Provider Supported Algorithms

Table 5-3 lists the Java algorithms supported by the SunPKCS11 provider and corresponding PKCS#11 mechanisms needed to support them. When multiple mechanisms are listed, they are given in the order of preference and any one of them is sufficient.

> **Note:**
>
> SunPKCS11 can be instructed to ignore mechanisms by using the `disabledMechanisms` and `enabledMechanisms` configuration directives (see SunPKCS11 Configuration).

For Elliptic Curve mechanisms, the SunPKCS11 provider will only use keys that use the `namedCurve` choice as encoding for the parameters and only allow the uncompressed point format. The SunPKCS11 provider assumes that a token supports all standard named domain parameters.

> **Note:**
>
> For Elliptic Curve (EC) names, the SunPKCS11 provider supports any EC name that the SunEC provider supports as long as the token supports it; see Table 4-27 in The SunEC Provider.

Table 5-3 Java Algorithms Supported by the SunPKCS11 Provider

Java Algorithm	PKCS#11 Mechanisms
Signature.MD2withRSA	CKM_MD2_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509

Table 5-3 (Cont.) Java Algorithms Supported by the SunPKCS11 Provider

Java Algorithm	PKCS#11 Mechanisms
Signature.MD5withRSA	CKM_MD5_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA1withRSA	CKM_SHA1_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA224withRSA	CKM_SHA224_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA256withRSA	CKM_SHA256_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA384withRSA	CKM_SHA384_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA512withRSA	CKM_SHA512_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA1withDSA	CKM_DSA_SHA1, CKM_DSA
Signature.NONEwithDSA	CKM_DSA
Signature.SHA1withECDSA	CKM_ECDSA_SHA1, CKM_ECDSA
Signature.SHA224withECDSA	CKM_ECDSA
Signature.SHA256withECDSA	CKM_ECDSA
Signature.SHA384withECDSA	CKM_ECDSA
Signature.SHA512withECDSA	CKM_ECDSA
Signature.NONEwithECDSA	CKM_ECDSA
Cipher.RSA/ECB/NoPadding	CKM_RSA_X_509
Cipher.RSA/ECB/PKCS1Padding	CKM_RSA_PKCS
Cipher.ARCFOUR	CKM_RC4
Cipher.DES/CBC/NoPadding	CKM_DES_CBC
Cipher.DES/CBC/PKCS5Padding	CKM_DES_CBC_PAD, CKM_DES_CBC
Cipher.DES/ECB/NoPadding	CKM_DES_ECB
Cipher.DES/ECB/PKCS5Padding	CKM_DES_ECB
Cipher.DESede/CBC/NoPadding	CKM_DES3_CBC
Cipher.DESede/CBC/PKCS5Padding	CKM_DES3_CBC_PAD, CKM_DES3_CBC
Cipher.DESede/ECB/NoPadding	CKM_DES3_ECB
Cipher.DESede/ECB/PKCS5Padding	CKM_DES3_ECB
Cipher.AES/CBC/NoPadding	CKM_AES_CBC
Cipher.AES/CBC/PKCS5Padding	CKM_AES_CBC_PAD, CKM_AES_CBC
Cipher.Blowfish/CBC/NoPadding	CKM_BLOWFISH_CBC
Cipher.Blowfish/CBC/PKCS5Padding	CKM_BLOWFISH_CBC
Cipher.AES/CTR/NoPadding	CKM_AES_CTR
Cipher.AES/ECB/NoPadding	CKM_AES_ECB
Cipher.AES/ECB/PKCS5Padding	CKM_AES_ECB
Cipher.AES_128/CBC/NoPadding	CKM_AES_CBC
Cipher.AES_128/ECB/NoPadding	CKM_AES_ECB
Cipher.AES_192/CBC/NoPadding	CKM_AES_CBC
Cipher.AES_192/ECB/NoPadding	CKM_AES_ECB
Cipher.AES_256/CBC/NoPadding	CKM_AES_CBC

Table 5-3 (Cont.) Java Algorithms Supported by the SunPKCS11 Provider

Java Algorithm	PKCS#11 Mechanisms
Cipher.AES_256/ECB/NoPadding	CKM_AES_ECB
KeyAgreement.ECDH	CKM_ECDH1_DERIVE
KeyAgreement.DiffieHellman	CKM_DH_PKCS_DERIVE
KeyPairGenerator.RSA	CKM_RSA_PKCS_KEY_PAIR_GEN
KeyPairGenerator.DSA	CKM_DSA_KEY_PAIR_GEN
KeyPairGenerator.EC	CKM_EC_KEY_PAIR_GEN
KeyPairGenerator.DiffieHellman	CKM_DH_PKCS_KEY_PAIR_GEN
KeyGenerator.ARCFOUR	CKM_RC4_KEY_GEN
KeyGenerator.DES	CKM_DES_KEY_GEN
KeyGenerator.DESede	CKM_DES3_KEY_GEN
KeyGenerator.AES	CKM_AES_KEY_GEN
KeyGenerator.Blowfish	CKM_BLOWFISH_KEY_GEN
Mac.HmacMD5	CKM_MD5_HMAC
Mac.HmacSHA1	CKM_SHA_1_HMAC
Mac.HmacSHA224	CKM_SHA224_HMAC
Mac.HmacSHA256	CKM_SHA256_HMAC
Mac.HmacSHA384	CKM_SHA384_HMAC
Mac.HmacSHA512	CKM_SHA512_HMAC
MessageDigest.MD2	CKM_MD2
MessageDigest.MD5	CKM_MD5
MessageDigest.SHA1	CKM_SHA_1
MessageDigest.SHA-224	CKM_SHA224
MessageDigest.SHA-256	CKM_SHA256
MessageDigest.SHA-384	CKM_SHA384
MessageDigest.SHA-512	CKM_SHA512
KeyFactory.RSA	Any supported RSA mechanism
KeyFactory.DSA	Any supported DSA mechanism
KeyFactory.EC	Any supported EC mechanism
KeyFactory.DiffieHellman	Any supported Diffie-Hellman mechanism
SecretKeyFactory.ARCFOUR	CKM_RC4
SecretKeyFactory.DES	CKM_DES_CBC
SecretKeyFactory.DESede	CKM_DES3_CBC
SecretKeyFactory.AES	CKM_AES_CBC
SecretKeyFactory.Blowfish	CKM_BLOWFISH_CBC
SecureRandom.PKCS11	CK_TOKEN_INFO has the CKF_RNG bit set
KeyStore.PKCS11	Always available

SunPKCS11 Provider KeyStore Requirements

The following describes the requirements placed by the SunPKCS11 provider's KeyStore implementation on the underlying native PKCS#11 library.

> **Note:**
>
> Changes may be made in future releases to maximize interoperability with as many existing PKCS#11 libraries as possible.

Read-Only Access

To map existing objects stored on a PKCS#11 token to KeyStore entries, the SunPKCS11 provider's KeyStore implementation performs the following operations.

1. A search for all private key objects on the token is performed by calling `C_FindObjects[Init|Final]`. The search template includes the following attributes:
   - CKA_TOKEN = true
   - CKA_CLASS = CKO_PRIVATE_KEY
2. A search for all certificate objects on the token is performed by calling `C_FindObjects[Init|Final]`. The search template includes the following attributes:
   - CKA_TOKEN = true
   - CKA_CLASS = CKO_CERTIFICATE
3. Each private key object is matched with its corresponding certificate by retrieving their respective CKA_ID attributes. A matching pair must share the same unique CKA_ID.
   For each matching pair, the certificate chain is built by following the issuer->subject path. From the end entity certificate, a call for `C_FindObjects[Init|Final]` is made with a search template that includes the following attributes:
   - CKA_TOKEN = true
   - CKA_CLASS = CKO_CERTIFICATE
   - CKA_SUBJECT = [DN of certificate issuer]

   This search is continued until either no certificate for the issuer is found, or until a self-signed certificate is found. If more than one certificate is found the first one is used.

   Once a private key and certificate have been matched (and its certificate chain built), the information is stored in a private key entry with the CKA_LABEL value from end entity certificate as the KeyStore alias.

   If the end entity certificate has no CKA_LABEL, then the alias is derived from the CKA_ID. If the CKA_ID can be determined to consist exclusively of printable characters, then a String alias is created by decoding the CKA_ID bytes using the UTF-8 charset. Otherwise, a hex String alias is created from the CKA_ID bytes ("0xFFFF...", for example).

If multiple certificates share the same CKA_LABEL, then the alias is derived from the CKA_LABEL plus the end entity certificate issuer and serial number (`"MyCert/CN=foobar/1234"`, for example).

4. Each certificate not part of a private key entry (as the end entity certificate) is checked whether it is trusted. If the CKA_TRUSTED attribute is true, then a KeyStore trusted certificate entry is created with the CKA_LABEL value as the KeyStore alias. If the certificate has no CKA_LABEL, or if multiple certificates share the same CKA_LABEL, then the alias is derived as described above.
   If the CKA_TRUSTED attribute is not supported then no trusted certificate entries are created.
5. Any private key or certificate object not part of a private key entry or trusted certificate entry is ignored.
6. A search for all secret key objects on the token is performed by calling `C_FindObjects[Init|Final]`. The search template includes the following attributes:
   - CKA_TOKEN = true
   - CKA_CLASS = CKO_SECRET_KEY

   A KeyStore secret key entry is created for each secret key object, with the CKA_LABEL value as the KeyStore alias. Each secret key object must have a unique CKA_LABEL.

Write Access

To create new KeyStore entries on a PKCS#11 token to KeyStore entries, the SunPKCS11 provider's KeyStore implementation performs the following operations.

1. When creating a KeyStore entry (during KeyStore.setEntry, for example), C_CreateObject is called with `CKA_TOKEN=true` to create token objects for the respective entry contents.
   Private key objects are stored with `CKA_PRIVATE=true`. The KeyStore alias (UTF8-encoded) is set as the CKA_ID for both the private key and the corresponding end entity certificate. The KeyStore alias is also set as the CKA_LABEL for the end entity certificate object.

   Each certificate in a private key entry's chain is also stored. The CKA_LABEL is not set for CA certificates. If a CA certificate is already in the token, a duplicate is not stored.

   Secret key objects are stored with `CKA_PRIVATE=true`. The KeyStore alias is set as the CKA_LABEL.
2. If an attempt is made to convert a session object to a token object (for example, if KeyStore.setEntry is called and the private key object in the specified entry is a session object), then C_CopyObject is called with `CKA_TOKEN=true`.
3. If multiple certificates in the token are found to share the same CKA_LABEL, then the write capabilities to the token are disabled.
4. Since the PKCS#11 specification does not allow regular applications to set `CKA_TRUSTED=true` (only token initialization applications may do so), trusted certificate entries can not be created.

Miscellaneous

In addition to the searches listed above, the following searches may be used by the SunPKCS11 provider's KeyStore implementation to perform internal functions. Specifically, `C_FindObjects[Init|Final]` may be called with any of the following attribute templates:

- ```
  CKA_TOKEN    true
  CKA_CLASS    CKO_CERTIFICATE
  CKA_SUBJECT  [subject DN]
  ```
- ```
  CKA_TOKEN    true
  CKA_CLASS    CKO_SECRET_KEY
  CKA_LABEL    [label]
  ```
- ```
  CKA_TOKEN    true
  CKA_CLASS    CKO_CERTIFICATE or CKO_PRIVATE_KEY
  CKA_ID       [cka_id]
  ```

Example Provider

The following is an example of a simple provider that demonstrates features of the `Provider` class.

```
package com.foo;

import java.io.*;
import java.lang.reflect.*;
import java.security.*;
import javax.crypto.*;

/**
 * Example provider that demonstrates some Provider class features.
 *
 *  . implement multiple different algorithms in a single class.
 *    Previously each algorithm needed to be implemented in a separate
class
 *    (e.g. one for SHA-256, one for SHA-384, etc.)
 *
 *  . multiple concurrent instances of the provider frontend class each
 *    associated with a different backend.
 *
 *  . it uses "unextractable" keys and lets the framework know which key
 *    objects it can and cannot support
 *
 * Note that this is only a simple example provider designed to demonstrate
 * several of the new features.  It is not explicitly designed for
efficiency.
 */
public final class ExampleProvider extends Provider {
```
```

```
    // reference to the crypto backend that implements all the algorithms
    final CryptoBackend cryptoBackend;

    public ExampleProvider(String name, CryptoBackend cryptoBackend) {
        super(name, 1.0, "JCA/JCE provider for " + name);
        this.cryptoBackend = cryptoBackend;
        // register the algorithms we support (SHA-256, SHA-384, DESede,
and AES)
        putService(new MyService
            (this, "MessageDigest", "SHA-256",
"com.foo.ExampleProvider$MyMessageDigest"));
        putService(new MyService
            (this, "MessageDigest", "SHA-384",
"com.foo.ExampleProvider$MyMessageDigest"));
        putService(new MyCipherService
            (this, "Cipher", "DES", "com.foo.ExampleProvider$MyCipher"));
        putService(new MyCipherService
            (this, "Cipher", "AES", "com.foo.ExampleProvider$MyCipher"));
    }

    // the API of our fictitious crypto backend
    static abstract class CryptoBackend {
        abstract byte[] digest(String algorithm, byte[] data);
        abstract byte[] encrypt(String algorithm, KeyHandle key, byte[]
data);
        abstract byte[] decrypt(String algorithm, KeyHandle key, byte[]
data);
        abstract KeyHandle createKey(String algorithm, byte[] keyData);
    }

    // the shell of the representation the crypto backend uses for keys
    private static final class KeyHandle {
        // fill in code
    }

    // we have our own ServiceDescription implementation that overrides
newInstance()
    // that calls the (Provider, String) constructor instead of the no-
args constructor
    private static class MyService extends Service {

        private static final Class[] paramTypes = {Provider.class,
String.class};

        MyService(Provider provider, String type, String algorithm,
                String className) {
            super(provider, type, algorithm, className, null, null);
        }

        public Object newInstance(Object param) throws
NoSuchAlgorithmException {
            try {
                // get the Class object for the implementation class
                Class clazz;
                Provider provider = getProvider();
```

```
                ClassLoader loader = provider.getClass().getClassLoader();
                if (loader == null) {
                    clazz = Class.forName(getClassName());
                } else {
                    clazz = loader.loadClass(getClassName());
                }
                // fetch the (Provider, String) constructor
                Constructor cons = clazz.getConstructor(paramTypes);
                // invoke constructor and return the SPI object
                Object obj = cons.newInstance(new Object[] {provider,
getAlgorithm()});
                return obj;
            } catch (Exception e) {
                throw new NoSuchAlgorithmException("Could not instantiate
service", e);
            }
        }
    }

    // custom ServiceDescription class for Cipher objects. See
supportsParameter() below
    private static class MyCipherService extends MyService {
        MyCipherService(Provider provider, String type, String algorithm,
                String className) {
            super(provider, type, algorithm, className);
        }
        // we override supportsParameter() to let the framework know which
        // keys we can support. We support instances of MySecretKey, if
they
        // are stored in our provider backend, plus SecretKeys with a RAW
encoding.
        public boolean supportsParameter(Object obj) {
            if (obj instanceof SecretKey == false) {
                return false;
            }
            SecretKey key = (SecretKey)obj;
            if (key.getAlgorithm().equals(getAlgorithm()) == false) {
                return false;
            }
            if (key instanceof MySecretKey) {
                MySecretKey myKey = (MySecretKey)key;
                return myKey.provider == getProvider();
            } else {
                return "RAW".equals(key.getFormat());
            }
        }
    }

    // our generic MessageDigest implementation. It implements all digest
    // algorithms in a single class. We only implement the bare minimum
    // of MessageDigestSpi methods
    private static final class MyMessageDigest extends MessageDigestSpi {
        private final ExampleProvider provider;
        private final String algorithm;
        private ByteArrayOutputStream buffer;
```

```
        MyMessageDigest(Provider provider, String algorithm) {
            super();
            this.provider = (ExampleProvider)provider;
            this.algorithm = algorithm;
            engineReset();
        }
        protected void engineReset() {
            buffer = new ByteArrayOutputStream();
        }
        protected void engineUpdate(byte b) {
            buffer.write(b);
        }
        protected void engineUpdate(byte[] b, int ofs, int len) {
            buffer.write(b, ofs, len);
        }
        protected byte[] engineDigest() {
            byte[] data = buffer.toByteArray();
            byte[] digest = provider.cryptoBackend.digest(algorithm, data);
            engineReset();
            return digest;
        }
    }

    // our generic Cipher implementation, only partially complete. It
implements
    // all cipher algorithms in a single class. We implement only as many
of the
    // CipherSpi methods as required to show how it could work
    private static abstract class MyCipher extends CipherSpi {
        private final ExampleProvider provider;
        private final String algorithm;
        private int opmode;
        private MySecretKey myKey;
        private ByteArrayOutputStream buffer;
        MyCipher(Provider provider, String algorithm) {
            super();
            this.provider = (ExampleProvider)provider;
            this.algorithm = algorithm;
        }
        protected void engineInit(int opmode, Key key, SecureRandom random)
                throws InvalidKeyException {
            this.opmode = opmode;
            myKey = MySecretKey.getKey(provider, algorithm, key);
            if (myKey == null) {
                throw new InvalidKeyException();
            }
            buffer = new ByteArrayOutputStream();
        }
        protected byte[] engineUpdate(byte[] b, int ofs, int len) {
            buffer.write(b, ofs, len);
            return new byte[0];
        }
        protected int engineUpdate(byte[] b, int ofs, int len, byte[] out,
int outOfs) {
            buffer.write(b, ofs, len);
```

```
            return 0;
        }
        protected byte[] engineDoFinal(byte[] b, int ofs, int len) {
            buffer.write(b, ofs, len);
            byte[] in = buffer.toByteArray();
            byte[] out;
            if (opmode == Cipher.ENCRYPT_MODE) {
                out = provider.cryptoBackend.encrypt(algorithm,
myKey.handle, in);
            } else {
                out = provider.cryptoBackend.decrypt(algorithm,
myKey.handle, in);
            }
            buffer = new ByteArrayOutputStream();
            return out;
        }
        // code for remaining CipherSpi methods goes here
    }

    // our SecretKey implementation. All our keys are stored in our crypto
    // backend, we only have an opaque handle available. There is no
    // encoded form of these keys.
    private static final class MySecretKey implements SecretKey {

        final String algorithm;
        final Provider provider;
        final KeyHandle handle;

        MySecretKey(Provider provider, String algorithm, KeyHandle handle)
{
            super();
            this.provider = provider;
            this.algorithm = algorithm;
            this.handle = handle;
        }
        public String getAlgorithm() {
            return algorithm;
        }
        public String getFormat() {
            return null; // this key has no encoded form
        }
        public byte[] getEncoded() {
            return null; // this key has no encoded form
        }
        // Convert the given key to a key of the specified provider, if
possible
        static MySecretKey getKey(ExampleProvider provider, String
algorithm, Key key) {
            if (key instanceof SecretKey == false) {
                return null;
            }
            // algorithm name must match
            if (!key.getAlgorithm().equals(algorithm)) {
                return null;
            }
```

```
            // if key is already an instance of MySecretKey and is stored
            // on this provider, return it right away
            if (key instanceof MySecretKey) {
                MySecretKey myKey = (MySecretKey)key;
                if (myKey.provider == provider) {
                    return myKey;
                }
            }
            // otherwise, if the input key has a RAW encoding, convert it
            if (!"RAW".equals(key.getFormat())) {
                return null;
            }
            byte[] encoded = key.getEncoded();
            KeyHandle handle = provider.cryptoBackend.createKey(algorithm,
encoded);
            return new MySecretKey(provider, algorithm, handle);
        }
    }
}
```

6

Java Authentication and Authorization Service (JAAS)

Java Authentication and Authorization Service (JAAS) Reference Guide describes Java Authentication and Authorization Service (JAAS), which enables you to authenticate users and securely determine who is currently executing Java code, and authorize users to ensure that they have the access control rights, or permissions, required to do the actions performed.

JAAS Tutorials provides tutorials about Java Authentication and Authorization Service (JAAS) authentication and authorization.

Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide shows you how to implement the `LoginModule` interface, which you plug into an application to provide a particular type of authentication.

Java Authentication and Authorization Service (JAAS) Reference Guide

JAAS can be used for two purposes:

- for *authentication* of users, to reliably and securely determine who is currently executing Java code, regardless of whether the code is running as an application, an applet, a bean, or a servlet; and
- for *authorization* of users to ensure they have the access control rights (permissions) required to do the actions performed.

JAAS implements a Java version of the standard Pluggable Authentication Module (PAM) framework.

Traditionally Java has provided codesource-based access controls (access controls based on *where* the code originated from and *who signed* the code). It lacked, however, the ability to additionally enforce access controls based on *who runs* the code. JAAS provides a framework that augments the Java security architecture with such support.

JAAS authentication is performed in a *pluggable* fashion. This permits applications to remain independent from underlying authentication technologies. New or updated authentication technologies can be plugged under an application without requiring modifications to the application itself. Applications enable the authentication process by instantiating a `LoginContext` object, which in turn references a `Configuration` to determine the authentication technology or technologies, or `LoginModule`(s), to be used in performing the authentication. Typical `LoginModule`s may prompt for and verify a user name and password. Others may read and verify a voice or fingerprint sample.

Once the user or service executing the code has been authenticated, the JAAS authorization component works in conjunction with the core Java SE access control model to protect access to sensitive resources. Access control decisions are based

both on the executing code's `CodeSource` *and* on the user or service running the code, who is represented by a `Subject` object. The `Subject` is updated by a `LoginModule` with relevant `Principals` and credentials if authentication succeeds.

Who Should Read This Document

This document is intended for experienced developers who require the ability to design applications constrained by a `CodeSource`-based and `Subject`-based security model. It is also intended to be read by `LoginModule` developers (developers implementing an authentication technology) prior to reading the Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide.

You may wish to first read JAAS Authentication Tutorial and JAAS Authorization Tutorial to get an overview of how to use JAAS and to see sample code in action, and then return to this document for further information.

Related Documentation

This document assumes you have already read the following:

- Java SE Platform Security Architecture
- Java SE Security Tutorial

A supplement to this guide is the JAAS LoginModule Developer's Guide, intended for experienced programmers who require the ability to write a `LoginModule` implementing an authentication technology.

The following **tutorials** for JAAS authentication and authorization can be run by everyone:

- JAAS Authentication Tutorial
- JAAS Authorization Tutorial

Similar tutorials for JAAS authentication and authorization, but which demonstrate the use of a Kerberos LoginModule and thus which require a Kerberos installation, can be found at

- JAAS Authentication
- JAAS Authorization

These two tutorials are a part of Introduction to JAAS and Java GSS-API Tutorials that utilize Kerberos as the underlying technology for authentication and secure communication.

Core Classes and Interfaces

The JAAS-related core classes and interfaces can be broken into three categories: Common, Authentication, and Authorization.

Common Classes

Common classes are those shared by both the JAAS authentication and authorization components.

The key JAAS class is `javax.security.auth.Subject`, which represents a grouping of related information for a single entity such as a person. It encompasses the entity's Principals, public credentials, and private credentials.

Note that the `java.security.Principal` interface is used to represent a Principal. Also note that a credential, as defined by JAAS, may be any Object.

Subject

To authorize access to resources, applications first need to authenticate the source of the request. The JAAS framework defines the term *subject* to represent the source of a request. A subject may be any entity, such as a person or a service. Once the subject is authenticated, a `javax.security.auth.Subject` is populated with associated identities, or Principals. A `Subject` may have many `Principals`. For example, a person may have a name `Principal` ("John Doe") and a SSN `Principal` ("123-45-6789"), which distinguish it from other subjects.

A `Subject` may also own security-related attributes, which are referred to as *credentials*; see the section Credentials. Sensitive credentials that require special protection, such as private cryptographic keys, are stored within a private credential `Set`. Credentials intended to be shared, such as public key certificates, are stored within a public credential `Set`. Different permissions (described below) are required to access and modify the different credential `Sets`.

Subjects are created using these constructors:

```
public Subject();

public Subject(boolean readOnly, Set principals,
               Set pubCredentials, Set privCredentials);
```

The first constructor creates a `Subject` with empty (non-null) `Sets` of `Principals` and credentials. The second constructor creates a `Subject` with the specified `Sets` of `Principals` and credentials. It also has a boolean argument which can be used to make the `Subject` read-only. In a read-only `Subject`, the `Principal` and credential `Sets` are immutable.

An application writer does not have to instantiate a `Subject`. If the application instantiates a `LoginContext` and does not pass a `Subject` to the `LoginContext` constructor, the `LoginContext` instantiates a new empty `Subject`. See the LoginContext section.

If a `Subject` was not instantiated to be in a read-only state, it can be set read-only by calling the following method:

```
public void setReadOnly();
```

A `javax.security.auth.AuthPermission` with target "setReadOnly" is required to invoke this method. Once in a read-only state, any attempt to add or remove `Principals` or credentials will result in an `IllegalStateException` being thrown. The following method may be called to test a `Subject`'s read-only state:

```
public boolean isReadOnly();
```

To retrieve the `Principals` associated with a Subject, two methods are available:

```
public Set getPrincipals();
public Set getPrincipals(Class c);
```

The first method returns all `Principals` contained in the `Subject`, while the second method only returns those `Principals` that are an instance of the specified Class `c`, or an instance of a subclass of Class `c`. An empty set will be returned if the `Subject` does not have any associated `Principals`.

To retrieve the public credentials associated with a `Subject`, these methods are available:

```
public Set getPublicCredentials();
public Set getPublicCredentials(Class c);
```

The behavior of these methods is similar to that for the `getPrincipals` methods, except in this case the public credentials are being obtained.

To access private credentials associated with a `Subject`, the following methods are available:

```
public Set getPrivateCredentials();
public Set getPrivateCredentials(Class c);
```

The behavior of these methods is similar to that for the `getPrincipals` and `getPublicCredentials` methods.

To modify or operate upon a `Subject`'s `PrincipalSet`, public credential `Set`, or private credential `Set`, callers use the methods defined in the `java.util.Set` class. The following example demonstrates this:

```
Subject subject;
Principal principal;
Object credential;

. . .

// add a Principal and credential to the Subject
subject.getPrincipals().add(principal);
subject.getPublicCredentials().add(credential);
```

Note: An `AuthPermission` with target "modifyPrincipals", "modifyPublicCredentials", or "modifyPrivateCredentials" is required to modify the respective `Sets`. Also note that only the sets returned via the `getPrincipals()`, `getPublicCredentials()`, and `getPrivateCredentials()` methods with no arguments are backed by the `Subject`'s respective internal sets. Therefore any modification to the returned set affects the internal sets as well. The sets returned via the `getPrincipals(Class c)`, `getPublicCredentials(Class c)`, and `getPrivateCredentials(Class c)` methods are not backed by the `Subject`'s respective internal sets. A new set is created and returned for each such method invocation. Modifications to these sets will not affect the `Subject`'s internal sets.

In order to iterate through a Set of private credentials, you need a `javax.security.auth.PrivateCredentialPermission` to access each credential. See the `PrivateCredentialPermission` API documentation for further information.

A `Subject` may be associated with an `AccessControlContext` (see the `doAs` and `doAsPrivileged` method descriptions below). The following method returns the `Subject` associated with the specified `AccessControlContext`, or `null` if no `Subject` is associated with the specified `AccessControlContext`.

```
public static Subject getSubject(final AccessControlContext acc);
```

An `AuthPermission` with target "getSubject" is required to call `Subject.getSubject`.

The `Subject` class also includes the following methods inherited from `java.lang.Object`.

```
public boolean equals(Object o);
public String toString();
public int hashCode();
```

The doAs methods for performing an action as a particular Subject

The following static methods may be called to perform an action as a particular `Subject`:

```
public static Object
    doAs(final Subject subject,
         final java.security.PrivilegedAction action);

public static Object
    doAs(final Subject subject,
         final java.security.PrivilegedExceptionAction action)
         throws java.security.PrivilegedActionException;
```

Both methods first associate the specified `subject` with the current Thread's `AccessControlContext`, and then execute the `action`. This achieves the effect of having the `action` run as the `subject`. The first method can throw runtime exceptions but normal execution has it returning an Object from the `run` method of its `action` argument. The second method behaves similarly except that it can throw a checked exception from its `PrivilegedExceptionAction run` method. An `AuthPermission` with target "doAs" is required to call the `doAs` methods.

Subject.doAs Example

Here is an example utilizing the first `doAs` method. Assume that someone named "Bob" has been authenticated by a `LoginContext` (see the LoginContext) and as a result a `Subject` was populated with a `Principal` of class `com.ibm.security.Principal`, and that `Principal` has the name "BOB". Also assume that a `SecurityManager` has been installed, and that the following exists in the access control policy (see Policy for more details on the policy file).

```
// grant "BOB" permission to read the file "foo.txt"
grant Principal com.ibm.security.Principal "BOB" {
```

```
    permission java.io.FilePermission "foo.txt", "read";
};
```

Here is the sample application code:

```
class ExampleAction implements java.security.PrivilegedAction {
    public Object run() {
        java.io.File f = new java.io.File("foo.txt");

        // the following call invokes a security check
        if (f.exists()) {
            System.out.println("File foo.txt exists");
        }
        return null;
    }
}

public class Example1 {
    public static void main(String[] args) {

        // Authenticate the subject, "BOB".
        // This process is described in the
        // LoginContext class.

        Subject bob;
        // Set bob to the Subject created during the
        // authentication process

        // perform "ExampleAction" as "BOB"
        Subject.doAs(bob, new ExampleAction());
    }
}
```

During execution, `ExampleAction` will encounter a security check when it makes a call to `f.exists()`. However, since `ExampleAction` is running as "BOB", and the policy (above) grants the necessary `FilePermission` to "BOB", the `ExampleAction` will pass the security check. If the `grant` statement in the policy is altered (adding an incorrect `CodeBase` or changing the `Principal` to "MOE", for example), then a `SecurityException` will be thrown.

The doAsPrivileged methods

The following methods also perform an action as a particular `Subject`.

```
public static Object doAsPrivileged(
    final Subject subject,
    final java.security.PrivilegedAction action,
    final java.security.AccessControlContext acc);

public static Object doAsPrivileged(
    final Subject subject,
    final java.security.PrivilegedExceptionAction action,
```

```
final java.security.AccessControlContext acc)
throws java.security.PrivilegedActionException;
```

An `AuthPermission` with target "doAsPrivileged" is required to call the `doAsPrivileged` methods.

doAs versus doAsPrivileged

The `doAsPrivileged` methods behave exactly the same as the `doAs` methods, except that instead of associating the provided `Subject` with the current Thread's `AccessControlContext`, they use the provided `AccessControlContext`. In this way, actions can be restricted by `AccessControlContexts` different from the current one.

An `AccessControlContext` contains information about all the code executed since the `AccessControlContext` was instantiated, including the code location and the permissions the code is granted by the policy. In order for an access control check to succeed, the policy must grant each code item referenced by the `AccessControlContext` the required permissions.

If the `AccessControlContext` provided to `doAsPrivileged` is `null`, then the action is not restricted by a separate `AccessControlContext`. One example where this may be useful is in a server environment. A server may authenticate multiple incoming requests and perform a separate `doAs` operation for each request. To start each `doAs` action "fresh," and without the restrictions of the current server `AccessControlContext`, the server can call `doAsPrivileged` and pass in a `nullAccessControlContext`.

Principals

As mentioned previously, once a `Subject` is authenticated, it is populated with associated identities, or `Principals`. A `Subject` may have many `Principals`. For example, a person may have a name `Principal` ("John Doe") and an SSN `Principal` ("123-45-6789"), which distinguish it from other `Subjects`. A `Principal` must implement the `java.security.Principal` and `java.io.Serializable` interfaces. See Subject for information about ways to update the `Principals` associated with a `Subject`.

Credentials

In addition to associated Principals, a Subject may own security-related attributes, which are referred to as credentials. A credential may contain information used to authenticate the subject to new services. Such credentials include passwords, Kerberos tickets, and public key certificates. Credentials might also contain data that simply enables the subject to perform certain activities. Cryptographic keys, for example, represent credentials that enable the subject to sign or encrypt data. Public and private credential classes are not part of the core JAAS class library. Any class, therefore, can represent a credential.

Public and private credential classes are not part of the core JAAS class library. Developers, however, may elect to have their credential classes implement two interfaces related to credentials: `Refreshable` and `Destroyable`.

Refreshable

The `javax.security.auth.Refreshable` **interface** provides the capability for a credential to refresh itself. For example, a credential with a particular time-restricted lifespan may implement this interface to allow callers to refresh the time period for which it is valid. The interface has two abstract methods:

```
boolean isCurrent();
```

This method determines whether the credential is current or valid.

```
void refresh() throws RefreshFailedException;
```

This method updates or extends the validity of the credential. The method implementation should perform an

`AuthPermission("refreshCredential")`

security check to ensure the caller has permission to refresh the credential.

Destroyable

The `javax.security.auth.Destroyable` **interface** provides the capability of destroying the contents within a credential. The interface has two abstract methods:

```
boolean isDestroyed();
```

Determines whether the credential has been destroyed.

```
void destroy() throws DestroyFailedException;
```

Destroys and clears the information associated with this credential. Subsequent calls to certain methods on this credential will result in an `IllegalStateException` being thrown. The method implementation should perform an `AuthPermission("destroyCredential")` security check to ensure the caller has permission to destroy the credential.

Authentication Classes and Interfaces

Authentication represents the process by which the identity of a subject is verified, and must be performed in a secure fashion; otherwise a perpetrator may impersonate others to gain access to a system. Authentication typically involves the subject demonstrating some form of evidence to prove its identity. Such evidence may be information only the subject would likely know or have (such as a password or fingerprint), or it may be information only the subject could produce (such as signed data using a private key).

To authenticate a subject (user or service), the following steps are performed:

1. An application instantiates a `LoginContext`.
2. The `LoginContext` consults a `Configuration` to load all of the `LoginModules` configured for that application.

3. The application invokes the `LoginContext`'s `login` method.
4. The `login` method invokes all of the loaded `LoginModules`. Each `LoginModule` attempts to authenticate the subject. Upon success, `LoginModules` associate relevant `Principals` and credentials with a `Subject` object that represents the subject being authenticated.
5. The `LoginContext` returns the authentication status to the application.
6. If authentication succeeded, the application retrieves the `Subject` from the `LoginContext`.

The authentication classes are described below.

LoginContext

The `javax.security.auth.login.LoginContext` class provides the basic methods used to authenticate subjects, and provides a way to develop an application independent of the underlying authentication technology. The `LoginContext` consults a `Configuration` to determine the authentication services, or `LoginModule(s)`, configured for a particular application. Therefore, different `LoginModules` can be plugged in under an application without requiring any modifications to the application itself.

`LoginContext` offers four constructors from which to choose:

```
    public LoginContext(String name) throws LoginException;

    public LoginContext(String name, Subject subject) throws
LoginException;

    public LoginContext(String name, CallbackHandler callbackHandler)
           throws LoginException

    public LoginContext(String name, Subject subject,
           CallbackHandler callbackHandler) throws LoginException
```

All of the constructors share a common parameter: name. This argument is used by the `LoginContext` as an index into the login `Configuration` to determine which `LoginModules` are configured for the application instantiating the `LoginContext`. Constructors that do not take a Subject as an input parameter instantiate a new `Subject`. Null inputs are disallowed for all constructors. Callers require an `AuthPermission` with target "createLoginContext.<*name*>" to instantiate a `LoginContext`. Here, <*name*> refers to the name of the login configuration entry that the application references in the name parameter for the `LoginContext` instantiation.

See CallbackHandler for information on what a `CallbackHandler` is and when you may need one.

Actual authentication occurs with a call to the following method:

```
    public void login() throws LoginException;
```

When `login` is invoked, all of the configured `LoginModules` are invoked to perform the authentication. If the authentication succeeded, the `Subject` (which may now hold

`Principals`, public credentials, and private credentials) can be retrieved by using the following method:

```
public Subject getSubject();
```

To logout a `Subject` and remove its authenticated `Principals` and credentials, the following method is provided:

```
public void logout() throws LoginException;
```

The following code sample demonstrates the calls necessary to authenticate and logout a Subject:

```
// let the LoginContext instantiate a new Subject
LoginContext lc = new LoginContext("entryFoo");
try {
    // authenticate the Subject
    lc.login();
    System.out.println("authentication successful");

    // get the authenticated Subject
    Subject subject = lc.getSubject();

    ...

    // all finished -- logout
    lc.logout();
} catch (LoginException le) {
    System.err.println("authentication unsuccessful: " +
        le.getMessage());
}
```

LoginModule

The `LoginModule` **interface** gives developers the ability to implement different kinds of authentication technologies that can be plugged in under an application. For example, one type of `LoginModule` may perform a user name/password-based form of authentication. Other `LoginModules` may interface to hardware devices such as smart cards or biometric devices.

Note: If you are an application writer, you do not need to understand the workings of `LoginModules`. All you have to know is how to write your application and specify configuration information (such as in a login configuration file) such that the application will be able to utilize the LoginModule specified by the configuration to authenticate the user.

If, on the other hand, you are a programmer who wishes to write a LoginModule implementing an authentication technology, see the Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide for detailed step-by-step instructions.

CallbackHandler

In some cases a `LoginModule` must communicate with the user to obtain authentication information. `LoginModules` use a `javax.security.auth.callback.CallbackHandler` for this purpose. Applications implement the `CallbackHandler` **interface** and pass it to the `LoginContext`, which forwards it directly to the underlying `LoginModules`. A `LoginModule` uses the `CallbackHandler` both to gather input from users (such as a password or smart card pin number) or to supply information to users (such as status information). By allowing the application to specify the `CallbackHandler`, underlying `LoginModules` can remain independent of the different ways applications interact with users. For example, the implementation of a `CallbackHandler` for a GUI application might display a window to solicit input from a user. On the other hand, the implementation of a `CallbackHandler` for a non-GUI tool might simply prompt the user for input directly from the command line.

`CallbackHandler`

is an interface with one method to implement:

```
void handle(Callback[] callbacks)
    throws java.io.IOException, UnsupportedCallbackException;
```

The `LoginModule` passes the `CallbackHandler handle` method an array of appropriate `Callbacks`, for example a `NameCallback` for the user name and a `PasswordCallback` for the password, and the `CallbackHandler` performs the requested user interaction and sets appropriate values in the `Callbacks`. For example, to process a `NameCallback`, the `CallbackHandler` may prompt for a name, retrieve the value from the user, and call the `NameCallback`'s `setName` method to store the name.

The `CallbackHandler` documentation has a lengthy example not included in this document that readers may want to examine.

Callback

The `javax.security.auth.callback` package contains the Callback **interface** as well as several implementations. LoginModules may pass an array of Callbacks directly to the handle method of a CallbackHandler.

Please consult the various Callback APIs for more information on their use.

Authorization Classes

To make JAAS authorization take place, granting access control permissions based not just on what code is running but also on who is running it, the following is required:

- The user must be authenticated, as described in the LoginContext section.
- The Subject that is the result of authentication must be associated with an access control context, as described in the Subject section.
- Principal-based entries must be configured in the security policy, as described below.

The `Policy` abstract class and the authorization-specific classes `AuthPermission` and `PrivateCredentialPermission` are described below.

Policy

The `java.security.Policy` class is an **abstract** class for representing the system-wide access control policy. The `Policy` API supports `Principal`-based queries.

As a default, the JDK provides a file-based subclass implementation, which was upgraded to support `Principal`-based `grant` entries in policy files.

Policy files and the structure of entries within them are described in Default Policy Implementation and Policy File Syntax.

AuthPermission

The `javax.security.auth.AuthPermission` class encapsulates the basic permissions required for JAAS. An `AuthPermission` contains a name (also referred to as a "target name") but no actions list; you either have the named permission or you don't.

In addition to its inherited methods (from the `java.security.Permission` class), an `AuthPermission` has two public constructors:

```
public AuthPermission(String name);
public AuthPermission(String name, String actions);
```

The first constructor creates a new `AuthPermission` with the specified name. The second constructor also creates a new `AuthPermission` object with the specified name, but has an additional actions argument which is currently unused and should be null. This constructor exists solely for the `Policy` object to instantiate new `Permission` objects. For most other code, the first constructor is appropriate.

Currently the `AuthPermission` object is used to guard access to the `Policy`, `Subject`, `LoginContext`, and `Configuration` objects. Refer to the `AuthPermission` Javadoc API documentation for the list of valid names that are supported.

PrivateCredentialPermission

The `javax.security.auth.PrivateCredentialPermission` class protects access to a `Subject`'s private credentials and provides one public constructor:

```
public PrivateCredentialPermission(String name, String actions);
```

JAAS Tutorials and Sample Programs

The JAAS Authentication and JAAS Authorization tutorials contain the following samples:

- `SampleAcn.java` is a sample application demonstrating JAAS authentication.
- `SampleAzn.java` is a sample application used by the authorization tutorial. It demonstrates both authentication and authorization.
- The Login Configuration File for the JAAS Authentication Tutorial describes `sample_jaas.config`, which is a sample login configuration file used by both tutorials.

- `sampleacn.policy` is a sample policy file granting permissions required by the code for the authentication tutorial.
- `sampleazn.policy` is a sample policy file granting permissions required by the code for the authorization tutorial.
- `SampleLoginModule.java` is the class specified by the tutorials' login configuration file (`sample_jaas.config`) as the class implementing the desired underlying authentication. `SampleLoginModule`'s user authentication consists of simply verifying that the name and password specified by the user have specific values.
- `SamplePrincipal.java` is a sample class implementing the `Principal` interface. It is used by `SampleLoginModule`.

See the tutorials for detailed information about the applications, the policy files, and the login configuration file.

Application writers do not need to understand the code for `SampleLoginModule.java` or `SamplePrincipal.java`, as explained in the tutorials. Programmers who wish to write LoginModules can learn how to do so by reading the Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide.

Appendix A: JAAS Settings in the java.security Security Properties File

A number of JAAS-related settings can be configured in the `java.security` master Security Properties file, which is located in the `conf/security` directory of the JDK.

JAAS adds two new security properties to `java.security`:

- `login.configuration.provider`
- `login.config.url.`*`n`*

The following pre-existing properties are also relevant for JAAS users:

- `policy.provider`
- `policy.url.`*`n`*

The following example demonstrates how to configure these properties. In this example, we leave the values provided in the default `java.security` file for the `policy.provider`, `policy.url.`*`n`*, and `login.configuration.provider` Security Properties. The default `java.security` file also lists a value for the `login.config.url.`*`n`* Security Property, but it is commented out. In the example below, it is not commented.

```
...

#
# Class to instantiate as the javax.security.auth.login.Configuration
# provider.
#
login.configuration.provider=sun.security.provider.ConfigFile

#
# Default login configuration file
```

```
#
#login.config.url.1=file:${user.home}/.java.login.config

#
# Class to instantiate as the system Policy. This is the name of the class
# that will be used as the Policy object. The system class loader is used
to
# locate this class.
#
policy.provider=sun.security.provider.PolicyFile

# The default is to have a single system-wide policy file,
# and a policy file in the user's home directory.
#
policy.url.1=file:${java.home}/conf/security/java.policy
policy.url.2=file:${user.home}/.java.policy

...
```

Note: Modifications made to this file may be overwritten by subsequent JDK updates. However, an alternate `java.security` properties file may be specified from the command line via the system property `java.security.properties=<URL>`. This properties file appends to the system properties file. If both properties files specify values for the same key, the value from command-line properties file is selected, as it is the last one loaded.

Also, specifying `java.security.properties==<URL>` (using two equals signs), then that properties file will completely override the system properties file.

To disable the ability to specify an additional properties file from the command line, set the key `security.overridePropertiesFile` to `false` in the system properties file. It is set to `true` by default.

Login Configuration Provider

The default JAAS login configuration implementation provided by Oracle gets its configuration information from files and expects the information to be provided in a specific format shown in the tutorials.

The default JAAS login configuration implementation can be replaced by specifying the alternative provider class implementation in the `login.configuration.provider` property.

For example:

```
login.configuration.provider=com.foo.Config
```

If the Security property `login.configuration.provider` is not found, or is left unspecified, then it is set to the default value:

```
login.configuration.provider=com.sun.security.auth.login.ConfigFile
```

Note that there is no means to dynamically set the login configuration provider from the command line.

Login Configuration URLs

If you are using a login configuration implementation that expects the configuration information to be specified in files (as does the default implementation from Oracle), the location of the login configuration file(s) can be statically set by specifying their respective URLs in the `login.config.url.`*n* property. '*n*' is a consecutively numbered integer starting with 1. If multiple configuration files are specified (if *n* >= 2), they will be read and unioned into one single configuration.

For example:

```
login.config.url.1=file:C:/config/.java.login.config
login.config.url.2=file:C:/users/foo/.foo.login.config
```

If the location of the configuration files is not set in the `java.security` properties file, and also is not specified dynamically from the command line (via the `-Djava.security.auth.login.config` option), JAAS attempts to load a default configuration from

```
file:${user.home}/.java.login.config
```

Policy Provider

The default policy implementation can be replaced by specifying the alternative provider class implementation in the `policy.provider` property.

For example:

```
policy.provider=com.foo.Policy
```

If the Security property `policy.provider` is not found, or is left unspecified, then the `Policy` is set to the default value:

```
policy.provider=sun.security.provider.PolicyFile
```

Note that there is no means to dynamically set the policy provider from the command line.

Policy File URLs

The location of the access control policy files can be statically set by specifying their respective URLs in the `auth.policy.url.`*n* property. *n* is a consecutively numbered integer starting with 1. If multiple policies are specified (if *n* >= 2), they will be read and unioned into one single policy.

For example:

```
policy.url.1=file:C:/policy/.java.policy
policy.url.2=file:C:/users/foo/.foo.policy
```

If the location of the policy file(s) is not set in the `java.security` properties file, and is not specified dynamically from the command line (via the `-Djava.security.policy` option), the access control policy defaults to the same policy as that of the system policy file installed with the JDK. That policy file

- grants all permissions to standard extensions
- allows anyone to listen on un-privileged ports
- allows any code to read certain "standard" properties that are not security-sensitive, such as the `os.name` and `file.separator` properties.

Appendix B: JAAS Login Configuration File

JAAS authentication is performed in a pluggable fashion, so Java applications can remain independent from underlying authentication technologies. Configuration information such as the desired authentication technology is specified at runtime. The source of the configuration information (for example, a file or a database) is up to the current `javax.security.auth.login.Configuration` implementation. The default `Configuration` implementation, `ConfigFile`, gets its configuration information from login configuration files. For details about the default login `Configuration` implementation provided with JAAS, see the `com.sun.security.auth.login.ConfigFile` class.

Login Configuration File Structure and Contents

A login configuration file consists of one or more entries, each specifying which underlying authentication technology should be used for a particular application or applications. The structure of each entry is the following:

```
<name used by application to refer to this entry> {
    <LoginModule> <flag> <LoginModule options>;
    <optional additional LoginModules, flags and options>;
};
```

Thus, each login configuration file entry consists of a name followed by one or more `LoginModule`-specific entries, where each `LoginModule`-specific entry is terminated by a semicolon and the entire group of `LoginModule`-specific entries is enclosed in braces. Each configuration file entry is terminated by a semicolon.

Example 6-1 Login Configuration File for JAAS Authentication Tutorial

As an example, the login configuration file used for the JAAS Authentication Tutorial tutorial contains just one entry, which is

```
Sample {
   sample.module.SampleLoginModule required debug=true;
};
```

Here, the entry is named `Sample` and that is the name that the JAAS Authentication tutorial application (`SampleAcn.java`) uses to refer to this entry. The entry specifies that the `LoginModule` to be used to do the user authentication is the `SampleLoginModule` in the `sample.module` package and that this `SampleLoginModule` is required to "succeed" in order for authentication to be considered successful. The

`SampleLoginModule` succeeds only if the name and password supplied by the user are the ones it expects (`testUser` and `testPassword`, respectively).

The **name** for an entry in a login configuration file is the name that applications use to refer to the entry when they instantiate a `LoginContext`, as described in JAAS Authentication Tutorial in the JAAS authentication tutorial. The name can be whatever name the application developer wishes to use. Here, the term "application" refers to whatever code does the JAAS login.

The specified `LoginModules` (described below) are used to control the authentication process. Authentication proceeds down the list in the exact order specified, as described in the `Configuration` class.

The subparts of each `LoginModule`-specific entry are the following:

- **LoginModule**: This specifies a class implementing the desired authentication technology. Specifically, the class must be a subclass of the `LoginModule` class, which is in the `javax.security.auth.spi` package. A typical `LoginModule` may prompt for and verify a user name and password, as is done by the `SampleLoginModule` (in the `sample.module` package) used for these tutorials. Any vendor can provide a LoginModule implementation that you can use. Some implementations are supplied with the JRE from Oracle. You can view the reference documentation for the various `LoginModules`, all in the `com.sun.security.auth` package:
  - `JndiLoginModule`
  - `KeyStoreLoginModule`
  - `Krb5LoginModule`
  - `NTLoginModule`
  - `UnixLoginModule`
- **flag**: The flag value indicates whether success of the preceding `LoginModule` is `required`, `requisite`, `sufficient`, or `optional`. If there is just one `LoginModule`-specific entry, as there is in our tutorials, then the flag for it should be "required". The options are described in more detail in the `Configuration` class.
- **LoginModule options**: If the specified `LoginModule` implementation has options that can be set, you specify any desired option values here. This is a space-separated list of values which are passed directly to the underlying `LoginModule`. Options are defined by the `LoginModule` itself, and control the behavior within it. For example, a `LoginModule` may define options to support debugging/testing capabilities.

  The correct way to specify options in the configuration file is by using a name-value pairing, for example `debug=true`, where the option name (in this case, `debug`) and value (in this case, `true`) should be separated by an equals symbol.

Example 6-2 Login Configuration File Demonstrating required, sufficient, requisite, and optional Flags

The following is a sample login configuration file that demonstrates the `required`, `sufficient`, `requisite`, and `optional` flags. See the `Configuration` class for more information about these flags.

```
Login1 {
       sample.SampleLoginModule required debug=true;
```

```
};

Login2 {
   sample.SampleLoginModule required;
   com.sun.security.auth.module.NTLoginModule sufficient;
   com.foo.SmartCard requisite debug=true;
   com.foo.Kerberos optional debug=true;
};
```

The application **Login1** only has one configured `LoginModule`, `SampleLoginModule`. Therefore, an attempt by **Login1** to authenticate a subject (user or service) will be successful if and only if the `SampleLoginModule` succeeds.

The authentication logic for the application **Login2** is easier to explain with the following table:

Table 6-1 Login2 Authentication Status

Module Class	Flag	Authentication Attempt 1	Authentication Attempt 2	Authentication Attempt 3	Authentication Attempt 4	Authentication Attempt 5	Authentication Attempt 6	Authentication Attempt 7	Authentication Attempt 8
`SampleLoginModule`	required	pass	pass	pass	pass	fail	fail	fail	fail
`NTLoginModule`	sufficient	pass	fail	fail	fail	pass	fail	fail	fail
SmartCard	requisite	*	pass	pass	fail	*	pass	pass	fail
Kerberos	optional	*	pass	fail	*	*	pass	fail	*
Overall Authentication	not applicable	pass	pass	pass	fail	fail	fail	fail	fail

* = trivial value due to control returning to the application because a previous *requisite* module failed or a previous *sufficient* module succeeded.

Where to Specify Which Login Configuration File Should Be Used

The configuration file to be used can be specified in one of two ways:

1. On the command line.

   You can use a `-Djava.security.auth.login.config` interpreter command line argument to specify the login configuration file that should be used. We use this approach for all the tutorials. For example, we run our `SampleAcn` application in the JAAS Authentication Tutorial using the following command, which specifies that the configuration file is the `sample_jaas.config` file in the current directory:

   ```
   java -Djava.security.auth.login.config==sample_jaas.config
   sample.SampleAcn
   ```

> **Note:**
>
> If you use a single equals sign (=) with the `java.security.auth.login.config` system property, then the configurations specified by both this system property and the `java.security` file are used.

2. In the Java Security Properties file.

   An alternate approach to specifying the location of the login configuration file is to indicate its URL as the value of a `login.config.url.`*n* property in the security properties file. The Security Properties file is the `java.security` file located in the `conf/security` directory of the JDK.

   Here, *n* indicates a consecutively-numbered integer starting with 1. Thus, if desired, you can specify more than one login configuration file by indicating one file's URL for the `login.config.url.1` property, a second file's URL for the `login.config.url.2` property, and so on. If more than one login configuration file is specified (that is, if *n* > 1), then the files are read and concatenated into a single configuration.

   Here is an example of what would need to be added to the security properties file in order to indicate the `sample_jaas.config` login configuration file used by this tutorial. This example assumes the file is in the `C:\AcnTest` directory on Windows:

   ```
   login.config.url.1=file:C:/AcnTest/sample_jaas.config
   ```

   (Note that URLs always use forward slashes, regardless of what operating system the user is running.)

JAAS Tutorials

This page links to two tutorials demonstrating various aspects of the use of JAAS (Java Authentication and Authorization Service):

- JAAS Authentication Tutorial: Explains how an application can authenticate users using JAAS.
- JAAS Authorization Tutorial: Explains how to enforce user-based access controls using JAAS.

The authentication technology used for these tutorials is very basic, just ensuring that the user specifies a particular name and password. Thus, these tutorials can be run by everyone.

JAAS Authentication Tutorial

JAAS can be used for two purposes:

- for *authentication* of users, to reliably and securely determine who is currently executing Java code, regardless of whether the code is running as an application, an applet, a bean, or a servlet; and
- for *authorization* of users to ensure they have the access control rights (permissions) required to do the actions performed.

This section provides a basic tutorial for the authentication component. The authorization component will be described in the JAAS Authorization Tutorial.

JAAS authentication is performed in a *pluggable* fashion. This permits Java applications to remain independent from underlying authentication technologies. New or updated technologies can be plugged in without requiring modifications to the application itself. An implementation for a particular authentication technology to be used is determined at runtime. The implementation is specified in a login configuration file. The authentication technology used for this tutorial is very basic, just ensuring that the user specifies a particular name and password.

The rest of this tutorial consists of the following sections:

1. The Authentication Tutorial Code
2. The Login Configuration
3. Running the Code
4. Running the Code with a Security Manager

If you want to first see the tutorial code in action, you can skip directly to Running the Code and then go back to the other sections to learn about coding and configuration file details.

The Authentication Tutorial Code

The code for this tutorial consists of three files:

- `SampleAcn.java` contains the sample application class (`SampleAcn`) and another class used to handle user input (`MyCallbackHandler`). **The code in this file is the only code you need to understand for this tutorial. Your application will only indirectly use the other source files.**
- `SampleLoginModule.java` is the class specified by the tutorial's login configuration file, `sample_jass.config`, described in The Login Configuration File for the JAAS Authentication Tutorial as the class implementing the desired underlying authentication. `SampleLoginModule`'s user authentication consists of simply verifying that the name and password specified by the user have specific values.
- `SamplePrincipal.java` is a sample class implementing the `java.security.Principal` interface. It is used by `SampleLoginModule`.

SampleAcn.java

Our authentication tutorial application code is contained in a single source file, `SampleAcn.java`. That file contains two classes:

- The SampleAcn Class
- The MyCallbackHandler Class

The SampleAcn Class

The `main` method of the `SampleAcn` class performs the authentication and then reports whether or not authentication succeeded.

The code for authenticating the user is very simple, consisting of just two steps:

1. Instantiating a LoginContext
2. Calling the LoginContext's login Method

First the basic code is shown, followed by The Complete SampleAcn Class Code, complete with the import statement it requires and error handling.

Instantiating a LoginContext

In order to authenticate a user, you first need a `javax.security.auth.login.LoginContext`. Here is the basic way to instantiate a LoginContext:

```
import javax.security.auth.login.*;
. . .
LoginContext lc =
    new LoginContext(<config file entry name>,
           <CallbackHandler to be used for user interaction>);
```

and here is the specific way our tutorial code does the instantiation:

```
import javax.security.auth.login.*;
. . .
LoginContext lc =
    new LoginContext("Sample",
          new MyCallbackHandler());
```

The arguments are the following:

1. **The name of an entry in the JAAS login configuration file**
   This is the name for the `LoginContext` to use to look up an entry for this application in the JAAS login configuration file, described in The Login Configuration. Such an entry specifies the class(es) that implement the desired underlying authentication technology(ies). The class(es) must implement the `LoginModule` interface, which is in the `javax.security.auth.spi` package.

   In our sample code, we use the `SampleLoginModule` supplied with this tutorial. The `SampleLoginModule` performs authentication by ensuring that the user types a particular name and password.

   The entry in the login configuration file we use for this tutorial, `sample_jass.config` (see The Login Configuration File for the JAAS Authentication Tutorial), has the name "Sample", so that is the name we specify as the first argument to the LoginContext constructor.

2. **A CallbackHandler instance**
   When a `LoginModule` needs to communicate with the user, for example to ask for a user name and password, it does not do so directly. That is because there are various ways of communicating with a user, and it is desirable for `LoginModules` to remain independent of the different types of user interaction. Rather, the `LoginModule` invokes a `javax.security.auth.callback.CallbackHandler` to perform the user interaction and obtain the requested information, such as the user name and password.

An instance of the particular `CallbackHandler` to be used is specified as the second argument to the `LoginContext` constructor. The `LoginContext` forwards that instance to the underlying `LoginModule` (in our case `SampleLoginModule`). An application typically provides its own `CallbackHandler` implementation. A simple `CallbackHandler`, `TextCallbackHandler`, is provided in the `com.sun.security.auth.callback` package to output information to and read input from the command line. However, we instead demonstrate the more typical case of an application providing its own `CallbackHandler` implementation, described in The MyCallbackHandler Class.

Calling the LoginContext's login Method

Once we have a `LoginContext lc`, we can call its `login` method to carry out the authentication process:

```
lc.login();
```

The `LoginContext` instantiates a new empty `javax.security.auth.Subject` object (which represents the user or service being authenticated; see Subject). The `LoginContext` constructs the configured `LoginModule` (in our case `SampleLoginModule`) and initializes it with this new `Subject` and `MyCallbackHandler`.

The `LoginContext`'s `login` method then calls methods in the `SampleLoginModule` to perform the login and authentication. The `SampleLoginModule` will utilize the `MyCallbackHandler` to obtain the user name and password. Then the `SampleLoginModule` will check that the name and password are the ones it expects.

If authentication is successful, the `SampleLoginModule` populates the `Subject` with a `Principal` representing the user. The `Principal` the `SampleLoginModule` places in the `Subject` is an instance of `SamplePrincipal`, which is a sample class implementing the `java.security.Principal` interface.

The calling application can subsequently retrieve the authenticated `Subject` by calling the `LoginContext`'s `getSubject` method, although doing so is not necessary for this tutorial.

The Complete SampleAcn Class Code

Now that you have seen the basic code required to authenticate the user, we can put it all together into the full class in `SampleAcn.java`, which includes relevant import statements and error handling:

SampleAcn.java

```
package sample;

import java.io.*;
import java.util.*;
import javax.security.auth.login.*;
import javax.security.auth.*;
import javax.security.auth.callback.*;

/**
 * This Sample application attempts to authenticate a user
 * and reports whether or not the authentication was successful.
 */
public class SampleAcn {
```

```
/**
 * Attempt to authenticate the user.
 *
 * @param args input arguments for this application.  These are ignored.
 */
 public static void main(String[] args) {

     // Obtain a LoginContext, needed for authentication. Tell it
     // to use the LoginModule implementation specified by the
     // entry named "Sample" in the JAAS login configuration
     // file and to also use the specified CallbackHandler.
     LoginContext lc = null;
     try {
         lc = new LoginContext("Sample", new MyCallbackHandler());
     } catch (LoginException le) {
         System.err.println("Cannot create LoginContext. "
             + le.getMessage());
         System.exit(-1);
     } catch (SecurityException se) {
         System.err.println("Cannot create LoginContext. "
             + se.getMessage());
         System.exit(-1);
     }

     // the user has 3 attempts to authenticate successfully
     int i;
     for (i = 0; i < 3; i++) {
         try {

             // attempt authentication
             lc.login();

             // if we return with no exception, authentication succeeded
             break;

         } catch (LoginException le) {

               System.err.println("Authentication failed:");
               System.err.println("  " + le.getMessage());
               try {
                   Thread.currentThread().sleep(3000);
               } catch (Exception e) {
                   // ignore
               }

         }
     }

     // did they fail three times?
     if (i == 3) {
         System.out.println("Sorry");
         System.exit(-1);
     }
```

```
        System.out.println("Authentication succeeded!");

    }
}


/**
 * The application implements the CallbackHandler.
 *
 * <p> This application is text-based.  Therefore it displays information
 * to the user using the OutputStreams System.out and System.err,
 * and gathers input from the user using the InputStream System.in.
 */
class MyCallbackHandler implements CallbackHandler {

    /**
     * Invoke an array of Callbacks.
     *
     * <p>
     *
     * @param callbacks an array of <code>Callback</code> objects which
contain
     *                  the information requested by an underlying security
     *                  service to be retrieved or displayed.
     *
     * @exception java.io.IOException if an input or output error occurs.
<p>
     *
     * @exception UnsupportedCallbackException if the implementation of
this
     *                  method does not support one or more of the
Callbacks
     *                  specified in the <code>callbacks</code> parameter.
     */
    public void handle(Callback[] callbacks)
    throws IOException, UnsupportedCallbackException {

        for (int i = 0; i < callbacks.length; i++) {
            if (callbacks[i] instanceof TextOutputCallback) {

                // display the message according to the specified type
                TextOutputCallback toc = (TextOutputCallback)callbacks[i];
                switch (toc.getMessageType()) {
                case TextOutputCallback.INFORMATION:
                    System.out.println(toc.getMessage());
                    break;
                case TextOutputCallback.ERROR:
                    System.out.println("ERROR: " + toc.getMessage());
                    break;
                case TextOutputCallback.WARNING:
                    System.out.println("WARNING: " + toc.getMessage());
                    break;
                default:
                    throw new IOException("Unsupported message type: " +
                                          toc.getMessageType());
```

```
                }

            } else if (callbacks[i] instanceof NameCallback) {

                // prompt the user for a username
                NameCallback nc = (NameCallback)callbacks[i];

                System.err.print(nc.getPrompt());
                System.err.flush();
                nc.setName((new BufferedReader
                        (new InputStreamReader(System.in))).readLine());

            } else if (callbacks[i] instanceof PasswordCallback) {

                // prompt the user for sensitive information
                PasswordCallback pc = (PasswordCallback)callbacks[i];
                System.err.print(pc.getPrompt());
                System.err.flush();
                pc.setPassword(System.console().readPassword());

            } else {
                throw new UnsupportedCallbackException
                        (callbacks[i], "Unrecognized Callback");
            }
        }
    }
}
```

The MyCallbackHandler Class

In some cases a `LoginModule` must communicate with the user to obtain authentication information. `LoginModules` use a `javax.security.auth.callback.CallbackHandler` for this purpose. An application can either use one of the sample implementations provided in the `com.sun.security.auth.callback` package or, more typically, write a `CallbackHandler` implementation. The application passes the `CallbackHandler` as an argument to the `LoginContext` instantiation. The `LoginContext` forwards the `CallbackHandler` directly to the underlying `LoginModules`.

The tutorial sample code supplies its own `CallbackHandler` implementation, the `MyCallbackHandler` class in page 6-20.

`CallbackHandler` is an interface with one method to implement:

```
void handle(Callback[] callbacks)
    throws java.io.IOException, UnsupportedCallbackException;
```

The `LoginModule` passes the `CallbackHandler` handle method an array of appropriate `javax.security.auth.callback.Callbacks`, for example a `NameCallback` for the user name and a `PasswordCallback` for the password, and the `CallbackHandler` performs the requested user interaction and sets appropriate values in the `Callbacks`.

The `MyCallbackHandler handle` method is structured as follows:

```
public void handle(Callback[] callbacks)
  throws IOException, UnsupportedCallbackException {

  for (int i = 0; i < callbacks.length; i++) {
    if (callbacks[i] instanceof TextOutputCallback) {

      // display a message according to a specified type
      . . .

    } else if (callbacks[i] instanceof NameCallback) {

      // prompt the user for a username
      . . .

    } else if (callbacks[i] instanceof PasswordCallback) {

      // prompt the user for a password
      . . .

    } else {
        throw new UnsupportedCallbackException
         (callbacks[i], "Unrecognized Callback");
    }
  }
}
```

A `CallbackHandler handle` method is passed an array of `Callback` instances, each of a particular type (`NameCallback`, `PasswordCallback`, etc.). It must handle each Callback, performing user interaction in a way that is appropriate for the executing application.

`MyCallbackHandler` handles three types of Callbacks: **`NameCallback`** to prompt the user for a user name, **`PasswordCallback`** to prompt for a password, and **`TextOutputCallback`** to report any error, warning, or other messages the `SampleLoginModule` wishes to send to the user.

The `handle` method handles a **`TextOutputCallback`** by extracting the message to be reported and then printing it to `System.out`, optionally preceded by additional wording that depends on the message type. The message to be reported is determined by calling the `TextOutputCallback`'s `getMessage` method and the type by calling its `getMessageType` method. Here is the code for handling a `TextOutputCallback`:

```
if (callbacks[i] instanceof TextOutputCallback) {

  // display the message according to the specified type
  TextOutputCallback toc = (TextOutputCallback)callbacks[i];
  switch (toc.getMessageType()) {
     case TextOutputCallback.INFORMATION:
        System.out.println(toc.getMessage());
        break;
     case TextOutputCallback.ERROR:
```

```
            System.out.println("ERROR: " + toc.getMessage());
            break;
        case TextOutputCallback.WARNING:
            System.out.println("WARNING: " + toc.getMessage());
            break;
        default:
            throw new IOException("Unsupported message type: " +
                toc.getMessageType());
    }
```

The `handle` method handles a **`NameCallback`** by prompting the user for a user name. It does this by printing the prompt to `System.err`. It then sets the name for use by the `SampleLoginModule` by calling the `NameCallback`'s `setName` method, passing it the name typed by the user:

```
} else if (callbacks[i] instanceof NameCallback) {

    // prompt the user for a username
    NameCallback nc = (NameCallback)callbacks[i];

    System.err.print(nc.getPrompt());
    System.err.flush();
    nc.setName((new BufferedReader
        (new InputStreamReader(System.in))).readLine());
```

Similarly, the `handle` method handles a **`PasswordCallback`** by printing a prompt to `System.err` to prompt the user for a password. It then sets the password for use by the `SampleLoginModule` by calling the `PasswordCallback`'s `setPassword` method, passing it the password typed by the user:

```
} else if (callbacks[i] instanceof PasswordCallback) {

    // prompt the user for sensitive information
    PasswordCallback pc = (PasswordCallback)callbacks[i];

    System.err.print(pc.getPrompt());
    System.err.flush();
    pc.setPassword(System.console().readPassword());
```

SampleLoginModule.java and SamplePrincipal.java

`SampleLoginModule.java` implements the `LoginModule` interface. `SampleLoginModule` is the class specified by the tutorial's login configuration file (see The Login Configuration File for the JAAS Authentication Tutorial) as the class implementing the desired underlying authentication. `SampleLoginModule`'s user authentication consists of simply verifying that the name and password specified by the user have specific values. This `SampleLoginModule` is specified by the tutorial's login configuration file as the `LoginModule` to use because (1) It performs a basic type of authentication suitable for any environment and thus is appropriate for a tutorial for all users, and (2) It provides an example `LoginModule` implementation for experienced programmers who require the ability to write a `LoginModule` implementing an authentication technology.

`SamplePrincipal.java` is a sample class implementing the `java.security.Principal` interface. If authentication is successful, the `SampleLoginModule` populates a `Subject` with a `SamplePrincipal` representing the user.

Important: If you are an application writer, you do not need to know how to write a `LoginModule` or a `Principal` implementation. You do not need to examine the `SampleLoginModule` or `SamplePrincipal` code. All you have to know is how to write your application and specify configuration information (such as in a login configuration file) such that the application will be able to utilize the `LoginModule` specified by the configuration to authenticate the user. You need to determine which `LoginModule`(s) you want to use and read the `LoginModule`'s documentation to learn about what options you can specify values for (in the configuration) to control the `LoginModule`'s behavior.

Any vendor can provide a LoginModule implementation that you can use. Some implementations are supplied with the JRE from Oracle, as listed in Appendix B: JAAS Login Configuration File.

Information for programmers who want to write a `LoginModule` can be found in Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide.

SampleLoginModule.java

```
package sample.module;

import java.util.*;
import java.io.IOException;
import javax.security.auth.*;
import javax.security.auth.callback.*;
import javax.security.auth.login.*;
import javax.security.auth.spi.*;
import sample.principal.SamplePrincipal;

/**
 * <p> This sample LoginModule authenticates users with a password.
 *
 * <p> This LoginModule only recognizes one user:       testUser
 * <p> testUser's password is:  testPassword
 *
 * <p> If testUser successfully authenticates itself,
 * a <code>SamplePrincipal</code> with the testUser's user name
 * is added to the Subject.
 *
 * <p> This LoginModule recognizes the debug option.
 * If set to true in the login Configuration,
 * debug messages will be output to the output stream, System.out.
 *
 */
public class SampleLoginModule implements LoginModule {

    // initial state
    private Subject subject;
    private CallbackHandler callbackHandler;
    private Map sharedState;
    private Map options;
```

```
    // configurable option
    private boolean debug = false;

    // the authentication status
    private boolean succeeded = false;
    private boolean commitSucceeded = false;

    // username and password
    private String username;
    private char[] password;

    // testUser's SamplePrincipal
    private SamplePrincipal userPrincipal;

    /**
     * Initialize this <code>LoginModule</code>.
     *
     * @param subject the <code>Subject</code> to be authenticated. <p>
     *
     * @param callbackHandler a <code>CallbackHandler</code> for
communicating
     *                  with the end user (prompting for user names and
     *                  passwords, for example). <p>
     *
     * @param sharedState shared <code>LoginModule</code> state. <p>
     *
     * @param options options specified in the login
     *                  <code>Configuration</code> for this particular
     *                  <code>LoginModule</code>.
     */
    public void initialize(Subject subject,
                   CallbackHandler callbackHandler,
                         Map<java.lang.String, ?> sharedState,
                         Map<java.lang.String, ?> options) {

        this.subject = subject;
        this.callbackHandler = callbackHandler;
        this.sharedState = sharedState;
        this.options = options;

        // initialize any configured options
        debug = "true".equalsIgnoreCase((String)options.get("debug"));
    }

    /**
     * Authenticate the user by prompting for a user name and password.
     *
     * @return true in all cases since this <code>LoginModule</code>
     *          should not be ignored.
     *
     * @exception FailedLoginException if the authentication fails. <p>
     *
     * @exception LoginException if this <code>LoginModule</code>
     *          is unable to perform the authentication.
```

```
     */
    public boolean login() throws LoginException {

        // prompt for a user name and password
        if (callbackHandler == null)
            throw new LoginException("Error: no CallbackHandler available
" +
                        "to garner authentication information from the
user");

        Callback[] callbacks = new Callback[2];
        callbacks[0] = new NameCallback("user name: ");
        callbacks[1] = new PasswordCallback("password: ", false);

        try {
            callbackHandler.handle(callbacks);
            username = ((NameCallback)callbacks[0]).getName();
            char[] tmpPassword =
((PasswordCallback)callbacks[1]).getPassword();
            if (tmpPassword == null) {
                // treat a NULL password as an empty password
                tmpPassword = new char[0];
            }
            password = new char[tmpPassword.length];
            System.arraycopy(tmpPassword, 0,
                        password, 0, tmpPassword.length);
            ((PasswordCallback)callbacks[1]).clearPassword();

        } catch (java.io.IOException ioe) {
            throw new LoginException(ioe.toString());
        } catch (UnsupportedCallbackException uce) {
            throw new LoginException("Error: " +
uce.getCallback().toString() +
                " not available to garner authentication information " +
                "from the user");
        }

        // print debugging information
        if (debug) {
            System.out.println("\t\t[SampleLoginModule] " +
                                "user entered user name: " +
                                username);
            System.out.print("\t\t[SampleLoginModule] " +
                                "user entered password: ");
            for (int i = 0; i < password.length; i++)
                System.out.print(password[i]);
            System.out.println();
        }

        // verify the username/password
        boolean usernameCorrect = false;
        boolean passwordCorrect = false;
        if (username.equals("testUser"))
            usernameCorrect = true;
        if (usernameCorrect &&
```

```
                    password.length == 12 &&
                    password[0] == 't' &&
                    password[1] == 'e' &&
                    password[2] == 's' &&
                    password[3] == 't' &&
                    password[4] == 'P' &&
                    password[5] == 'a' &&
                    password[6] == 's' &&
                    password[7] == 's' &&
                    password[8] == 'w' &&
                    password[9] == 'o' &&
                    password[10] == 'r' &&
                    password[11] == 'd') {

                    // authentication succeeded!!!
                    passwordCorrect = true;
                    if (debug)
                        System.out.println("\t\t[SampleLoginModule] " +
                                           "authentication succeeded");
                    succeeded = true;
                    return true;
                } else {

                    // authentication failed -- clean out state
                    if (debug)
                        System.out.println("\t\t[SampleLoginModule] " +
                                           "authentication failed");
                    succeeded = false;
                    username = null;
                    for (int i = 0; i < password.length; i++)
                        password[i] = ' ';
                    password = null;
                    if (!usernameCorrect) {
                        throw new FailedLoginException("User Name Incorrect");
                    } else {
                        throw new FailedLoginException("Password Incorrect");
                    }
                }
            }

            /**
             * This method is called if the LoginContext's
             * overall authentication succeeded
             * (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL
        LoginModules
             * succeeded).
             *
             * If this LoginModule's own authentication attempt
             * succeeded (checked by retrieving the private state saved by the
             * <code>login</code> method), then this method associates a
             * <code>SamplePrincipal</code>
             * with the <code>Subject</code> located in the
             * <code>LoginModule</code>.  If this LoginModule's own
             * authentication attempted failed, then this method removes
             * any state that was originally saved.
```

```
     *
     * @exception LoginException if the commit fails.
     *
     * @return true if this LoginModule's own login and commit
     *          attempts succeeded, or false otherwise.
     */
    public boolean commit() throws LoginException {
        if (succeeded == false) {
            return false;
        } else {
            // add a Principal (authenticated identity)
            // to the Subject

            // assume the user we authenticated is the SamplePrincipal
            userPrincipal = new SamplePrincipal(username);
            if (!subject.getPrincipals().contains(userPrincipal))
                subject.getPrincipals().add(userPrincipal);

            if (debug) {
                System.out.println("\t\t[SampleLoginModule] " +
                                "added SamplePrincipal to Subject");
            }

            // in any case, clean out state
            username = null;
            for (int i = 0; i < password.length; i++)
                password[i] = ' ';
            password = null;

            commitSucceeded = true;
            return true;
        }
    }

    /**
     * This method is called if the LoginContext's
     * overall authentication failed.
     * (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL
LoginModules
     * did not succeed).
     *
     * If this LoginModule's own authentication attempt
     * succeeded (checked by retrieving the private state saved by the
     * <code>login</code> and <code>commit</code> methods),
     * then this method cleans up any state that was originally saved.
     *
     * @exception LoginException if the abort fails.
     *
     * @return false if this LoginModule's own login and/or commit attempts
     *          failed, and true otherwise.
     */
    public boolean abort() throws LoginException {
        if (succeeded == false) {
            return false;
        } else if (succeeded == true && commitSucceeded == false) {
```

```
            // login succeeded but overall authentication failed
            succeeded = false;
            username = null;
            if (password != null) {
                for (int i = 0; i < password.length; i++)
                    password[i] = ' ';
                password = null;
            }
            userPrincipal = null;
        } else {
            // overall authentication succeeded and commit succeeded,
            // but someone else's commit failed
            logout();
        }
        return true;
    }

    /**
     * Logout the user.
     *
     * This method removes the <code>SamplePrincipal</code>
     * that was added by the <code>commit</code> method.
     *
     * @exception LoginException if the logout fails.
     *
     * @return true in all cases since this <code>LoginModule</code>
     *          should not be ignored.
     */
    public boolean logout() throws LoginException {

        subject.getPrincipals().remove(userPrincipal);
        succeeded = false;
        succeeded = commitSucceeded;
        username = null;
        if (password != null) {
            for (int i = 0; i < password.length; i++)
                password[i] = ' ';
            password = null;
        }
        userPrincipal = null;
        return true;
    }
}
```

SamplePrincipal.java

```
package sample.principal;

import java.security.Principal;

/**
 * This class implements the <code>Principal</code> interface
 * and represents a Sample user.
 *
```

```
 * Principals such as this <code>SamplePrincipal</code>
 * may be associated with a particular <code>Subject</code>
 * to augment that <code>Subject</code> with an additional
 * identity.  Refer to the <code>Subject</code> class for more information
 * on how to achieve this.  Authorization decisions can then be based upon
 * the Principals associated with a <code>Subject</code>.
 *
 * @see java.security.Principal
 * @see javax.security.auth.Subject
 */
public class SamplePrincipal implements Principal, java.io.Serializable {

    /**
     * @serial
     */
    private String name;

    /**
     * Create a SamplePrincipal with a Sample username.
     *
     * @param name the Sample username for this user.
     *
     * @exception NullPointerException if the <code>name</code>
     *                  is <code>null</code>.
     */
    public SamplePrincipal(String name) {
        if (name == null)
            throw new NullPointerException("illegal null input");

        this.name = name;
    }

    /**
     * Return the Sample username for this <code>SamplePrincipal</code>.
     *
     * @return the Sample username for this <code>SamplePrincipal</code>
     */
    public String getName() {
        return name;
    }

    /**
     * Return a string representation of this <code>SamplePrincipal</code>.
     *
     * @return a string representation of this <code>SamplePrincipal</
code>.
     */
    public String toString() {
        return("SamplePrincipal:  " + name);
    }

    /**
     * Compares the specified Object with this <code>SamplePrincipal</code>
     * for equality.  Returns true if the given object is also a
     * <code>SamplePrincipal</code> and the two SamplePrincipals
```

```
     * have the same username.
     *
     * @param o Object to be compared for equality with this
     *          <code>SamplePrincipal</code>.
     *
     * @return true if the specified Object is equal equal to this
     *          <code>SamplePrincipal</code>.
     */
    public boolean equals(Object o) {
        if (o == null)
            return false;

        if (this == o)
            return true;

        if (!(o instanceof SamplePrincipal))
            return false;
        SamplePrincipal that = (SamplePrincipal)o;

        if (this.getName().equals(that.getName()))
            return true;
        return false;
    }

    /**
     * Return a hash code for this <code>SamplePrincipal</code>.
     *
     * @return a hash code for this <code>SamplePrincipal</code>.
     */
    public int hashCode() {
        return name.hashCode();
    }
}
```

The Login Configuration

JAAS authentication is performed in a pluggable fashion, so applications can remain independent from underlying authentication technologies. A system administrator determines the authentication technologies, or `LoginModules`, to be used for each application and configures them in a login `Configuration`. The source of the configuration information (for example, a file or a database) is up to the current `javax.security.auth.login.Configuration` implementation. The default `Configuration` implementation from Oracle reads configuration information from configuration files, as described in the `ConfigFile` class.

See Appendix B: JAAS Login Configuration File for information as to what a login configuration file is, what it contains, and how to specify which login configuration file should be used.

The Login Configuration File for the JAAS Authentication Tutorial

As noted, the login configuration file we use for this tutorial, `sample_jass.config`, contains just one entry, which is

```
Sample {
  sample.module.SampleLoginModule required debug=true;
};
```

This entry is named "Sample" and that is the name that our tutorial application, `SampleAcn`, uses to refer to this entry. The entry specifies that the `LoginModule` to be used to do the user authentication is the `SampleLoginModule` in the `sample.module` package and that this `SampleLoginModule` is required to "succeed" in order for authentication to be considered successful. The `SampleLoginModule` succeeds only if the name and password supplied by the user are the one it expects ("testUser" and "testPassword", respectively).

The `SampleLoginModule` also defines a "debug" option that can be set to `true` as shown. If this option is set to `true`, `SampleLoginModule` outputs extra information about the progress of authentication. A `LoginModule` can define as many options as it wants. The `LoginModule` documentation should specify the possible option names and values you can set in your configuration file.

Running the Code

To execute our JAAS authentication tutorial code, all you have to do is

1. Place the following file into a directory:
   - `sample_jass.config` login configuration file (see The Login Configuration File for the JAAS Authentication Tutorial)
2. Create a subdirectory named `sample` of that top-level directory, and place the following into it (note the `SampleAcn` and `MyCallbackHandler` classes, both in `SampleAcn.java`, are in a package named `sample`):
   - `SampleAcn.java` application source file
3. Create a subdirectory of the `sample` directory and name it `module`. Place the following into it (note the `SampleLoginModule` class is in a package named `sample.module`):
   - `SampleLoginModule.java` source file
4. Create another subdirectory of the `sample` directory and name it `principal`. Place the following into it (note the `SamplePrincipal` class is in a package named `sample.principal`):
   - `SamplePrincipal.java` source file
5. While in the top-level directory, compile `SampleAcn.java`, `SampleLoginModule.java`, and `SamplePrincipal.java`:

   ```
   javac sample/SampleAcn.java sample/module/SampleLoginModule.java
   sample/principal/SamplePrincipal.java
   ```

   (Type all that on one line.)

6. Execute the `SampleAcn` application, specifying
   - by `-Djava.security.auth.login.config==sample_jaas.config` that the login configuration file to be used is `sample_jaas.config`.

The full command is below.

```
java -Djava.security.auth.login.config==sample_jaas.config sample.SampleAcn
```

You will be prompted for your user name and password, and the `SampleLoginModule` specified in the login configuration file will check to ensure these are correct. The `SampleLoginModule` expects `testUser` for the user name and `testPassword` for the password.

You will see some messages output by `SampleLoginModule` as a result of the `debug` option being set to `true` in the login configuration file. Then, if your login is successful, you will see the following message output by `SampleAcn`:

```
Authentication succeeded!
```

If the login is not successful (for example, if you misspell the password), you will see

```
Authentication failed:
```

followed by a reason for the failure. For example, if you mistype the password, you may see a message like the following:

```
Authentication failed:
  Password Incorrect
```

`SampleAcn` gives you three chances to successfully log in.

Running the Code with a Security Manager

When a Java program is run with a security manager installed, the program is not allowed to access resources or otherwise perform security-sensitive operations unless it is explicitly granted permission to do so by the security policy in effect. (See Permissions in the JDK.) The permission must be granted by an entry in a policy file (see Default Policy Implementation and Policy File Syntax.)

Most browsers install a security manager, so *applets* typically run under the scrutiny of a security manager. *Applications*, on the other hand, do not, since a security manager is not automatically installed when an application is running. Thus an application, like our `SampleAcn` application, by default has full access to resources.

To run an application with a security manager, simply invoke the interpreter with a `-Djava.security.manager` argument included on the command line.

If you try invoking `SampleAcn` with a security manager but without specifying any policy file, you will get the following (unless you have a default policy setup elsewhere that grants the required permissions or grants `AllPermission`):

```
% java -Djava.security.manager \
 -Djava.security.auth.login.config==sample_jaas.config sample.SampleAcn
```

```
Exception in thread "main" java.security.AccessControlException:
  access denied (
  javax.security.auth.AuthPermission createLoginContext.Sample)
```

As you can see, you get an `AccessControlException`, because we haven't created and used a policy file granting our code the permission that is required in order to be allowed to create a `LoginContext`.

Here are the complete steps required in order to be able to run our `SampleAcn` application with a security manager installed. You can skip the first five steps if you have already done them, as described in Running the Code.

1. Place the following file into a directory:
   - `sample_jass.config` login configuration file (see The Login Configuration File for the JAAS Authentication Tutorial)
2. Create a subdirectory named `sample` of that top-level directory, and place the following into it (note the `SampleAcn` and `MyCallbackHandler` classes, both in `SampleAcn.java`, are in a package named `sample`):
   - `SampleAcn.java` application source file
3. Create a subdirectory of the `sample` directory and name it `module`. Place the following into it (note the `SampleLoginModule` class is in a package named `sample.module`):
   - `SampleLoginModule.java` source file
4. Create another subdirectory of the `sample` directory and name it `principal`. Place the following into it (note the `SamplePrincipal` class is in a package named `sample.principal`):
   - `SamplePrincipal.java` source file
5. While in the top-level directory, compile `SampleAcn.java`, `SampleLoginModule.java`, and `SamplePrincipal.java`:

   ```
   javac sample/SampleAcn.java sample/module/SampleLoginModule.java
   sample/principal/SamplePrincipal.java
   ```

   (Type all that on one line.)
6. Create a JAR file containing `SampleAcn.class` and `MyCallbackHandler.class`:

   ```
   jar -cvf SampleAcn.jar sample/SampleAcn.class sample/
   MyCallbackHandler.class
   ```

   (Type all that on one line.) This command creates a JAR file, `SampleAcn.jar`, and places the `SampleAcn.class` and `MyCallbackHandler.class` files inside it.
7. Create a JAR file containing `SampleLoginModule.class` and `SamplePrincipal.class`:

   ```
   jar -cvf SampleLM.jar sample/module/SampleLoginModule.class sample/
   principal/SamplePrincipal.class
   ```
8. Create a policy file granting the required permissions.

   The permission that is needed by code attempting to instantiate a `LoginContext` is a `javax.security.auth.AuthPermission` with target `createLoginContext.<entry name>`. Here, `<entry name>` refers to the name of the login configuration file entry that the application references in its instantiation of

`LoginContext`. The name used by our `SampleAcn` application's `LoginContext` instantiation is `Sample`, as you can see in the code:

```
LoginContext lc =
    new LoginContext("Sample",
          new MyCallbackHandler());
```

Thus, the permission that needs to be granted to `SampleAcn.jar` is

```
permission javax.security.auth.AuthPermission
  "createLoginContext.Sample";
```

The `SampleLM.jar` file also needs to be granted a permission. The documentation for a `LoginModule` should tell you what permissions it needs to be granted. In the case of `SampleLoginModule`, it needs a `javax.security.auth.AuthPermission` with target `modifyPrincipals` in order to populate a `Subject` with a `Principal`:

```
permission javax.security.auth.AuthPermission
  "modifyPrincipals";
```

Copy the policy file `sampleacn.policy` to the same directory as that in which you stored `SampleAcn.java`, etc. The policy file contains the following `grant` statement to grant `SampleAcn.jar` (in the current directory) its required permission:

```
grant codebase "file:./SampleAcn.jar" {
   permission javax.security.auth.AuthPermission
"createLoginContext.Sample";
};
```

The policy file also contains the following `grant` statement to grant `SampleLM.jar` (also in the current directory) its required permission:

```
grant codebase "file:./SampleLM.jar" {
   permission javax.security.auth.AuthPermission "modifyPrincipals";
};
```

Note: Policy files and the structure of entries within them are described in Default Policy Implementation and Policy File Syntax. Permissions are described in Permissions in the JDK.

Execute the `SampleAcn` application, specifying

a. by an appropriate `-classpath` clause that classes should be searched for in the `SampleAcn.jar` and `SampleLM.jar` JAR files,

b. by `-Djava.security.manager` that a security manager should be installed,

c. by `-Djava.security.policy==sampleacn.policy` that the policy file to be used is `sampleacn.policy`, and

d. by `-Djava.security.auth.login.config==sample_jaas.config` that the login configuration file to be used is `sample_jaas.config`.

Below are the full commands to use for both Windows and Solaris, Linux, and macOS systems. The only difference is that on Windows systems you use semicolons to separate class path items, while you use colons for that purpose on Solaris, Linux, and macOS systems.

Here is the full command for Windows systems:

```
java -classpath SampleAcn.jar;SampleLM.jar
 -Djava.security.manager
 -Djava.security.policy==sampleacn.policy
 -Djava.security.auth.login.config==sample_jaas.config
 sample.SampleAcn
```

Here is the full command for Solaris, Linux, and macOS systems:

```
java -classpath SampleAcn.jar:SampleLM.jar
 -Djava.security.manager
 -Djava.security.policy==sampleacn.policy
 -Djava.security.auth.login.config==sample_jaas.config
 sample.SampleAcn
```

Type all that on one line. Multiple lines are used here for legibility. If the command is too long for your system, you may need to place it in a `.bat` file (for Microsoft Windows) or a `.sh` file (for Solaris, Linux, and macOS) and then run that file to execute the command.

Since the specified policy file contains an entry granting the code the required permissions, execution should proceed without any exceptions indicating a required permission was not granted. You will be prompted for a user name and password (use `testUser` and `testPassword`), and the `SampleLoginModule` specified in the login configuration file will check the name and password. If your login is successful, you will see the message `Authentication succeeded!` and if not, you will see `Authentication failed:` followed by a reason for the failure.

sampleacn.policy

```
/* grant the sample LoginModule permissions */
grant codebase "file:./SampleLM.jar" {
    permission javax.security.auth.AuthPermission "modifyPrincipals";
};

grant codebase "file:./SampleAcn.jar" {
   permission javax.security.auth.AuthPermission
"createLoginContext.Sample";
};
```

JAAS Authorization Tutorial

This tutorial expands the program and policy file developed in the JAAS Authentication Tutorial tutorial to demonstrate the JAAS authorization component, which ensures the authenticated caller has the access control rights (permissions) required to do subsequent security-sensitive operations. Since the authorization component requires

that the user authentication first be completed, please read the JAAS Authentication Tutorial tutorial first if you have not already done so.

The rest of this tutorial consists of the following sections:

- What is JAAS Authorization?
- How is JAAS Authorization Performed?
  - How Do You Make Principal-Based Policy File Statements?
  - How Do You Associate a Subject with an Access Control Context?
- The Authorization Tutorial Code
- The Login Configuration File for the JAAS Authorization Tutorial
- The Policy File
- Running the Authorization Tutorial Code

If you want to first see the tutorial code in action, you can skip directly to Running the Authorization Tutorial Code and then go back to the other sections to learn more.

What is JAAS Authorization?

JAAS authorization extends the existing Java security architecture that uses a security policy (see Default Policy Implementation and Policy File Syntax) to specify what access rights are granted to executing code. That architecture is *code-centric*. That is, the permissions are granted based on code characteristics: where the code is coming from and whether it is digitally signed and if so by whom. We saw an example of this in the `sampleacn.policy` file used in the JAAS Authentication Tutorial tutorial. That file contains the following:

```
grant codebase "file:./SampleAcn.jar" {

   permission javax.security.auth.AuthPermission
                    "createLoginContext.Sample";
};
```

This grants the code in the `SampleAcn.jar` file, located in the current directory, the specified permission. (No signer is specified, so it doesn't matter whether the code is signed or not.)

JAAS authorization augments the existing code-centric access controls with new *user-centric* access controls. Permissions can be granted based not just on what code is running but also on *who* is running it.

When an application uses JAAS authentication to authenticate the user (or other entity such as a service), a `Subject` is created as a result. The purpose of the `Subject` is to represent the authenticated user. A `Subject` is comprised of a set of `Principal`s, where each `Principal` represents an identity for that user. For example, a `Subject` could have a name `Principal` ("Susan Smith") and a Social Security Number `Principal` ("987-65-4321"), thereby distinguishing this `Subject` from other `Subject`s.

Permissions can be granted in the policy to specific `Principal`s. After the user has been authenticated, the application can associate the `Subject` with the current access control context. For each subsequent security-checked operation (a local file

access, for example), the Java runtime will automatically determine whether the policy grants the required permission only to a specific `Principal` and if so, the operation will be allowed only if the `Subject` associated with the access control context contains the designated `Principal`.

How is JAAS Authorization Performed?

To make JAAS authorization take place, the following is required:

- The user must be authenticated, as described in the JAAS Authentication Tutorial tutorial.
- `Principal`-based entries must be configured in the security policy; see How Do You Make Principal-Based Policy File Statements?
- The `Subject` that is the result of authentication must be associated with the current access control context; see How Do You Associate a Subject with an Access Control Context?.

How Do You Make Principal-Based Policy File Statements?

Policy file `grant` statements (see Default Policy Implementation and Policy File Syntax can now optionally include one or more `Principal` fields. Inclusion of a `Principal` field indicates that the user or other entity represented by the specified `Principal`, executing the specified code, has the designated permissions.

Thus, the basic format of a `grant` statement is now

```
grant <signer(s) field>, <codeBase URL>
  <Principal field(s)> {
    permission perm_class_name "target_name", "action";
    ....
    permission perm_class_name "target_name", "action";
  };
```

where each of the `signer`, `codeBase` and `Principal` fields is optional and the order between the fields doesn't matter.

A `Principal` field looks like the following:

```
Principal Principal_class "principal_name"
```

That is, it is the word `Principal` (where case doesn't matter) followed by the (fully qualified) name of a `Principal` class and a principal name.

A `Principal` class is a class that implements the `java.security.Principal` interface. All `Principal` objects have an associated name that can be obtained by calling their `getName` method. The format used for the name is dependent on each `Principal` implementation.

The type of `Principal` placed in the `Subject` created by the basic authentication mechanism used by this tutorial is `SamplePrincipal`, so that is what should be used as the *`Principal_class`* part of our `grant` statement's `Principal` designation. User names for `SamplePrincipals` are of the form *name*, and the only user name accepted

for this tutorial is `testUser`, so the *`principal_name`* designation to use in the `grant` statement is `testUser`.

It is possible to include more than one `Principal` field in a `grant` statement. If multiple `Principal` fields are specified, then the permissions in that `grant` statement are granted only if the `Subject` associated with the current access control context contains *all* of those `Principals`.

To grant the same set of permissions to different `Principals`, create multiple `grant` statements where each lists the permissions and contains a single `Principal` field designating one of the `Principals`.

The policy file for this tutorial includes one `grant` statement with a `Principal` field:

```
grant codebase "file:./SampleAction.jar",
        Principal sample.principal.SamplePrincipal "testUser" {

   permission java.util.PropertyPermission "java.home", "read";
   permission java.util.PropertyPermission "user.home", "read";
   permission java.io.FilePermission "foo.txt", "read";
};
```

This specifies that the indicated permissions are granted to the specified `Principal` executing the code in `SampleAction.jar`. (Note: the `SamplePrincipal` class is in the `sample.principal` package.)

How Do You Associate a Subject with an Access Control Context?

To create and associate a Subject with the current access control context, you need the following:

- The user must first be authenticated, as described in JAAS Authentication Tutorial.
- The static `doAs` method from the `Subject` class must be called, passing it an authenticated `Subject` and a `java.security.PrivilegedAction` or `java.security.PrivilegedExceptionAction`. (See Appendix A: API for Privileged Blocks in Permissions in the JDK for a comparison of PrivilegedAction and PrivilegedExceptionAction.) The `doAs` method associates the provided Subject with the current access control context and then invokes the `run` method from the action. The `run` method implementation contains all the code to be executed as the specified Subject. The action thus executes as the specified Subject.
  The static `doAsPrivileged` method from the Subject class may be called instead of the `doAs` method, as will be done for this tutorial. In addition to the parameters passed to `doAs`, `doAsPrivileged` requires a third parameter: an AccessControlContext. Unlike `doAs`, which associates the provided Subject with the current access control context, `doAsPrivileged` associates the Subject with the provided access control context or with an empty access control context if the parameter passed in is `null`, as is the case for this tutorial. See doAs vs. doAsPrivileged in the JAAS Reference Guide for a comparison of those methods.

The Authorization Tutorial Code

The code for this tutorial consists of four files:

- `SampleAzn.java` is exactly the same as the `SampleAcn.java` application file from the JAAS Authentication Tutorial tutorial except for the additional code needed to call `Subject.doAsPrivileged`.
- `SampleAction.java` contains the `SampleAction` class. This class implements `PrivilegedAction` and has a `run` method that contains all the code we want to be executed with `Principal`-based authorization checks.
- `SampleLoginModule.java` is the class specified by the tutorial's login configuration file (see The Login Configuration File for the JAAS Authorization Tutorial) as the class implementing the desired underlying authentication. `SampleLoginModule`'s user authentication consists of simply verifying that the name and password specified by the user have specific values. This class was also used by the JAAS Authentication Tutorial tutorial and will not be discussed further here.
- `SamplePrincipal.java` is a sample class implementing the `java.security.Principal` interface. It is used by `SampleLoginModule`. This class was also used by the JAAS Authentication tutorial and will not be discussed further here.

The `SampleLoginModule.java` and `SamplePrincipal.java` files were also used in the JAAS Authentication Tutorial tutorial, so they are not described further here. The other source files are described below.

SampleAzn.java

Like `SampleAcn`, the `SampleAzn` class instantiates a `LoginContext lc` and calls its `login` method to perform the authentication. If successful, the authenticated `Subject` (which includes a `SamplePrincipal` representing the user) is obtained by calling the `LoginContext`'s `getSubject` method:

```
Subject mySubject = lc.getSubject();
```

After providing the user some information about the `Subject`, such as which `Principals` it has, the `main` method then calls `Subject.doAsPrivileged`, passing it the authenticated `Subject mySubject`, a `PrivilegedAction` (`SampleAction`) and a `null AccessControlContext`, as described in the following.

The `SampleAction` class is instantiated via the following:

```
PrivilegedAction action = new SampleAction();
```

The call to `Subject.doAsPrivileged` is performed via:

```
Subject.doAsPrivileged(mySubject, action, null);
```

The `doAsPrivileged` method invokes execution of the `run` method in the `PrivilegedAction action` (`SampleAction`) to initiate execution of the rest of the code, which is considered to be executed on behalf of the `Subject mySubject`.

Passing `null` as the `AccessControlContext` (third) argument to `doAsPrivileged` indicates that `mySubject` should be associated with a new empty `AccessControlContext`. The result is that security checks occurring during execution of `SampleAction` will only require permissions for the `SampleAction` code

itself (or other code it invokes), running as `mySubject`. Note that the caller of `doAsPrivileged` (and the callers on the execution stack at the time `doAsPrivileged` was called) do not require any permissions while the action executes.

SampleAzn.java

```
package sample;

import java.io.*;
import java.util.*;
import java.security.Principal;
import java.security.PrivilegedAction;
import javax.security.auth.*;
import javax.security.auth.callback.*;
import javax.security.auth.login.*;
import javax.security.auth.spi.*;
import com.sun.security.auth.*;

/**
 * This Sample application attempts to authenticate a user
 * and executes a SampleAction as that user.
 *
 * If the user successfully authenticates itself,
 * the username and number of Credentials is displayed.
 */
public class SampleAzn {

    /**
     * Attempt to authenticate the user.
     *
     * @param args input arguments for this application.  These are
ignored.
     */
     public static void main(String[] args) {

        // Obtain a LoginContext, needed for authentication. Tell it
        // to use the LoginModule implementation specified by the
        // entry named "Sample" in the JAAS login configuration
        // file and to also use the specified CallbackHandler.
        LoginContext lc = null;
        try {
            lc = new LoginContext("Sample", new MyCallbackHandler());
        } catch (LoginException le) {
            System.err.println("Cannot create LoginContext. "
                + le.getMessage());
            System.exit(-1);
        } catch (SecurityException se) {
            System.err.println("Cannot create LoginContext. "
                + se.getMessage());
            System.exit(-1);
        }

        // the user has 3 attempts to authenticate successfully
        int i;
        for (i = 0; i < 3; i++) {
```

```
            try {

                // attempt authentication
                lc.login();

                // if we return with no exception, authentication succeeded
                break;

            } catch (LoginException le) {

                  System.err.println("Authentication failed:");
                  System.err.println("  " + le.getMessage());
                  try {
                      Thread.currentThread().sleep(3000);
                  } catch (Exception e) {
                      // ignore
                  }

            }
        }

        // did they fail three times?
        if (i == 3) {
            System.out.println("Sorry");
            System.exit(-1);
        }

        System.out.println("Authentication succeeded!");

        Subject mySubject = lc.getSubject();

        // let's see what Principals we have
        Iterator principalIterator = mySubject.getPrincipals().iterator();
        System.out.println("Authenticated user has the following
Principals:");
        while (principalIterator.hasNext()) {
            Principal p = (Principal)principalIterator.next();
            System.out.println("\t" + p.toString());
        }

        System.out.println("User has " +
                        mySubject.getPublicCredentials().size() +
                        " Public Credential(s)");

        // now try to execute the SampleAction as the authenticated Subject
        PrivilegedAction action = new SampleAction();
        Subject.doAsPrivileged(mySubject, action, null);

        System.exit(0);
    }
}

/**
 * A CallbackHandler implemented by the application.
 *
```

```
 * This application is text-based.  Therefore it displays information
 * to the user using the OutputStreams System.out and System.err,
 * and gathers input from the user using the InputStream System.in.
 */
class MyCallbackHandler implements CallbackHandler {

    /**
     * Invoke an array of Callbacks.
     *
     * @param callbacks an array of <code>Callback</code> objects which
contain
     *                  the information requested by an underlying security
     *                  service to be retrieved or displayed.
     *
     * @exception java.io.IOException if an input or output error occurs.
<p>
     *
     * @exception UnsupportedCallbackException if the implementation of
this
     *                  method does not support one or more of the
Callbacks
     *                  specified in the <code>callbacks</code> parameter.
     */
    public void handle(Callback[] callbacks)
    throws IOException, UnsupportedCallbackException {

        for (int i = 0; i < callbacks.length; i++) {
            if (callbacks[i] instanceof TextOutputCallback) {

                // display the message according to the specified type
                TextOutputCallback toc = (TextOutputCallback)callbacks[i];
                switch (toc.getMessageType()) {
                case TextOutputCallback.INFORMATION:
                    System.out.println(toc.getMessage());
                    break;
                case TextOutputCallback.ERROR:
                    System.out.println("ERROR: " + toc.getMessage());
                    break;
                case TextOutputCallback.WARNING:
                    System.out.println("WARNING: " + toc.getMessage());
                    break;
                default:
                    throw new IOException("Unsupported message type: " +
                                          toc.getMessageType());
                }

            } else if (callbacks[i] instanceof NameCallback) {

                // prompt the user for a username
                NameCallback nc = (NameCallback)callbacks[i];

                System.err.print(nc.getPrompt());
                System.err.flush();
                nc.setName((new BufferedReader
                        (new InputStreamReader(System.in))).readLine());
```

```
            } else if (callbacks[i] instanceof PasswordCallback) {

                // prompt the user for sensitive information
                PasswordCallback pc = (PasswordCallback)callbacks[i];
                System.err.print(pc.getPrompt());
                System.err.flush();
                pc.setPassword(System.console().readPassword());

            } else {
                throw new UnsupportedCallbackException
                        (callbacks[i], "Unrecognized Callback");
            }
        }
    }
}
```

SampleAction.java

`SampleAction.java` contains the `SampleAction` class. This class implements `java.security.PrivilegedAction` and has a `run` method that contains all the code we want to be executed as the Subject `mySubject`. For this tutorial, we will perform three operations, each of which cannot be done unless code has been granted required permissions. We will:

- Read and print the value of the `java.home` system property,
- Read and print the value of the `user.home` system property, and
- Determine whether or not a file named `foo.txt` exists in the current directory.

Here is the code:

SampleAction.java

```
package sample;

import java.io.File;
import java.security.PrivilegedAction;

/**
 * This is a Sample PrivilegedAction implementation, designed to be
 * used with the Sample application.
 *
 */
public class SampleAction implements PrivilegedAction {

    /**
     * This Sample PrivilegedAction performs the following operations:
     * <ul>
     *   <li>Access the System property, <i>java.home</i></li>
     *   <li> Access the System property, <i>user.home</i></li>
     *   <li> Access the file, <i>foo.txt</i></li>
     * </ul>
     *
     * @return <code>null</code> in all cases.
```

```
     *
     * @exception SecurityException if the caller does not have permission
     *           to perform the operations listed above.
     */
    public Object run() {
        System.out.println("\nYour java.home property: "
                           +System.getProperty("java.home"));

        System.out.println("\nYour user.home property: "
                           +System.getProperty("user.home"));

        File f = new File("foo.txt");
        System.out.print("\nfoo.txt does ");
        if (!f.exists())
            System.out.print("not ");
        System.out.println("exist in the current working directory.");
        return null;
    }
}
```

The Login Configuration File for the JAAS Authorization Tutorial

The login configuration file used for this tutorial can be exactly the same as that used by the JAAS Authentication Tutorial tutorial. Thus we can use the `sample_jaas.config` file, which contains just one entry:

```
Sample {
  sample.module.SampleLoginModule required debug=true;
};
```

This entry is named `Sample` and that is the name that both our tutorial applications `SampleAcn` and `SampleAzn` use to refer to it. The entry specifies that the `LoginModule` to be used to do the user authentication is the `SampleLoginModule` in the `sample.module` package and that this `SampleLoginModule` is required to "succeed" in order for authentication to be considered successful. The `SampleLoginModule` succeeds only if the name and password supplied by the user are the one it expects (`testUser` and `testPassword`, respectively).

The `SampleLoginModule` also defines a `debug` option that can be set to `true` as shown. If this option is set to `true`, `SampleLoginModule` outputs extra information about the progress of authentication.

The Policy File

The application for this authorization tutorial consists of two classes, `SampleAzn` and `SampleAction`. The code in each class contains some security-sensitive operations and thus relevant permissions are required in a policy file in order for the operations to be executed.

The `LoginModule` used by this tutorial, `SampleLoginModule`, also contains an operation requiring a permission.

The permissions required by each of these classes are described below, followed by the full policy file.

Permissions Required by SampleAzn

The main method of the `SampleAzn` class does two operations for which permissions are required. It

- creates a `LoginContext`, and
- calls the `doAsPrivileged` static method of the Subject class.

The `LoginContext` creation is exactly the same as was done in the authentication tutorial, and it thus needs the same `javax.security.auth.AuthPermission` permission with target "`createLoginContext.Sample`".

In order to call the `doAsPrivileged` method of the `Subject` class, you need to have a `javax.security.auth.AuthPermission` with target "`doAsPrivileged`".

Assuming the `SampleAzn` class is placed in a JAR file named `SampleAzn.jar`, these permissions can be granted to the `SampleAzn` code via the following `grant` statement in the policy file:

```
grant codebase "file:./SampleAzn.jar" {
   permission javax.security.auth.AuthPermission
                    "createLoginContext.Sample";
   permission javax.security.auth.AuthPermission "doAsPrivileged";
};
```

Permissions Required by SampleAction

The `SampleAction` code does three operations for which permissions are required. It

- reads the value of the `java.home` system property.
- reads the value of the `user.home` system property.
- checks to see whether or not a file named `foo.txt` exists in the current directory.

The permissions required for these operations are the following:

```
permission java.util.PropertyPermission "java.home", "read";
permission java.util.PropertyPermission "user.home", "read";
permission java.io.FilePermission "foo.txt", "read";
```

We need to grant these permissions to the code in `SampleAction.class`, which we will place in a JAR file named `SampleAction.jar`. However, for this particular `grant` statement we want to grant the permissions not just to the *code* but to a specific user executing the code, to demonstrate how to restrict access to a particular user.

Thus, as explained in How Do You Make Principal-Based Policy File Statements?, our `grant` statement looks like the following:

```
grant codebase "file:./SampleAction.jar", Principal
sample.principal.SamplePrincipal "testUser" {
    permission java.util.PropertyPermission "java.home", "read";
    permission java.util.PropertyPermission "user.home", "read";
    permission java.io.FilePermission "foo.txt", "read";
};
```

Permissions Required by SampleLoginModule

The `SampleLoginModule` code does one operation for which permissions are required. It needs a `javax.security.auth.AuthPermission` with target "`modifyPrincipals`" in order to populate a `Subject` with a `Principal`. The grant statement is the following:

```
grant codebase "file:./SampleLM.jar" {
    permission javax.security.auth.AuthPermission "modifyPrincipals";
};
```

The Full Policy File

The full policy file is `sampleazn.policy`:

sampleazn.policy

```
/* grant the sample LoginModule permissions */

grant codebase "file:./SampleAction.jar", Principal
sample.principal.SamplePrincipal "testUser" {
    permission java.util.PropertyPermission "java.home", "read";
    permission java.util.PropertyPermission "user.home", "read";
    permission java.io.FilePermission "foo.txt", "read";
};

grant codebase "file:./SampleLM.jar" {
    permission javax.security.auth.AuthPermission "modifyPrincipals";
};

grant codebase "file:./SampleAcn.jar" {
   permission javax.security.auth.AuthPermission
"createLoginContext.Sample";
};
```

Running the Authorization Tutorial Code

To execute our JAAS authorization tutorial code, all you have to do is

1. Place the following files into a directory:
   - `sample_jaas.config` login configuration file (see The Login Configuration File for the JAAS Authorization Tutorial)
   - `sampleazn.policy` policy file
2. Create a subdirectory named `sample` of that top-level directory, and place the following into it (note the `SampleAzn` and `SampleAction` classes are in a package named `sample`):
   - `SampleAzn.java` source file
   - `SampleAction.java` source file

3. Create a subdirectory of the `sample` directory and name it `module`. Place the following into it (note the `SampleLoginModule` class is in a package named `sample.module`):
   - `SampleLoginModule.java` source file
4. Create another subdirectory of the `sample` directory and name it `principal`. Place the following into it (note the `SamplePrincipal` class is in a package named `sample.principal`):
   - `SamplePrincipal.java` source file
5. While in the top-level directory, compile all the source files:

   ```
   javac sample/SampleAction.java sample/SampleAzn.java sample/module/
   SampleLoginModule.java sample/principal/SamplePrincipal.java
   ```

   (Type all that on one line.)
6. Create a JAR file named `SampleAzn.jar` containing `SampleAzn.class` and `MyCallbackHandler.class` (Note the sources for both these classes are in `SampleAzn.java`):

   ```
   jar -cvf SampleAzn.jar sample/SampleAzn.class sample/
   MyCallbackHandler.class
   ```

   (Type all that on one line.)
7. Create a JAR file named `SampleAction.jar` containing `SampleAction.class`:

   ```
   jar -cvf SampleAction.jar sample/SampleAction.class
   ```
8. Create a JAR file containing `SampleLoginModule.class` and `SamplePrincipal.class`:

   ```
   jar -cvf SampleLM.jar sample/module/SampleLoginModule.class sample/
   principal/SamplePrincipal.class
   ```
9. Execute the `SampleAzn` application, specifying
   a. by an appropriate `-classpath` clause that classes should be searched for in the `SampleAzn.jar`, `SampleAction.jar`, and `SampleLM.jar` JAR files,
   b. by `-Djava.security.manager` that a security manager should be installed,
   c. by `-Djava.security.policy==sampleazn.policy` that the policy file to be used is `sampleazn.policy`, and
   d. by `-Djava.security.auth.login.config==sample_jaas.config` that the login configuration file to be used is `sample_jaas.config`.

   Below are the full commands to use for both Windows and Solaris, Linux, and macOS systems. The only difference is that on Windows systems you use semicolons to separate class path items, while you use colons for that purpose on Solaris, Linux, and macOS systems.

   Here is the full command for Windows systems:

   ```
   java -classpath SampleAzn.jar;SampleAction.jar;SampleLM.jar
    -Djava.security.manager
    -Djava.security.policy==sampleazn.policy
    -Djava.security.auth.login.config==sample_jaas.config sample.SampleAzn
   ```

Here is the full command for Solaris, Linux, and macOS systems:

```
java -classpath SampleAzn.jar:SampleAction.jar:SampleLM.jar
 -Djava.security.manager
 -Djava.security.policy==sampleazn.policy
 -Djava.security.auth.login.config==sample_jaas.config sample.SampleAzn
```

Type the full command on one line. Multiple lines are used here for legibility. If the command is too long for your system, you may need to place it in a `.bat` file (for Windows) or a `.sh` file (for Solaris, Linux, and macOS) and then run that file to execute the command.

You will be prompted for a user name and password (use `testUser` and `testPassword`), and the `SampleLoginModule` specified in the login configuration file will check the name and password. If your login is successful, you will see the message `Authentication succeeded!` and if not, you will see `Authentication failed:` followed by a reason for the failure.

Once authentication is successfully completed, the rest of the program (in `SampleAction`) will be executed on behalf of you, the user, requiring you to have been granted appropriate permissions. The `sampleazn.policy` policy file grants you the required permissions, so you will see a display of the values of your `java.home` and `user.home` system properties and a statement as to whether or not you have a file named `foo.txt` in the current directory.

Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide

JAAS provides subject-based authorization on authenticated identities. This document focuses on the authentication aspect of JAAS, specifically the `LoginModule` interface.

Who Should Read This Document

This document is intended for experienced programmers who require the ability to write a `LoginModule` implementing an authentication technology.

Related Documentation

This document assumes you have already read the following:

- Java Authentication and Authorization Service (JAAS) Reference Guide

It also discusses various classes and interfaces in the JAAS API. See the Javadoc API documentation for the JAAS API specification for more detailed information:

- `javax.security.auth`
- `javax.security.auth.callback`
- `javax.security.auth.kerberos`
- `javax.security.auth.login`
- `javax.security.auth.spi`
- `javax.security.auth.x500`

The following packages contain supported `LoginModule` examples:

- `com.sun.security.auth`
- `com.sun.security.auth.callback`
- `com.sun.security.auth.login`
- `com.sun.security.auth.module`

The following **tutorials** for JAAS authentication and authorization can be run by everyone:

- JAAS Authentication Tutorial
- JAAS Authorization Tutorial

Similar tutorials for JAAS authentication and authorization, but which demonstrate the use of a Kerberos LoginModule and thus which require a Kerberos installation, can be found at

- JAAS Authentication Tutorial
- JAAS Authorization Tutorial

These two tutorials are a part of Introduction to JAAS and Java GSS-API Tutorials that utilize Kerberos as the underlying technology for authentication and secure communication.

Introduction to LoginModule

Authentication technology providers must implement the `LoginModule` interface. `LoginModules` are plugged in under applications to provide a particular type of authentication.

While applications write to the `LoginContext` Application Programming Interface (API), authentication technology providers implement the `LoginModule` interface. A `Configuration` specifies the `LoginModule`(s) to be used with a particular login application. Different `LoginModules` can be plugged in under the application without requiring any modifications to the application itself.

The `LoginContext` is responsible for reading the `Configuration` and instantiating the specified `LoginModules`. Each `LoginModule` is initialized with a `Subject`, a `CallbackHandler`, shared `LoginModule` state, and `LoginModule`-specific options.

The **`Subject`** represents the user or service currently being authenticated and is updated by a `LoginModule` with relevant `Principal`s and credentials if authentication succeeds. `LoginModules` use the **`CallbackHandler`** to communicate with users (to prompt for user names and passwords, for example), as described in the login method description. Note that the `CallbackHandler` may be null. A `LoginModule` that requires a `CallbackHandler` to authenticate the `Subject` may throw a `LoginException` if it was initialized with a `null` `CallbackHandler`. `LoginModules` optionally use the **shared state** to share information or data among themselves.

The **`LoginModule`-specific options** represent the options configured for this `LoginModule` in the login `Configuration`. The options are defined by the `LoginModule` itself and control the behavior within it. For example, a `LoginModule` may define options to support debugging/testing capabilities. Options are defined using a key-value syntax, such as *debug=true*. The `LoginModule` stores the options as a `Map` so

that the values may be retrieved using the key. Note that there is no limit to the number of options a `LoginModule` chooses to define.

The calling application sees the authentication process as a single operation invoked via a call to the `LoginContext`'s `login` method. However, the authentication process within each `LoginModule` proceeds in two distinct phases. In the first phase of authentication, the `LoginContext`'s `login` method invokes the `login` method of each `LoginModule` specified in the `Configuration`. The `login` method for a `LoginModule` performs the actual authentication (prompting for and verifying a password for example) and saves its authentication status as private state information. Once finished, the `LoginModule`'s `login` method returns `true` (if it succeeded) or `false` (if it should be ignored), or it throws a `LoginException` to specify a failure. In the failure case, the `LoginModule` must not retry the authentication or introduce delays. The responsibility of such tasks belongs to the application. If the application attempts to retry the authentication, each `LoginModule`'s `login` method will be called again.

In the second phase, if the `LoginContext`'s overall authentication succeeded (calls to the relevant *required*, *requisite*, *sufficient* and *optional* `LoginModules`' `login` methods succeeded), then the `commit` method for each `LoginModule` gets invoked. (For an explanation of the `LoginModule` flags *required*, *requisite*, *sufficient* and *optional*, please consult the `Configuration` documentation and Appendix B: JAAS Login Configuration File in the JAAS Reference Guide.) The `commit` method for a `LoginModule` checks its privately saved state to see if its own authentication succeeded. If the overall `LoginContext` authentication succeeded and the `LoginModule`'s own authentication succeeded, then the `commit` method associates the relevant `Principals` (authenticated identities) and credentials (authentication data such as cryptographic keys) with the `Subject`.

If the `LoginContext`'s overall authentication failed (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL `LoginModules`' `login` methods did not succeed), then the `abort` method for each `LoginModule` gets invoked. In this case, the `LoginModule` removes/destroys any authentication state originally saved.

Logging out a `Subject` involves only one phase. The `LoginContext` invokes the `LoginModule`'s `logout` method. The `logout` method for the `LoginModule` then performs the logout procedures, such as removing `Principals` or credentials from the `Subject`, or logging session information.

Steps to Implement a LoginModule

The following are the steps required to implement and test a `LoginModule`:

Step 1: Understand the Authentication Technology

First, understand the authentication technology to be implemented by your new `LoginModule` provider and determine its requirements.

1. Determine whether or not your `LoginModule` will require some form of user interaction (retrieving a user name and password, for example). If so, you will need to become familiar with the `CallbackHandler` interface and the `javax.security.auth.callback` package.

   In that package you will find several possible `Callback` implementations to use. (Alternatively, you can create your own `Callback` implementations.) The `LoginModule` will invoke the `CallbackHandler` specified by the application itself

and passed to the `LoginModule`'s `initialize` method. The `LoginModule` passes the `CallbackHandler` an array of appropriate `Callbacks`. See LoginModule.login Method in Step 3: Implement the LoginModule Interface.

> **Note:**
>
> It is possible for `LoginModule` implementations not to have any end-user interactions. Such `LoginModules` would not need to access the `callback` package.

2. Determine what configuration options you want to make available to the user, who specifies configuration information in whatever form the current Configuration implementation expects (for example, in files). For each option, decide the option name and possible values.

   For example, if a `LoginModule` may be configured to consult a particular authentication server host, decide on the option's key name ("auth_server", for example), as well as the possible server hostnames valid for that option ("server_one.example.com" and "server_two.example.com", for example).

Step 2: Name the LoginModule Implementation

Decide on the proper package and class name for your `LoginModule`.

For example, a `LoginModule` developed by IBM might be called `com.ibm.auth.Module` where `com.ibm.auth` is the package name and `Module` is the name of the `LoginModule` class implementation.

Step 3: Implement the LoginModule Interface

The `LoginModule` interface specifies five abstract methods that you must implement:

- `initialize`
- `login`
- `commit`
- `abort`
- `logout`

In addition to these methods, a `LoginModule` implementation must provide a public constructor with no arguments. This allows for its proper instantiation by a `LoginContext`. Note that if no constructor is provided in your `LoginModule` implementation, a default no-argument constructor is automatically inherited from the `Object` class.

> **Note:**
>
> If you don't implement the `LoginModule` interface, then a `LoginException` will be thrown when you try to use your login module.

LoginModule.initialize Method

```
public void initialize(
    Subject subject,
    CallbackHandler handler,
    Map<java.lang.String, ?> sharedState,
    Map<java.lang.String, ?> options);
```

The `initialize` method is called to initialize the `LoginModule` with the relevant authentication and state information.

This method is called by a `LoginContext` immediately after this `LoginModule` has been instantiated, and prior to any calls to its other public methods. The method implementation should store away the provided arguments for future use.

The `initialize` method may additionally peruse the provided *sharedState* to determine what additional authentication state it was provided by other `LoginModules`, and may also traverse through the provided *options* to determine what configuration options were specified to affect the `LoginModule`'s behavior. It may save option values in variables for future use.

Below is a list of options commonly supported by `LoginModules`. Note that the following is simply a guideline. Modules are free to support a subset (or none) of the following options.

- `tryFirstPass` - If `true`, the first `LoginModule` in the stack saves the password entered, and subsequent `LoginModules` also try to use it. If authentication fails, the `LoginModules` prompt for a new password and retry the authentication.
- `useFirstPass` - If `true`, the first `LoginModule` in the stack saves the password entered, and subsequent `LoginModules` also try to use it. `LoginModules` do not prompt for a new password if authentication fails (authentication simply fails).
- `tryMappedPass` - If `true`, the first `LoginModule` in the stack saves the password entered, and subsequent `LoginModules` attempt to map it into their service-specific password. If authentication fails, the `LoginModules` prompt for a new password and retry the authentication.
- `useMappedPass` - If `true`, the first `LoginModule` in the stack saves the password entered, and subsequent `LoginModules` attempt to map it into their service-specific password. `LoginModules` do not prompt for a new password if authentication fails (authentication simply fails).
- `moduleBanner` - If `true`, then when invoking the `CallbackHandler`, the `LoginModule` provides a `TextOutputCallback` as the first `Callback`, which describes the `LoginModule` performing the authentication.
- `debug` - If `true`, instructs a `LoginModule` to output debugging information.

The `initialize` method may freely ignore state or options it does not understand, although it would be wise to log such an event if it does occur.

Note that the `LoginContext` invoking this `LoginModule` (and the other configured `LoginModules`, as well), all share the same references to the provided `Subject` and `sharedState`. Modifications to the `Subject` and `sharedState` will, therefore, be seen by all.

LoginModule.login Method

```
boolean login() throws LoginException;
```

The `login` method is called to authenticate a `Subject`. This is phase 1 of authentication.

This method implementation should perform the actual authentication. For example, it may cause prompting for a user name and password, and then attempt to verify the password against a password database. Another example implementation may inform the user to insert their finger into a fingerprint reader, and then match the input fingerprint against a fingerprint database.

If your `LoginModule` requires some form of user interaction (retrieving a user name and password, for example), it should not do so directly. That is because there are various ways of communicating with a user, and it is desirable for `LoginModules` to remain independent of the different types of user interaction. Rather, the `LoginModule`'s `login` method should invoke the `handle` method of the `CallbackHandler` interface passed to the `initialize` method to perform the user interaction and set appropriate results, such as the user name and password. The `LoginModule` passes the `CallbackHandler` an array of appropriate `Callbacks`, for example a `NameCallback` for the user name and a `PasswordCallback` for the password, and the `CallbackHandler` performs the requested user interaction and sets appropriate values in the `Callbacks`. For example, to process a `NameCallback`, the `CallbackHandler` may prompt for a name, retrieve the value from the user, and call the `NameCallback`'s `setName` method to store the name.

The authentication process may also involve communication over a network. For example, if this method implementation performs the equivalent of a *kinit* in Kerberos, then it would need to contact the KDC. If a password database entry itself resides in a remote naming service, then that naming service needs to be contacted, perhaps via the Java Naming and Directory Interface (JNDI). Implementations might also interact with an underlying operating system. For example, if a user has already logged into an operating system like Solaris, Linux, macOS, or Windows, this method might simply import the underlying operating system's identity information.

The `login` method should

1. Determine whether or not this `LoginModule` should be ignored. One example of when it should be ignored is when a user attempts to authenticate under an identity irrelevant to this `LoginModule` (if a user attempts to authenticate as *root* using NIS, for example). If this `LoginModule` should be ignored, `login` should return `false`. Otherwise, it should do the following:
2. Call the `CallbackHandler` `handle` method if user interaction is required.
3. Perform the authentication.
4. Store the authentication result (success or failure).
5. If authentication succeeded, save any relevant state information that may be needed by the `commit` method.
6. Return `true` if authentication succeeds, or throw a `LoginException` such as `FailedLoginException` if authentication fails.

Note that the `login` method implementation should not associate any new `Principal` or credential information with the saved `Subject` object. This method merely performs

the authentication, and then stores away the authentication result and corresponding authentication state. This result and state will later be accessed by the `commit` or `abort` method. Note that the result and state should typically not be saved in the *sharedState* `Map`, as they are not intended to be shared with other `LoginModules`.

An example of where this method might find it useful to store state information in the *sharedState* `Map` is when `LoginModules` are configured to share passwords. In this case, the entered password would be saved as shared state. By sharing passwords, the user only enters the password once, and can still be authenticated to multiple `LoginModules`. The standard conventions for saving and retrieving names and passwords from the *sharedState* `Map` are the following:

- `javax.security.auth.login.name` - Use this as the shared state map key for saving/retrieving a name. The value should be a `String`.
- `javax.security.auth.login.password` - Use this as the shared state map key for saving/retrieving a password. The value should be a `char` array.

If authentication fails, the `login` method should not retry the authentication. This is the responsibility of the application. Multiple `LoginContext login` method calls by an application are preferred over multiple login attempts from within `LoginModule.login()`.

LoginModule.commit Method

```
boolean commit() throws LoginException;
```

The `commit` method is called to commit the authentication process. This is phase 2 of authentication when phase 1 succeeds. It is called if the `LoginContext`'s overall authentication succeeded (that is, if the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL `LoginModules` succeeded.)

This method should access the authentication result and corresponding authentication state saved by the `login` method.

If the authentication result denotes that the `login` method failed, then this `commit` method should remove/destroy any corresponding state that was originally saved.

If the saved result instead denotes that this `LoginModule`'s `login` method succeeded, then the corresponding state information should be accessed to build any relevant `Principal` and credential information. Such `Principals` and credentials should then be added to the `Subject` stored away by the `initialize` method.

After adding `Principals` and credentials, dispensable state fields should be destroyed expeditiously. Likely fields to destroy would be user names and passwords stored during the authentication process.

The `commit` method should save private state indicating whether the commit succeeded or failed.

The following chart depicts what a `LoginModule`'s `commit` method should return. The different boxes represent the different situations that may occur. For example, the top-left corner box depicts what the `commit` method should return if both the previous call to `login` succeeded and the `commit` method itself succeeded.

Table 6-2 LoginModule.commit Method Return Values

Login Status	COMMIT: SUCCESS	COMMIT: FAILURE
LOGIN: SUCCESS	return TRUE	throw EXCEPTION
LOGIN: FAILURE	return FALSE	return FALSE

LoginModule.abort Method

```
boolean abort() throws LoginException;
```

The `abort` method is called to abort the authentication process. This is phase 2 of authentication when phase 1 fails. It is called if the `LoginContext`'s overall authentication failed.

This method first accesses this `LoginModule`'s authentication result and corresponding authentication state saved by the `login` (and possibly `commit`) methods, and then clears out and destroys the information. Sample state to destroy would be user names and passwords.

If this `LoginModule`'s authentication attempt failed, then there shouldn't be any private state to clean up.

The following charts depict what a `LoginModule`'s `abort` method should return. This first chart assumes that the previous call to `login` succeeded. For instance, the `abort` method should return TRUE if both the previous call to `login` and `commit` succeeded, and the `abort` method itself also succeeded.

Table 6-3 LoginModule.abort Method Return Values: Login Succeeded

Login Status	ABORT: SUCCESS	ABORT: FAILURE
COMMIT: SUCCESS	return TRUE	throw EXCEPTION
COMMIT: FAILURE	return TRUE	throw EXCEPTION

The second chart depicts what a `LoginModule`'s abort method should return, assuming that the previous call to login failed. For instance, the `abort` method should return FALSE if the previous call to `login` failed, the previous call to `commit` succeeded, and the `abort` method itself also succeeded.

Table 6-4 LoginModule.abort Method Return Values: Login Failed

Login Status	ABORT: SUCCESS	ABORT: FAILURE
COMMIT: SUCCESS	return FALSE	return FALSE
COMMIT: FAILURE	return FALSE	return FALSE

LoginModule.logout Method

```
boolean logout() throws LoginException;
```

The `logout` method is called to log out a `Subject`.

This method removes `Principals`, and removes/destroys credentials associated with the `Subject` during the `commit` operation. This method should not touch those `Principals` or credentials previously existing in the `Subject`, or those added by other `LoginModules`.

If the `Subject` has been marked *read-only* (the `Subject`'s `isReadOnly` method returns *true*), then this method should only destroy credentials associated with the `Subject` during the `commit` operation (removing the credentials is not possible). If the `Subject` has been marked as *read-only* and the credentials associated with the `Subject` during the `commit` operation are not destroyable (they do not implement the `Destroyable` interface), then this method may throw a `LoginException`.

The `logout` method should return `true` if logout succeeds, or otherwise throw a `LoginException`.

Step 4: Choose or Write a Sample Application

Either choose an existing sample application for your testing, or write a new one.

See Java Authentication and Authorization Service (JAAS) Reference Guide for information about application requirements and a sample application you can use for your testing.

Step 5: Compile the LoginModule and Application

Compile your new `LoginModule` and the application you will use for testing.

Step 6: Prepare for Testing

Step 6a: Place Your `LoginModule` and Application Code in JAR Files

Place your `LoginModule` and application code in separate JAR files, in preparation for referencing the JAR files in the policy in Step 6b: Set `LoginModule` and Application JAR File Permissions. Here is a sample command for creating a JAR file:

```
jar cvf <JAR file name> <list of classes, separated by spaces>
```

This command creates a JAR file with the specified name containing the specified classes.

For more information on the **jar** tool, see jar.

Step 6b: Set `LoginModule` and Application JAR File Permissions

If your `LoginModule` and/or application performs security-sensitive tasks that will trigger security checks (making network connections, reading or writing files on a local disk, etc.), it will need to be granted the required permissions if it is run while a security manager is installed; see Permissions in the JDK.

Since `LoginModules` usually associate `Principals` and credentials with an authenticated Subject, some types of permissions a `LoginModule` will typically require are `AuthPermission`s with target names "modifyPrincipals", "modifyPublicCredentials", and "modifyPrivateCredentials".

A sample statement granting permissions to a `LoginModule` whose code is in `MyLM.jar` appears below. Such a statement could appear in a policy file. In this example, the `MyLM.jar` file is assumed to be in the `/localWork` directory.

```
grant codeBase "file:/localWork/MyLM.jar" {
    permission javax.security.auth.AuthPermission "modifyPrincipals";
    permission javax.security.auth.AuthPermission
"modifyPublicCredentials";
    permission javax.security.auth.AuthPermission
"modifyPrivateCredentials";
};
```

> **Note:**
>
> Since a `LoginModule` is always invoked within an `AccessController.doPrivileged` call, it should not have to call `doPrivileged` itself. If it does, it may inadvertently open up a security hole. For example, a `LoginModule` that invokes the application-provided `CallbackHandler` inside a `doPrivileged` call opens up a security hole by permitting the application's `CallbackHandler` to gain access to resources it would otherwise not have been able to access.

Step 6c: Create a Configuration Referencing the LoginModule

Because JAAS supports a pluggable authentication architecture, your new `LoginModule` can be used without requiring modifications to existing applications. Only the login `Configuration` needs to be updated in order to indicate use of a new `LoginModule`.

The default `Configuration` implementation from Oracle reads configuration information from configuration files, as described in `ConfigFile`.

Create a configuration file to be used for testing. For example, to configure the previously-mentioned hypothetical IBM `LoginModule` for an application, the configuration file might look like this:

```
AppName {
    com.ibm.auth.Module REQUIRED debug=true;
};
```

where `AppName` should be whatever name the application uses to refer to this entry in the login configuration file. The application specifies this name as the first argument to the `LoginContext` constructor.

Step 7: Test Use of the LoginModule

Test your application and its use of the `LoginModule`. When you run the application, specify the login configuration file to be used. For example, suppose your application is named `MyApp`, it is located in `MyApp.jar`, and your configuration file is `test.conf`.

You could run the application and specify the configuration file via the following:

```
java -classpath MyApp.jar
 -Djava.security.auth.login.config=test.conf MyApp
```

Type all that on one line. Multiple lines are used here for legibility.

To specify a policy file named `my.policy` and run the application with a security manager installed, do the following:

```
java -classpath MyApp.jar -Djava.security.manager
 -Djava.security.policy=my.policy
 -Djava.security.auth.login.config=test.conf MyApp
```

Again, type all that on one line.

You may want to configure the `LoginModule` with a *debug* option to help ensure that it is working correctly.

Debug your code and continue testing as needed. If you have problems, review the steps above and ensure they are all completed.

Be sure to vary user input and the `LoginModule` options specified in the configuration file.

Be sure to also include testing using different installation options (e.g., placing the `LoginModule` on the class path or module path) and execution environments (with or without a security manager running). In particular, in order to ensure your `LoginModule` works when a security manager is installed and the `LoginModule`, you need to test such an installation and execution environment, after granting required permissions, as described in Step 6b: Set `LoginModule` and Application JAR File Permissions.

1. If you find during testing that your `LoginModule` or application needs modifications, make the modifications, recompile (Step 5: Compile the LoginModule and Application).
2. Place the updated code in a JAR file (Step 6a: Place Your `LoginModule` and Application Code in JAR Files).
3. If needed fix or add to the permissions (Step 6b: Set `LoginModule` and Application JAR File Permissions).
4. If needed modify the login configuration file (Step 6c: Create a Configuration Referencing the LoginModule).
5. Re-run the application and repeat these steps as needed.

Step 8: Document Your LoginModule Implementation

Write documentation for clients of your `LoginModule`.

Example documentation you may want to include is:

- A README or User Guide describing
  1. The authentication process employed by your `LoginModule` implementation.
  2. Information on how to install the `LoginModule`.
  3. Configuration options accepted by the `LoginModule`. For each option, specify the option name and possible values (or types of values), as well as the behavior the option controls.
  4. The permissions required by your `LoginModule` when it is run with a security manager.
- An example `Configuration` file that references your new `LoginModule`.
- An example policy file granting your `LoginModule` the required permissions.
- API documentation. Putting Javadoc comments into your source code as you write it will make the Javadoc API documentation easy to generate.

Step 9: Make LoginModule JAR File and Documents Available

Make your `LoginModule` JAR file and documentation available to clients.

7

Java Generic Security Services (Java GSS-API)

Java Generic Security Services (Java GSS-API) is used for securely exchanging messages between communicating applications.

Introduction to JAAS and Java GSS-API Tutorials is a series of tutorials demonstrating various aspects of the use of Java Authentication and Authorization Service (JAAS) and Java GSS-API.

Single Sign-on Using Kerberos in Java discusses how to use Single Sign-On based on the Kerberos V5 protocol.

Advanced Security Programming in Java SE Authentication, Secure Communication and Single Sign-On shows you how to use the Java SE GSS APIs to build applications that authenticate their users, communicate securely with other applications and services, and configure your applications in a Kerberos environment to achieve Single Sign-On.

The Kerberos 5 GSS-API Mechanism describes and lists security features regarding Java Generic Security Services (Java GSS) for Kerberos 5.

Introduction to JAAS and Java GSS-API Tutorials

This page links to a series of tutorials demonstrating various aspects of the use of JAAS (Java Authentication and Authorization Service) and Java GSS-API.

JAAS can be used for two purposes:

- for *authentication* of users, to reliably and securely determine who is currently executing Java code, and
- for *authorization* of users to ensure they have the access control rights (permissions) required to do security-sensitive operations.

Java GSS-API is used for *securely exchanging messages* between communicating applications. The Java GSS-API contains the Java bindings for the Generic Security Services Application Program Interface (GSS-API) defined in RFC 5653. GSS-API offers application programmers uniform access to security services atop a variety of underlying security mechanisms, including Kerberos.

Note: JSSE is another API that can be used for secure communication. For the differences between the two, see When to Use Java GSS-API Versus JSSE.

The reason both JAAS and Java GSS-API tutorials are presented together is because JAAS authentication is typically performed prior to secure communication using Java GSS-API. Thus JAAS and Java GSS-API are related and often used together. However, it is possible for applications to use JAAS without Java GSS-API, and it is also possible to use Java GSS-API without JAAS. Furthermore, JAAS itself can be used simply for authentication or for both authentication and authorization.

The following tutorials provide working examples for all of the scenarios described above.

1. Use of Java GSS-API for Secure Message Exchanges Without JAAS Programming

   Demonstrates the use of the Java GSS-API for secure message exchanges between a client application and a server application.

2. JAAS Authentication

   Explains how an application can authenticate users using JAAS.

3. JAAS Authorization

   Explains how to enforce user-based access controls using JAAS.

4. Use of JAAS Login Utility

   Describes a utility program that authenticates a user using JAAS and executes any application as that user. The appropriate user-based access controls are enforced while the application executes. This utility, as a convenience, essentially performs the operations described in the JAAS Authentication and JAAS Authorization tutorials on your behalf. Therefore it is possible to skip directly to this tutorial if you do not need to know how to perform JAAS authentication and authorization directly.

5. Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges

   The most comprehensive tutorial. The Login utility is used to authenticate a service user and to start up a server application as that user. The Login utility is also used to authenticate a client user and to start up a client application as that user. Finally the client and server applications, on behalf of their authenticated client and service users, exchange secure messages using the Java GSS-API.

6. More Things You Can Do with Java GSS-API and JAAS

   Shows additional operations the server application in the previous tutorial can perform once communication has been established with the client application.

All applications in all tutorials in this series utilize Kerberos Version 5 as the underlying technology for authentication and secure communication. See Kerberos Requirements. The term "Kerberos" used throughout the tutorials is meant to refer to Kerberos Version 5.

When to Use Java GSS-API Versus JSSE

Java GSS-API and JSSE provide you with the same basic set of security features:

1. Client-server authentication
2. Encryption and integrity protection of transmitted data

However, there are some key differences between the two. This document lists some of them to help you decide which might be more appropriate in your environment:

1. **Kerberos Single Sign-On Support**

   GSS-API contains support for Kerberos as a mandatory security mechanism. This means that if your desktop has Kerberos support, you can write Java GSS-API based applications that never prompt the user for a password.

   RFC 2712 defined Kerberos Cipher Suites for TLS, but the document is out of date and does not support modern encryption types like AES.

2. **Communications API**

   JSSE supports a socket-based API. JSSE sockets extend the socket classes found in `java.net` and JSSE socket factories extend the socket factories found in `javax.net`. Thus, if your application is written such that its security needs to be configured via a socket factory, then JSSE might be more appropriate for you. JSSE sockets need to use some reliable transport. Typically, implementations use TCP.

   Java GSS-API, on the other hand, is a token-based API that relies on the application to do the communication. This means that the application can use TCP sockets, UDP datagrams, or any other channel that will allow it to transport Java GSS-API generated tokens. If your application has varying communication protocol needs, then Java GSS-API might be more appropriate for you. Java GSS-API can read and write its tokens using input and output streams. However, you will need to set up the streams yourself.

3. **Credential Delegation**

   Java GSS-API allows the client to delegate its credentials to the server when using Kerberos. If your application will be deployed in a multi-tier environment where intermediaries need to impersonate clients when talking to backend layers, Java GSS-API might be more appropriate for you.

4. **Selective Encryption**

   Because Java GSS-API is token-based, you can choose to selectively encrypt certain messages but not all. If your application needs to intersperse plaintext and ciphertext messages, Java GSS-API might be more appropriate for you.

5. **Protocol Requirements**

   JSSE provides an implementation of the TLS protocol defined in RFC 5246. Java GSS-API provides an implementation of the GSS-API framework defined in RFC 5653, as well as an implementation of the Kerberos Version 5 mechanism defined in RFC 1964. (On Microsoft Windows platforms, this may be known as SSPI with Kerberos.) Some servers such as HTTPS servers will require you to use TLS, in which case JSSE will be appropriate for you. Other servers such as LDAP servers using SASL might need GSS-API with Kerberos, in which case Java GSS-API will be appropriate for you.

Use of Java GSS-API for Secure Message Exchanges Without JAAS Programming

This tutorial presents two sample applications demonstrating the use of the Java GSS-API for secure exchanges of messages between communicating applications, in this case a client application and a server application.

Java GSS-API uses what is called a "security mechanism" to provide these services. The GSS-API implementation contains support for the Kerberos V5 mechanism in addition to any other vendor-specific choices. The Kerberos V5 mechanism is used for this tutorial.

In order to perform authentication between the client and server and to establish cryptographic keys for secure communication, a GSS-API mechanism needs access to certain credentials for the local entity on each side of the connection. In our case, the credential used on the client side consists of a Kerberos ticket, and on the server side, it consists of a long-term Kerberos secret key. Kerberos tickets can optionally

include the host address and IPv4 and IPv6 host addresses are both supported. Java GSS-API requires that the mechanism obtain these credentials from the Subject associated with the thread's access control context.

To populate a Subject with such credentials, client and server applications typically will first perform JAAS authentication using a Kerberos `LoginModule`. The JAAS Authentication tutorial demonstrates how to do this. The JAAS Authorization tutorial then demonstrates how to associate the authenticated Subject with the thread's access control context. A utility has also been written as a convenience to automatically perform those operations on your behalf. The Use of JAAS Login Utility tutorial demonstrates how to use the Login utility.

For this tutorial, we will not have the client and server perform JAAS authentication, nor will we have them use the Login utility. Instead, we will rely on setting the system property `javax.security.auth.useSubjectCredsOnly` to `false`, which allows us to relax the restriction of requiring a GSS mechanism to obtain necessary credentials from an existing Subject, set up by JAAS. See The useSubjectCredsOnly System Property.

> **Note:**
>
> This is a simplified introductory tutorial. For example, we do not include any policy files or run the sample code using a security manager. In real life, code using Java GSS-API should be run with a security manager, so that security-sensitive operations would not be allowed unless the required permissions were explicitly granted.

There is another tutorial, Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges, that is just like the tutorial you are reading except that it utilizes the Login utility, policy files, and a more complex login configuration file. A login configuration file (see Appendix B: JAAS Login Configuration File), required whenever JAAS authentication is done, specifies the desired authentication module.

As with all tutorials in this series, the underlying technology used to support authentication and secure communication for the applications in this tutorial is Kerberos V5. See Kerberos Requirements.

- Overview of the Client and Server Applications
- The SampleClient and SampleServer Code
- Kerberos User and Service Principal Names
- The Login Configuration File
- The useSubjectCredsOnly System Property
- Running the SampleClient and SampleServer Programs

If you want to first see the tutorial code in action, you can skip directly to Running the SampleClient and SampleServer Programs and then go back to the other sections to learn more.

Overview of the Client and Server Applications

The applications for this tutorial are named `SampleClient` and `SampleServer`.

Here is a summary of execution of the `SampleClient` and `SampleServer` applications:

1. Run the `SampleServer` application. `SampleServer`
   a. Reads its argument, the port number that it should listen on for client connections.
   b. Creates a `ServerSocket` for listening for client connections on that port.
   c. Listens for a connection.
2. Run the `SampleClient` application (possibly on a different machine). `SampleClient`
   a. Reads its arguments: (1) The name of the Kerberos principal that represents `SampleServer` (see Kerberos User and Service Principal Names), (2) the name of the host (machine) on which `SampleServer` is running, and (3) the port number on which `SampleServer` listens for client connections.
   b. Attempts a socket connection with the `SampleServer`, using the host and port it was passed as arguments.
3. The socket connection is accepted by `SampleServer` and both applications initialize a `DataInputStream` and a `DataOutputStream` from the socket input and output streams, to be used for future data exchanges.
4. `SampleClient` and `SampleServer` each instantiate a `GSSContext` and follow a protocol for establishing a shared context that will enable subsequent secure data exchanges.
5. `SampleClient` and `SampleServer` can now securely exchange messages.
6. When `SampleClient` and `SampleServer` are done exchanging messages, they perform clean-up operations.

The actual code and further details are presented in the following sections.

The SampleClient and SampleServer Code

The entire code for both the `SampleClient.java` and `SampleServer.java` programs resides in their `main` methods and can be broken down into the following subparts:

1. Obtaining the Command-Line Arguments
2. Establishing a Socket Connection for Message Exchanges
3. Establishing a Security Context
4. Exchanging Messages Securely
5. Clean Up

> **Note:**
>
> The Java GSS-API classes utilized by these programs (`GSSManager`, `GSSContext`, `GSSName`, `GSSCredential`, `MessageProp`, and `Oid`) are found in the `org.ietf.jgss` package.

Obtaining the Command-Line Arguments

The first thing both our client and server `main` methods do is read the command-line arguments.

Arguments Read by SampleClient

`SampleClient` expects three arguments:

1. A service principal name – The name of the Kerberos principal that represents `SampleServer` (see Kerberos User and Service Principal Names).
2. A host name – The machine on which `SampleServer` is running.
3. A port number – The port number of the port on which `SampleServer` listens for connections.

Here is the code for reading the command-line arguments:

```
if (args.length < 3) {
    System.out.println("Usage: java <options> Login SampleClient "
       + " <servicePrincipal> <hostName> <port>");
    System.exit(-1);
}

String server = args[0];
String hostName = args[1];
int port = Integer.parseInt(args[2]);
```

Argument Read by SampleServer

`SampleServer` expects just one argument:

- A local port number – The port number used by `SampleServer` for listening for connections with clients. This number should be the same as the port number specified when running the `SampleClient` program.

Here is the code for reading the command-line argument:

```
if (args.length != 1) {
    System.out.println(
        "Usage: java <options> Login SampleServer <localPort>");
    System.exit(-1);
}

int localPort = Integer.parseInt(args[0]);
```

Establishing a Socket Connection for Message Exchanges

Java GSS-API provides methods for creating and interpreting tokens (opaque byte data). The tokens contain messages to be securely exchanged between two peers, but the method of actual token transfer is up to the peers. For our `SampleClient` and `SampleServer` applications, we establish a socket connection between the client and server and exchange data using the socket input and output streams.

SampleClient Code for Socket Connection

`SampleClient` was passed as arguments the name of the host machine `SampleServer` is running on, as well as the port number on which `SampleServer` will be listening for connections, so `SampleClient` has all it needs to establish a socket connection with `SampleServer`. It uses the following code to set up the connection and initialize a `DataInputStream` and a `DataOutputStream` for future data exchanges:

```
Socket socket = new Socket(hostName, port);

DataInputStream inStream =
  new DataInputStream(socket.getInputStream());
DataOutputStream outStream =
  new DataOutputStream(socket.getOutputStream());

System.out.println("Connected to server "
   + socket.getInetAddress());
```

SampleServer Code for Socket Connection

The `SampleServer` application was passed as an argument the port number to be used for listening for connections from clients. It creates a `ServerSocket` for listening on that port:

```
ServerSocket ss = new ServerSocket(localPort);
```

The `ServerSocket` can then wait for and accept a connection from a client, and then initialize a `DataInputStream` and a `DataOutputStream` for future data exchanges with the client :

```
Socket socket = ss.accept();

DataInputStream inStream =
    new DataInputStream(socket.getInputStream());
DataOutputStream outStream =
    new DataOutputStream(socket.getOutputStream());

System.out.println("Got connection from client "
    + socket.getInetAddress());
```

The `accept` method waits until a client (in our case, `SampleClient`) requests a connection on the host and port of the `SampleServer`, which `SampleClient` does via

```
Socket socket = new Socket(hostName, port);
```

When the connection is requested and established, the `accept` method returns a new `Socket` object bound to a new port. The server can communicate with the client over this new socket and continue to listen for other client connection requests on the `ServerSocket` bound to the original port. Thus, a server program typically has a loop which can handle multiple connection requests.

The basic loop structure for our `SampleServer` is the following:

```
while (true) {

    Socket socket = ss.accept();

    <Establish input and output streams for the connection>;
    <Establish a context with the client>;
    <Exchange messages with the client>;
    <Clean up>;
}
```

Client connections are queued at the original port, so with this program structure used by `SampleServer`, the interaction with the first client making a connection has to complete before the next connection can be accepted. The server could actually service multiple clients simultaneously through the use of threads – one thread per client connection, as in

```
while (true) {
    <accept a connection>;
    <create a thread to handle the client>;
}
```

Establishing a Security Context

Before two applications can use Java GSS-API to securely exchange messages between them, they must establish a joint security context using their credentials. (Note: In the case of `SampleClient`, the credentials were established when the Login utility authenticated the user on whose behalf the `SampleClient` was run, and similarly for `SampleServer`.) The security context encapsulates shared state information that might include, for example, cryptographic keys. One use of such keys might be to encrypt messages to be exchanged, if encryption is requested.

As part of the security context establishment, the context initiator (in our case, `SampleClient`) is authenticated to the acceptor (`SampleServer`), and may require that the acceptor also be authenticated back to the initiator, in which case we say that "mutual authentication" took place.

Both applications create and use a `GSSContext` object to establish and maintain the shared information that makes up the security context.

The instantiation of the context object is done differently by the context initiator and the context acceptor. After the initiator instantiates a `GSSContext`, it may choose to set various context options that will determine the characteristics of the desired security context, for example, specifying whether or not mutual authentication should take place. After all the desired characteristics have been set, the initiator calls the `initSecContext` method, which produces a token required by the acceptor's `acceptSecContext` method.

While Java GSS-API methods exist for preparing tokens to be exchanged between applications, *it is the responsibility of the applications to actually transfer the tokens between them.* So after the initiator has received a token from its call to `initSecContext`, it sends that token to the acceptor. The acceptor calls `acceptSecContext`, passing it the token. The `acceptSecContext` method may in turn return a token. If it does, the acceptor should send that token to the initiator, which

should then call `initSecContext` again and pass it this token. Each time `initSecContext` or `acceptSecContext` returns a token, the application that called the method should send the token to its peer and that peer should pass the token to its appropriate method (`acceptSecContext` or `initSecContext`). This continues until the context is fully established (which is the case when the context's `isEstablished` method returns `true`).

The context establishment code for our sample applications is described in the following:

- Context Establishment by SampleClient
- Context Establishment by SampleServer

Context Establishment by SampleClient

In our client/server scenario, SampleClient is the context initiator. Here are the basic steps it takes to establish a security context:

1. SampleClient GSSContext Instantiation
2. SampleClient Setting of Desired Options
3. SampleClient Context Establishment Loop: Loops while the context is not yet established, each time calling `initSecContext`, sending any returned token to `SampleServer`, and receiving a token (if any) from `SampleServer`.

SampleClient GSSContext Instantiation

A `GSSContext` is created by instantiating a `GSSContext` and then calling one of its `createContext` methods. The `GSSManager` class serves as a factory for other important GSS API classes. It can create instances of classes implementing the `GSSContext`, `GSSCredential`, and `GSSName` interfaces.

`SampleClient` obtains an instance of the default `GSSManager` subclass by calling the `GSSManager` static method `getInstance`:

```
GSSManager manager = GSSManager.getInstance();
```

The default `GSSManager` subclass is one whose `create*` methods (`createContext`, etc.) return classes whose implementations support Kerberos as the underlying technology.

The `GSSManager` factory method for creating a context on the initiator's side has the following signature:

```
GSSContext createContext(GSSName peer, Oid mech,
            GSSCredential myCred, int lifetime);
```

The arguments are described below, followed by the complete call to `createContext`.

The GSSName peer Argument

The peer in our client/server paradigm is the server. For the `peer` argument, we need a `GSSName` for the service principal representing the server. (See Kerberos User and Service Principal Names.) A `String` for the service principal name is passed as the first argument to `SampleClient`, which places the argument into its local `String` variable named `server`. The `GSSManager` `manager` is used to instantiate a `GSSName`

by calling one of its `createName` methods. `SampleClient` calls the `createName` method with the following signature:

```
GSSName createName(String nameStr, Oid nameType);
```

`SampleClient` passes the `server` `String` for the `nameStr` argument.

The second argument is an `Oid`. An `Oid` represents a Universal Object Identifier. Oids are hierarchically globally-interpretable identifiers used within the GSS-API framework to identify mechanisms and name types. The structure and encoding of Oids is defined in the ISOIEC-8824 and ISOIEC-8825 standards. The `Oid` passed to the `createName` method is specifically a name type `Oid` (not a mechanism Oid).

In GSS-API, string names are often mapped from a mechanism-independent format into a mechanism-specific format. Usually, an Oid specifies what name format the string is in so that the mechanism knows how to do this mapping. Passing in a `null` `Oid` indicates that the name is already in a native format that the mechanism uses. This is the case for the `server` String; it is in the appropriate format for a Kerberos Version 5 name. Thus, `SampleClient` passes a `null` for the `Oid`. Here is the call:

```
GSSName serverName = manager.createName(server, null);
```

The Oid mech Argument

The second argument to the `GSSManager` `createContext` method is an Oid representing the mechanism to be used for the authentication between the client and the server during context establishment and for subsequent secure communication between them.

Our tutorial will use Kerberos V5 as the security mechanism. The Oid for the Kerberos V5 mechanism is defined in RFC 1964 as "1.2.840.113554.1.2.2" so we create such an `Oid`:

```
Oid krb5Oid = new Oid("1.2.840.113554.1.2.2");
```

`SampleClient` passes `krb5Oid` as the second argument to `createContext`.

The GSSCredential myCred Argument

The third argument to the `GSSManager` `createContext` method is a `GSSCredential` representing the caller's credentials. If you pass `null` for this argument, as `SampleClient` does, the default credentials are used.

The int lifetime Argument

The final argument to the `GSSManager` `createContext` method is an `int` specifying the desired lifetime, in seconds, for the context that is created. `SampleClient` passes `GSSContext.DEFAULT_LIFETIME` to request a default lifetime.

The Complete createContext Call

Now that we have all the required arguments, here is the call `SampleClient` makes to create a `GSSContext`:

```
GSSContext context =
    manager.createContext(serverName,
                          krb5Oid,
```

```
                                        null,
                                        GSSContext.DEFAULT_LIFETIME);
```

SampleClient Setting of Desired Options

After instantiating a context, and prior to actually establishing the context with the context acceptor, the context initiator may choose to set various options that determine the desired security context characteristics. Each such option is set by calling a `request` method on the instantiated context. Most `request` methods take a `boolean` argument for indicating whether or not the feature is requested. It is not always possible for a request to be satisfied, so whether or not it was can be determined after context establishment by calling one of the `get` methods.

`SampleClient` requests the following:

1. **Mutual authentication**. The context initiator is always authenticated to the acceptor. If the initiator requests mutual authentication, then the acceptor is also authenticated to the initiator.
2. **Confidentiality**. Requesting confidentiality means that you request the *enabling* of encryption for the context method named `wrap`. Encryption is actually used only if the `MessageProp` object passed to the `wrap` method requests privacy.
3. **Integrity**. This requests integrity for the `wrap` and `getMIC` methods. When integrity is requested, a cryptographic tag known as a Message Integrity Code (MIC) will be generated when calling those methods. When `getMIC` is called, the generated MIC appears in the returned token. When `wrap` is called, the MIC is packaged together with the message (the original message or the result of encrypting the message, depending on whether confidentiality was applied) all as part of one token. You can subsequently verify the MIC against the message to ensure that the message has not been modified in transit.

The `SampleClient` code for making these requests on the `GSSException context` is the following:

```
context.requestMutualAuth(true);  // Mutual authentication
context.requestConf(true);  // Will use encryption later
context.requestInteg(true); // Will use integrity later
```

After the context is established, the client must explicitly check the context states by calling the accesor methods, like `getMutualAuthState`, `getConfState`, or `getIntegState`, and destroy the security context if any of them do not match the desired state.

> **Note:**
>
> When using the default `GSSManager` implementation and the Kerberos mechanism, these requests will always be granted.

SampleClient Context Establishment Loop

After `SampleClient` has instantiated a `GSSContext` and specified the desired context options, it can actually establish the security context with `SampleServer`. To do so, `SampleClient` has a loop. Each loop iteration

1. Calls the context's `initSecContext` method. If this is the first call, the method is passed a `null` token. Otherwise, it is passed the token most recently sent to `SampleClient` by `SampleServer` (a token generated by a `SampleServer` call to `acceptSecContext`).
2. Sends the token returned by `initSecContext` (if any) to `SampleServer`. The first call to `initSecContext` always produces a token. The last call might not return a token.
3. Checks to see if the context is established. If not, `SampleClient` receives another token from `SampleServer` and then starts the next loop iteration.

The tokens returned by `initSecContext` or received from `SampleServer` are placed in a byte array. Tokens should be treated by `SampleClient` and `SampleServer` as opaque data to be passed between them and interpreted by Java GSS-API methods.

The `initSecContext` arguments are a byte array containing a token, the starting offset into that array of where the token begins, and the token length. For the first call, `SampleClient` passes a null token, since no token has yet been received from `SampleServer`.

To exchange tokens with `SampleServer`, `SampleClient` uses the `DataInputStream` `inStream` and `DataOutputStream` `outStream` it previously set up using the input and output streams for the socket connection made with `SampleServer`. Note that whenever a token is written, the number of bytes in the token is written first, followed by the token itself. The reasons are discussed in the introduction to the The SampleClient and SampleServer Message Exchanges section.

Here is the `SampleClient` context establishment loop, followed by code displaying information about who the client and server are and whether or not mutual authentication actually took place:

```
byte[] token = new byte[0];

while (!context.isEstablished()) {

    // token is ignored on the first call
    token = context.initSecContext(token, 0, token.length);

    // Send a token to the server if one was generated by
    // initSecContext
    if (token != null) {
        System.out.println("Will send token of size "
                   + token.length + " from initSecContext.");
        outStream.writeInt(token.length);
        outStream.write(token);
        outStream.flush();
    }

    // If the client is done with context establishment
    // then there will be no more tokens to read in this loop
    if (!context.isEstablished()) {
        token = new byte[inStream.readInt()];
        System.out.println("Will read input token of size "
                   + token.length
                   + " for processing by initSecContext");
        inStream.readFully(token);
```

```
    }
}

System.out.println("Context Established! ");
System.out.println("Client is " + context.getSrcName());
System.out.println("Server is " + context.getTargName());
if (context.getMutualAuthState())
    System.out.println("Mutual authentication took place!");
```

Context Establishment by SampleServer

In our client/server scenario, `SampleServer` is the context acceptor. Here are the basic steps it takes to establish a security context:

1. SampleServer GSSContext Instantiation
2. SampleClient Context Establishment Loop: Loops while the context is not yet established, each time receiving a token from `SampleClient`, calling `acceptSecContext` and passing it the token, and sending any returned token to `SampleClient`.

SampleServer GSSContext Instantiation

As described in SampleClient GSSContext Instantiation, a `GSSContext` is created by instantiating a `GSSManager` and then calling one of its `createContext` methods.

Like `SampleClient`, `SampleServer` obtains an instance of the default `GSSManager` subclass by calling the `GSSManager` static method `getInstance`:

```
GSSManager manager = GSSManager.getInstance();
```

The `GSSManager` factory method for creating a context on the acceptor's side has the following signature:

```
GSSContext createContext(GSSCredential myCred);
```

If you pass `null` for the `GSSCredential` argument, as `SampleServer` does, the default credentials are used. The context is instantiated via the following:

```
GSSContext context = manager.createContext((GSSCredential)null);
```

SampleServer Context Establishment Loop

After SampleServer has instantiated a GSSContext, it can establish the security context with SampleClient. To do so, SampleServer has a loop that continues until the context is established. Each loop iteration does the following:

1. Receives a token from SampleClient. This token is the result of a SampleClient `initSecContext` call.
2. Calls the context's `acceptSecContext` method, passing it the token just received.
3. If `acceptSecContext` returns a token, then SampleServer sends this token to SampleClient and then starts the next loop iteration if the context is not yet established.

The tokens returned by `acceptSecContext` or received from SampleClient are placed in a byte array.

The `acceptSecContext` arguments are a byte array containing a token, the starting offset into that array of where the token begins, and the token length.

To exchange tokens with SampleClient, SampleServer uses the DataInputStream `inStream` and DataOutputStream `outStream` it previously set up using the input and output streams for the socket connection made with SampleClient.

Here is the SampleServer context establishment loop:

```
byte[] token = null;

while (!context.isEstablished()) {

    token = new byte[inStream.readInt()];
    System.out.println("Will read input token of size "
       + token.length
       + " for processing by acceptSecContext");
    inStream.readFully(token);

    token = context.acceptSecContext(token, 0, token.length);

    // Send a token to the peer if one was generated by
    // acceptSecContext
    if (token != null) {
        System.out.println("Will send token of size "
           + token.length
           + " from acceptSecContext.");
        outStream.writeInt(token.length);
        outStream.write(token);
        outStream.flush();
    }
}

System.out.print("Context Established! ");
System.out.println("Client is " + context.getSrcName());
System.out.println("Server is " + context.getTargName());
if (context.getMutualAuthState())
    System.out.println("Mutual authentication took place!");
```

Exchanging Messages Securely

Once a security context has been established between `SampleClient` and `SampleServer`, they can use the context to securely exchange messages.

- GSSContext Methods for Message Exchange
- The SampleClient and SampleServer Message Exchanges

GSSContext Methods for Message Exchange

Two types of methods exist for preparing messages for secure exchange: `wrap` and `getMIC`. There are actually two `wrap` methods (and two `getMIC` methods), where the differences between the two are the indication of where the input message is (a byte array or an input stream) and where the output should go (to a byte array return value or to an output stream).

These methods for preparing messages for exchange, and the corresponding methods for interpretation by the peer of the resulting tokens, are described below.

wrap

The `wrap` method is the primary method for message exchanges.

The signature for the `wrap` method called by `SampleClient` is the following:

```
byte[] wrap (byte[] inBuf, int offset, interface len,
                MessageProp msgProp)
```

You pass `wrap` a message (in `inBuf`), the offset into `inBuf` where the message begins (`offset`), and the length of the message (`len`). You also pass a `MessageProp`, which is used to indicate the desired QOP (Quality-of-Protection) and to specify whether or not privacy (encryption) is desired. A QOP value selects the cryptographic integrity and encryption (if requested) algorithm(s) to be used. The algorithms corresponding to various QOP values are specified by the provider of the underlying mechanism. For example, the values for Kerberos V5 are defined in RFC 1964 in section 4.2. It is common to specify 0 as the QOP value to request the default QOP.

The `wrap` method returns a token containing the message and a cryptographic Message Integrity Code (MIC) over it. The message placed in the token will be encrypted if the `MessageProp` indicates privacy is desired. You do not need to know the format of the returned token; it should be treated as opaque data. You send the returned token to your peer application, which calls the `unwrap` method to "unwrap" the token to get the original message and to verify its integrity.

getMIC

If you simply want to get a token containing a cryptographic Message Integrity Code (MIC) for a supplied message, you call `getMIC`. A sample reason you might want to do this is to confirm with your peer that you both have the same data, by just transporting a MIC for that data without incurring the cost of transporting the data itself to each other.

The signature for the `getMIC` method called by `SampleServer` is the following:

```
byte[] getMIC (byte[] inMsg, int offset, int len,
            MessageProp msgProp)
```

You pass `getMIC` a message (in `inMsg`), the offset into `inMsg` where the message begins (`offset`), and the length of the message (`len`). You also pass a `MessageProp`, which is used to indicate the desired QOP (Quality-of-Protection). It is common to specify 0 as the QOP value to request the default QOP.

If you have a token created by `getMIC` and the message used to calculate the MIC (or a message purported to be the message on which the MIC was calculated), you can call the `verifyMIC` method to verify the MIC for the message. If the verification is successful (that is, if a `GSSException` is not thrown), it proves that the message is exactly the same as it was when the MIC was calculated. A peer receiving a message from an application typically expects a MIC as well, so that they can verify the MIC and be assured the message has not been modified or corrupted in transit. Note: If you know ahead of time that you will want the MIC as well as the message then it is more convenient to use the `wrap` and `unwrap` methods. But there could be situations where the message and the MIC are received separately.

The signature for the `verifyMIC` corresponding to the `getMIC` shown above is the following:

```
void verifyMIC (byte[] inToken, int tokOffset, int tokLen,
        byte[] inMsg, int msgOffset, int msgLen,
        MessageProp msgProp);
```

This verifies the MIC contained in the `inToken` (of length `tokLen`, starting at offset `tokOffset`) over the message contained in `inMsg` (of length `msgLen`, starting at offset `msgOffset`). The `MessageProp` is used by the underlying mechanism to return information to the caller, such as the QOP indicating the strength of protection that was applied to the message.

The SampleClient and SampleServer Message Exchanges

The message exchanges between SampleClient and SampleServer are summarized below, followed by the coding details.

These steps are the "standard" steps used for verifying a GSS-API client and server. A group at MIT has written a GSS-API client and a GSS-API server that have become fairly popular test programs for checking interoperability between different implementations of the GSS-API library. (These GSS-API sample applications can be downloaded as a part of the Kerberos distribution available from MIT at http://web.mit.edu/kerberos.) This client and server from MIT follow the protocol that once the context is established, the client sends a message across and it expects back the MIC on that message. If you implement a GSS-API library, it is common practice to test it by running either the client or server using your library implementation against a corresponding peer server or client that uses another GSS-API library implementation. If both library implementations conform to the standards, then the two peers will be able to communicate successfully.

One implication of testing your client or server against ones written in C (like the MIT ones) is the way tokens must be exchanged. C implementations of GSS-API do not include stream-based methods. In the absence of stream-based methods on your peer, when you write a token you must first write the number of bytes and then write the token. Similarly, when you are reading a token, you first read the number of bytes and then read the token. This is what `SampleClient` and `SampleServer` do.

Here is the summary of the `SampleClient` and `SampleServer` message exchanges:

1. `SampleClient` calls `wrap` to encrypt and calculate a MIC for a message.
2. `SampleClient` sends the token returned from `wrap` to `SampleServer`.
3. `SampleServer` calls `unwrap` to obtain the original message and verify its integrity.
4. `SampleServer` calls `getMIC` to calculate a MIC on the decrypted message.
5. `SampleServer` sends the token returned by `getMIC` (which contains the MIC) to `SampleClient`.
6. `SampleClient` calls `verifyMIC` to verify that the MIC sent by `SampleServer` is a valid MIC for the original message.

SampleClient Code to Encrypt the Message and Send It

The `SampleClient` code for encrypting a message, calculating a MIC for it, and sending the result to `SampleServer` is the following:

```
byte[] messageBytes = "Hello There!\0".getBytes();

/*
 * The first MessageProp argument is 0 to request
 * the default Quality-of-Protection.
 * The second argument is true to request
 * privacy (encryption of the message).
 */
MessageProp prop =  new MessageProp(0, true);

/*
 * Encrypt the data and send it across. Integrity protection
 * is always applied, irrespective of encryption.
 */
token = context.wrap(messageBytes, 0, messageBytes.length,
    prop);
System.out.println("Will send wrap token of size "
    + token.length);
outStream.writeInt(token.length);
outStream.write(token);
outStream.flush();
```

SampleServer Code to Unwrap Token, Calculate MIC, and Send It

The following `SampleServer` code reads the wrapped token sent by `SampleClient` and "unwraps" it to obtain the original message and have its integrity verified. The unwrapping in this case includes decryption since the message was encrypted.

> **Note:**
>
> Here, the integrity check is expected to succeed. But note that in general if an integrity check fails, it signifies that the message was changed in transit. If the `unwrap` method encounters an integrity check failure, it throws a `GSSException` with major error code `GSSException.BAD_MIC`.

```
/*
 * Create a MessageProp which unwrap will use to return
 * information such as the Quality-of-Protection that was
 * applied to the wrapped token, whether or not it was
 * encrypted, etc. Since the initial MessageProp values
 * are ignored, it doesn't matter what they are set to.
 */
MessageProp prop = new MessageProp(0, false);

/*
 * Read the token. This uses the same token byte array
 * as that used during context establishment.
 */
token = new byte[inStream.readInt()];
```

```
System.out.println("Will read token of size "
    + token.length);
inStream.readFully(token);

byte[] bytes = context.unwrap(token, 0, token.length, prop);
String str = new String(bytes);
System.out.println("Received data \""
    + str + "\" of length " + str.length());
System.out.println("Encryption applied: "
    + prop.getPrivacy());
```

Next, `SampleServer` generates a MIC for the decrypted message and sends it to `SampleClient`. This is not really necessary but simply illustrates generating a MIC on the decrypted message, which should be exactly the same as the original message `SampleClient` wrapped and sent to `SampleServer`. When `SampleServer` generates this and sends it to `SampleClient`, and `SampleClient` verifies it, this proves to `SampleClient` that the decrypted message `SampleServer` has is in fact exactly the same as the original message from `SampleClient`.

```
/*
 * First reset the QOP of the MessageProp to 0
 * to ensure the default Quality-of-Protection
 * is applied.
 */
prop.setQOP(0);

token = context.getMIC(bytes, 0, bytes.length, prop);

System.out.println("Will send MIC token of size "
                   + token.length);
outStream.writeInt(token.length);
outStream.write(token);
outStream.flush();
```

SampleClient Code to Verify the MIC

The following `SampleClient` code reads the MIC calculated by `SampleServer` on the decrypted message and then verifies that the MIC is a MIC for the original message, proving that the decrypted message `SampleServer` has is the same as the original message:

```
token = new byte[inStream.readInt()];
System.out.println("Will read token of size " + token.length);
inStream.readFully(token);

/*
 * Recall messageBytes is the byte array containing
 * the original message and prop is the MessageProp
 * already instantiated by SampleClient.
 */
context.verifyMIC(token, 0, token.length,
          messageBytes, 0, messageBytes.length,
          prop);
```

```
System.out.println("Verified received MIC for message.");
```

Clean Up

When `SampleClient` and `SampleServer` have finished exchanging messages, they need to perform cleanup operations. Both contain the following code to

- close the socket connection and
- release system resources and cryptographic information stored in the context object and then invalidate the context.

```
socket.close();
context.dispose();
```

Kerberos User and Service Principal Names

Since the underlying authentication and secure communication technology used by this tutorial is Kerberos V5, we use Kerberos-style principal names wherever a user or service is called for (see Principals).

For example, when you run `SampleClient` you are asked to provide your **user name**. Your Kerberos-style user name is simply the user name you were assigned for Kerberos authentication. It consists of a base user name (like `mjones`) followed by an "`@`" and your realm (like `mjones@KRBNT-OPERATIONS.EXAMPLE.COM`).

A server program like `SampleServer` is typically considered to offer a "service" and to be run on behalf of a particular "**service principal**." A service principal name for `SampleServer` is needed in several places:

- When you run `SampleServer`, and `SampleClient` attempts a connection to it, the underlying Kerberos mechanism will attempt to authenticate to the Kerberos KDC. It prompts you to log in. You should log in as the appropriate service principal.
- When you run `SampleClient`, one of the arguments is the service principal name. This is needed so `SampleClient` can initiate establishment of a security context with the appropriate service.
- If the `SampleClient` and `SampleServer` programs were run with a security manager (they're not for this tutorial), the client and server policy files would each require a `ServicePermission` with name equal to the service principal name and action equal to "initiate" or "accept" (for initiating or accepting establishment of a security context).

Throughout this document, and in the accompanying login configuration file, `service_principal@your_realm`, is used as a placeholder to be replaced by the actual name to be used in your environment. *Any* Kerberos principal can actually be used for the service principal name. So **for the purposes of trying out this tutorial, you could use your user name as both the client user name and the service principal name.**

In a production environment, system administrators typically like servers to be run as specific principals only and may assign a particular name to be used. Often the Kerberos-style service principal name assigned is of the form

```
service_name/machine_name@realm;
```

For example, an nfs service run on a machine named `raven` in the realm named `KRBNT-OPERATIONS.EXAMPLE.COM` could have the service principal name

```
nfs/raven@KRBNT-OPERATIONS.EXAMPLE.COM
```

Such multi-component names are not required, however. Single-component names, just like those of user principals, can be used. For example, an installation might use the same ftp service principal `ftp@realm` for all ftp servers in that realm, while another installation might have different ftp principals for different ftp servers, such as `ftp/host1@realm` and `ftp/host2@realm` on machines `host1` and `host2`, respectively.

When the Realm Is Required in Principal Names

If the realm of a user or service principal name is the default realm (see Kerberos Requirements), you can leave off the realm when you are logging into Kerberos (that is, when you are prompted for your user name). Thus, for example, if your user name is `mjones@KRBNT-OPERATIONS.EXAMPLE.COM`, and you run `SampleClient`, when it requests your user name you could just specify `mjones`, leaving off the realm. The name is interpreted in the context of being a Kerberos principal name and the default realm is appended, as needed.

You can also leave off the realm if a principal name will be converted to a `GSSName` by a `GSSManager` `createName` method. For example, when you run `SampleClient`, one of the arguments is the server service principal name. You can specify the name without including the realm, because `SampleClient` passes the name to such a `createName` method, which appends the default realm as needed.

It is recommended that you always include realms when principal names are used in login configuration files and policy files, because the behavior of the parsers for such files may be implementation-dependent; they may or may not append the default realm before such names are utilized and subsequent actions may fail if there is no realm in the name.

The Login Configuration File

For this tutorial, we are letting the underlying Kerberos mechanism obtain credentials of the users running `SampleClient` and `SampleServer`, rather than invoking JAAS methods directly (as in the JAAS Authentication and JAAS Authorization tutorials) or indirectly (for example, via the Login utility described in the Use of JAAS Login Utility tutorial and in the Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges tutorial).

The default Kerberos mechanism implementation supplied by Oracle actually prompts for a Kerberos name and password and authenticates the specified user (or service) to the Kerberos KDC. The mechanism relies on JAAS to perform this authentication.

JAAS supports a pluggable authentication framework, meaning that any type of authentication module can be plugged under a calling application. A login configuration

specifies the login module to be used for a particular application. The default JAAS implementation from Oracle requires that the login configuration information be specified in a file. (Note: Some other vendors might not have file-based implementations.) See Appendix B: JAAS Login Configuration File for information as to what a login configuration file is, what it contains, and how to specify which login configuration file should be used.

For this tutorial, the Kerberos login module `com.sun.security.auth.module.Krb5LoginModule` is specified in the configuration file. This login module prompts for a Kerberos name and password and attempts to authenticate to the Kerberos KDC.

Both `SampleClient` and `SampleServer` can use the same login configuration file, if that file contains two entries, one entry for the client side and one for the server side.

The `bcsLogin.conf` login configuration file used for this tutorial is the following:

```
com.sun.security.jgss.initiate  {
  com.sun.security.auth.module.Krb5LoginModule required;
};

com.sun.security.jgss.accept  {
  com.sun.security.auth.module.Krb5LoginModule required storeKey=true
};
```

Entries with these two names (`com.sun.security.jgss.initiate` and `com.sun.security.jgss.accept`) are used by Oracle implementations of GSS-API mechanisms when they need new credentials. Since the mechanism used in this tutorial is the Kerberos V5 mechanism, a Kerberos login module will need to be invoked in order to obtain these credentials. Thus we list `Krb5LoginModule` as a required module in these entries. The `com.sun.security.jgss.initiate` entry specifies the configuration for the client side and the `com.sun.security.jgss.accept` entry for the server side.

The `Krb5LoginModule` succeeds only if the attempt to log in to the Kerberos KDC as a specified entity is successful. When running `SampleClient` or `SampleServer`, the user will be prompted for a name and password.

The `SampleServer` entry **`storeKey=true`** indicates that a secret key should be calculated from the password provided during login and it should be stored in the private credentials of the `Subject` created as a result of login. This key is subsequently utilized during mutual authentication when establishing a security context between `SampleClient` and `SampleServer`.

The `Krb5LoginModule` Javadoc API documentation describes the configuration options that the `Krb5LoginModule` class supports.

The useSubjectCredsOnly System Property

For this tutorial, we set the system property `javax.security.auth.useSubjectCredsOnly` to `false`, which allows us to relax the usual restriction of requiring a GSS mechanism to obtain necessary credentials from an existing Subject, set up by JAAS. When this restriction is relaxed, it allows the mechanism to obtain credentials from some vendor-specific location. For example, some vendors might choose to use the operating system's cache if one exists, while others might choose to read from a protected file on disk.

When this restriction is relaxed, Oracle's Kerberos mechanism still looks for the credentials in the Subject associated with the thread's access control context, but if it doesn't find any there, it performs JAAS authentication using a Kerberos module to obtain new ones. The Kerberos module prompts you for a Kerberos principal name and password. Note that Kerberos mechanism implementations from other vendors may behave differently when this property is set to `false`. Consult their documentation to determine their implementation's behavior.

Running the SampleClient and SampleServer Programs

To execute the `SampleClient` and `SampleServer` programs, do the following:

- Prepare SampleServer for Execution
- Prepare SampleClient for Execution
- Execute SampleServer
- Execute SampleClient

Prepare SampleServer for Execution

To prepare `SampleServer` for execution, do the following:

1. Copy the following files into a directory accessible by the machine on which you will run `SampleServer`:
   - The `SampleServer.java` source file.
   - The `bcsLogin.conf` login configuration file.
2. Compile `SampleServer.java`:

   ```
   javac SampleServer.java
   ```

Prepare SampleClient for Execution

To prepare `SampleClient` for execution, do the following:

1. Copy the following files into a directory accessible by the machine on which you will run `SampleClient`:
   - The `SampleClient.java` source file.
   - The `bcsLogin.conf` login configuration file.
2. Compile `SampleClient.java`:

   ```
   javac SampleClient.java
   ```

Execute SampleServer

It is important to execute `SampleServer` before `SampleClient` because `SampleClient` will try to make a socket connection to `SampleServer` and that will fail if `SampleServer` is not yet running and accepting socket connections.

To execute `SampleServer`, be sure to run it on the machine it is expected to be run on. This machine name (host name) is specified as an argument to `SampleClient`. The

service principal name appears in several places, including the login configuration file and the policy files.

Go to the directory in which you have prepared `SampleServer` for execution. Execute `SampleServer`, specifying

- by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified.

  For example, if your realm is `KRBNT-OPERATIONS.EXAMPLE.COM` you'd put `-Djava.security.krb5.realm=KRBNT-OPERATIONS.EXAMPLE.COM`.

- by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified.

  For example, if your KDC is `samplekdc.example.com`you'd put `-Djava.security.krb5.kdc=samplekdc.example.com`.

- by `-Djavax.security.auth.useSubjectCredsOnly=false` that the underlying mechanism can decide how to get credentials. See The useSubjectCredsOnly System Property.

- by `-Djava.security.auth.login.config=bcsLogin.conf` that the login configuration file to be used is `bcsLogin.conf`.

The only argument required by `SampleServer` is one specifying the port number to be used for listening for client connections. Choose a high port number unlikely to be used for anything else. An example would be something like 4444.

Below is the full command to use for both Microsoft Windows and Solaris, Linux, and macOS systems.

> **Note:**
>
> **Important: In this command, you must replace `<port_number>` with an appropriate port number, `<your_realm>` with your Kerberos realm, and `<your_kdc>` with your Kerberos KDC.**
>
> The `java.security.krb5.kdc` system property interprets the ":" symbol as a separation character for multiple KDCs. If the KDC is not listening on the default port (88), you must provide the default realm and its KDC(s) in a `krb5.conf` file, then set the system property `java.security.krb5.kdc.conf` with the name of this file:
>
> `-Djava.security.krb5.conf=<your_krb5.conf_file>`

Here is the command:

```
java -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djavax.security.auth.useSubjectCredsOnly=false
 -Djava.security.auth.login.config=bcsLogin.conf
 SampleServer <port_number>
```

The full command should appear on one line (or, on Solaris, Linux, or macOS, on multiple lines where each line but the last is terminated with " \" indicating that there is more to come). Multiple lines are used here just for legibility. Since this command is very long, you may need to place it in a .bat file (for Windows) or a .sh file (for Solaris, Linux, or macOS) and then run that file to execute the command.

The `SampleServer` code will listen for socket connections on the specified port. When prompted, type the Kerberos name and password for the service principal. The underlying Kerberos authentication mechanism specified in the login configuration file will log the service principal into Kerberos.

For login troubleshooting suggestions, see Troubleshooting.

Execute SampleClient

To execute `SampleClient`, first go to the directory in which you have prepared `SampleClient` for execution. Execute `SampleClient`, specifying

- by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified.
- by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified.
- by `-Djavax.security.auth.useSubjectCredsOnly=false` that the underlying mechanism can decide how to get credentials.
- by `-Djava.security.auth.login.config=bcsLogin.conf` that the login configuration file to be used is `bcsLogin.conf`.

The `SampleClient` arguments are (1) the Kerberos name of the service principal that represents `SampleServer` (see Kerberos User and Service Principal Names, (2) the name of the host (machine) on which `SampleServer` is running, and (3) the port number on which `SampleServer` is listening for client connections.

Below is the full command to use for both Windows and Solaris, Linux, and macOS systems.

> **Note:**
>
> **Important: In this command, you must replace *<service_principal>*, *<host>*, *<port_number>*, *<your_realm>*, and *<your_kdc>* with appropriate values** (and note that the port number must be the same as the port number passed as an argument to `SampleServer`). These values need not be placed in quotes.

Here is the command:

```
java -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djavax.security.auth.useSubjectCredsOnly=false
 -Djava.security.auth.login.config=bcsLogin.conf
 SampleClient <service_principal> <host> <port_number>
```

Type the full command on one line. Multiple lines are used here for legibility. As with the command for executing `SampleServer`, if the command is too long to type directly into your command window, place it in a .bat file (Windows) or a .sh file (Solaris, Linux, and macOS) and then execute that file.

When prompted, type your Kerberos user name and password. The underlying Kerberos authentication mechanism specified in the login configuration file will log you into Kerberos. The `SampleClient` code requests a socket connection with `SampleServer`. Once `SampleServer` accepts the connection, `SampleClient` and `SampleServer` establish a shared context and then exchange messages as described in this tutorial.

For login troubleshooting suggestions, see Troubleshooting.

JAAS Authentication

JAAS can be used for two purposes:

- for *authentication* of users, to reliably and securely determine who is currently executing Java code, regardless of whether the code is running as an application, an applet, a bean, or a servlet; and
- for *authorization* of users to ensure they have the access control rights (permissions) required to do the actions performed.

This section provides a basic tutorial for the authentication component. The authorization component will be described in the JAAS Authorization tutorial.

JAAS authentication is performed in a *pluggable* fashion. This permits Java applications to remain independent from underlying authentication technologies. New or updated technologies can be plugged in without requiring modifications to the application itself. An implementation for a particular authentication technology to be used is determined at runtime. The implementation is specified in a login configuration file. The authentication technology used for this tutorial is Kerberos. (See Kerberos Requirements.)

The rest of this tutorial consists of the following sections:

1. The Authentication Tutorial Code
2. The Login Configuration
3. Running the Code
4. Running the Code with a Security Manager

If you want to first see the tutorial code in action, you can skip directly to Running the Code and then go back to the other sections to learn about coding and configuration file details.

The Authentication Tutorial Code

Our authentication tutorial code is contained in a single source file, `JaasAcn.java`. This file's `main` method performs the authentication and then reports whether or not authentication succeeded.

The code for authenticating the user is very simple, consisting of just two steps:

1. Instantiating a LoginContext

2. Calling the LoginContext's login Method

Instantiating a LoginContext

In order to authenticate a user, you first need a `javax.security.auth.login.LoginContext`. Here is the basic way to instantiate a `LoginContext`:

```
import javax.security.auth.login.*;
. . .
LoginContext lc =
    new LoginContext(<config file entry name>,
           <CallbackHandler to be used for user interaction>);
```

and here is the specific way our tutorial code does the instantiation:

```
import javax.security.auth.login.*;
import com.sun.security.auth.callback.TextCallbackHandler;
. . .
LoginContext lc =
    new LoginContext("JaasSample",
          new TextCallbackHandler());
```

The arguments are the following:

1. **The name of an entry in the JAAS login configuration file**
   This is the name for the `LoginContext` to use to look up an entry for this application in the JAAS login configuration file, described in The Login Configuration. Such an entry specifies the class(es) that implement the desired underlying authentication technology(ies). The class(es) must implement the `LoginModule` interface, which is in the `javax.security.auth.spi` package.

   In our sample code, we use the `Krb5LoginModule` in the `com.sun.security.auth.module` package, which performs Kerberos authentication.

   The entry in the login configuration file we use for this tutorial (see `jaas.conf`) has the name "JaasSample", so that is the name we specify as the first argument to the `LoginContext` constructor.

2. **A `CallbackHandler` instance.**
   When a `LoginModule` needs to communicate with the user, for example to ask for a user name and password, it does not do so directly. That is because there are various ways of communicating with a user, and it is desirable for `LoginModules` to remain independent of the different types of user interaction. Rather, the `LoginModule` invokes a `CallbackHandler` to perform the user interaction and obtain the requested information, such as the user name and password. (`CallbackHandler` is an interface in the `javax.security.auth.callback` package.)

   An instance of the particular `CallbackHandler` to be used is specified as the second argument to the `LoginContext` constructor. The `LoginContext` forwards that instance to the underlying `LoginModule` (in our case `Krb5LoginModule`). An application typically provides its own `CallbackHandler` implementation. A simple `CallbackHandler`,

`TextCallbackHandler`, is provided in the `com.sun.security.auth.callback` package to output information to and read input from the command line.

Calling the LoginContext's login Method

Once we have a `LoginContext lc`, we can call its `login` method to carry out the authentication process:

```
lc.login();
```

The LoginContext instantiates a new empty `javax.security.auth.Subject` object (which represents the user or service being authenticated; see Subject). The `LoginContext` constructs the configured `LoginModule` (in our case `Krb5LoginModule`) and initializes it with this new `Subject` and `TextCallbackHandler`.

The `LoginContext`'s `login` method then calls methods in the `Krb5LoginModule` to perform the login and authentication. The `Krb5LoginModule` will utilize the `TextCallbackHandler` to obtain the user name and password. Then the `Krb5LoginModule` will use this information to get the user credentials from the Kerberos KDC. See the Kerberos reference documentation.

If authentication is successful, the `Krb5LoginModule` populates the `Subject` with (1) a Kerberos `Principal` representing the user and (2) the user's credentials (TGT).

The calling application can subsequently retrieve the authenticated `Subject` by calling the `LoginContext`'s `getSubject` method, although doing so is not necessary for this tutorial.

The Login Configuration

JAAS authentication is performed in a pluggable fashion, so applications can remain independent from underlying authentication technologies. A system administrator determines the authentication technologies, or `LoginModules`, to be used for each application and configures them in a login `Configuration`. The source of the configuration information (for example, a file or a database) is up to the current `javax.security.auth.login.Configuration` implementation. The default `Configuration` implementation from Oracle reads configuration information from configuration files, as described in `com.sun.security.auth.login.ConfigFile`.

See Appendix B: JAAS Login Configuration File for information as to what a login configuration file is, what it contains, and how to specify which login configuration file should be used.

The Login Configuration File for This Tutorial

As noted, the login configuration file we use for this tutorial, jass.conf, contains just one entry, which is

```
JaasSample {
  com.sun.security.auth.module.Krb5LoginModule required;
};
```

This entry is named `JaasSample` and that is the name that our tutorial application, `JaasAcn`, uses to refer to this entry. The entry specifies that the `LoginModule` to be used to do the user authentication is the `Krb5LoginModule` in the `com.sun.security.auth.module` package and that this `Krb5LoginModule` is required to "succeed" in order for authentication to be considered successful. The `Krb5LoginModule` succeeds only if the name and password supplied by the user are successfully used to log the user into the Kerberos KDC.

See the `Krb5LoginModule` Javadoc API documentation for information about all the possible options that can be passed to `Krb5LoginModule`.

Running the Code

To execute our JAAS authentication tutorial code, all you have to do is

1. Place the `JaasAcn.java` application source file and the `jaas.conf` login configuration file into a directory.
2. Compile `JaasAcn.java`:

   ```
   javac JaasAcn.java
   ```

3. Execute the `JaasAcn` application, specifying
   - by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified.

     For example, if your realm is `KRBNT-OPERATIONS.EXAMPLE.COM` you'd put `-Djava.security.krb5.realm=KRBNT-OPERATIONS.EXAMPLE.COM`.
   - by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified.

     For example, if your KDC is `samplekdc.example.com` you'd put `-Djava.security.krb5.kdc=samplekdc.example.com`.
   - by `-Djava.security.auth.login.config=jaas.conf` that the login configuration file to be used is `jaas.conf`.

The full command is below.

> **Note:**
>
> **Be sure to replace *<your_realm>* with your Kerberos realm, and *<your_kdc>* with your Kerberos KDC.**

```
java -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djava.security.auth.login.config=jaas.conf JaasAcn
```

Type all that on one line. Multiple lines are used here for legibility.

You will be prompted for your Kerberos user name and password, and the underlying Kerberos authentication mechanism specified in the login configuration file will log you into Kerberos. If your login is successful, you will see the following message:

```
Authentication succeeded!
```

If the login is not successful (for example, if you misspell your password), you will see

```
Authentication failed:
```

followed by a reason for the failure. For example, if you mistype your user name, you may see a message like the following (where the formatting is slightly modified here to increase legibility):

```
Authentication failed:
  Kerberos Authentication Failed:
    javax.security.auth.login.LoginException:
      KrbException: Client not found in Kerberos database
```

For login troubleshooting suggestions, see Troubleshooting.

After fixing any problems, re-run the program to try again.

Running the Code with a Security Manager

When a Java program is run with a security manager installed, the program is not allowed to access resources or otherwise perform security-sensitive operations unless it is explicitly granted permission (see Permissions in the JDK to do so by the security policy in effect. The permission must be granted by an entry in a policy file (see Default Policy Implementation and Policy File Syntax).

Most browsers install a security manager, so *applets* typically run under the scrutiny of a security manager. *Applications*, on the other hand, do not, since a security manager is not automatically installed when an application is running. Thus an application, like our `JaasAcn` application, by default has full access to resources.

To run an application with a security manager, simply invoke the interpreter with a `-Djava.security.manager` argument included on the command line.

If you try invoking `JaasAcn` with a security manager but without specifying any policy file, you will get the following (unless you have a default policy setup elsewhere that grants the required permissions or grants `AllPermission`):

```
% java -Djava.security.manager \
 -Djava.security.krb5.realm=<your_realm> \
 -Djava.security.krb5.kdc=<your_kdc> \
 -Djava.security.auth.login.config=jaas.conf JaasAcn
Exception in thread "main" java.security.AccessControlException:
  access denied (
  javax.security.auth.AuthPermission createLoginContext.JaasSample)
```

As you can see, you get an `AccessControlException`, because we haven't created and used a policy file granting our code the permission that is required in order to be allowed to create a `LoginContext`.

Here are the complete steps required in order to be able to run our `JaasAcn` application with a security manager installed. You can skip the first two steps if you have already done them, as described in Running the Code.

1. Place the `JaasAcn.java` application source file and the `jaas.conf` login configuration file into a directory.
2. Compile `JaasAcn.java`:

   ```
   javac JaasAcn.java
   ```

3. Create a JAR file containing `JaasAcn.class`:

   ```
   jar -cvf JaasAcn.jar JaasAcn.class
   ```

   This command creates a JAR file, `JaasAcn.jar`, and places the `JaasAcn.class` file inside it.

4. Create a policy file granting the code in the JAR file the required permission. The permission that is needed by code attempting to instantiate a `LoginContext` is a `javax.security.auth.AuthPermission` with target `createLoginContext.<entry name>`. Here, `<entry name>` refers to the name of the login configuration file entry that the application references in its instantiation of `LoginContext`. The name used by our `JaasAcn` application's `LoginContext` instantiation is `JaasSample`, as you can see in the code:

   ```
   LoginContext lc =
       new LoginContext("JaasSample",
             new TextCallbackHandler());
   ```

   Thus, the permission that needs to be granted to `JaasAcn.jar` is

   ```
   permission javax.security.auth.AuthPermission
     "createLoginContext.JaasSample";
   ```

   Copy the policy file `jaasacn.policy` to the same directory as that in which you stored `JaasAcn.java`, etc. This is a text file containing the following `grant` statement to grant `JaasAcn.jar` (in the current directory) the required permission:

   ```
   grant codebase "file:./JaasAcn.jar" {
      permission javax.security.auth.AuthPermission
                        "createLoginContext.JaasSample";
   };
   ```

   Note: Policy files and the structure of entries within them are described in Default Policy Implementation and Policy File Syntax. Permissions are described in Permissions in the JDK.

5. Execute the `JaasAcn` application, specifying

a. by an appropriate `-classpath` clause that classes should be searched for in the `JaasAcn.jar` JAR file,

b. by `-Djava.security.manager` that a security manager should be installed,

c. by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified. For example, if your realm is `KRBNT-OPERATIONS.EXAMPLE.COM`you'd put `-Djava.security.krb5.realm=KRBNT-OPERATIONS.EXAMPLE.COM`.

d. by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified. For example, if your KDC is `samplekdc.example.com`you'd put `-Djava.security.krb5.kdc=samplekdc.example.com`.

e. by `-Djava.security.policy=jaasacn.policy` that the policy file to be used is `jaasacn.policy`, and

f. by `-Djava.security.auth.login.config=jaas.conf` that the login configuration file to be used is `jaas.conf`.

The full command is below.

> **Note:**
>
> **Be sure to replace `<your_realm>` with your Kerberos realm, and `<your_kdc>` with your Kerberos KDC.**

```
java -classpath JaasAcn.jar -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djava.security.policy=jaasacn.policy
 -Djava.security.auth.login.config=jaas.conf JaasAcn
```

Type all that on one line. Multiple lines are used here for legibility. If the command is too long for your system, you may need to place it in a .bat file (for Windows) or a .sh file (for Solaris, Linux, and macOS) then run that file to execute the command.

Since the specified policy file contains an entry granting the code the required permission, `JaasAcn` will be allowed to instantiate a `LoginContext` and continue execution. You will be prompted for your Kerberos user name and password, and the underlying Kerberos authentication mechanism specified in the login configuration file will log you into Kerberos. If your login is successful, you will see the message "Authentication succeeded!" and if not, you will see "Authentication failed:" followed by a reason for the failure.

For login troubleshooting suggestions, see Troubleshooting.

JAAS Authorization

This tutorial expands the program and policy file developed in the JAAS Authentication tutorial to demonstrate the JAAS authorization component, which ensures the authenticated caller has the access control rights (permissions) required to do

subsequent security-sensitive operations. Since the authorization component requires that the user authentication first be completed, please read the JAAS Authentication tutorial first if you have not already done so.

The rest of this tutorial consists of the following sections:

- What is JAAS Authorization?
- How Is JAAS Authorization Performed?
  - How Do You Make Principal-Based Policy File Statements?
  - How Do You Associate a Subject with an Access Control Context?
- The Authorization Tutorial Code
- The Login Configuration File
- The Policy File
- Running the Authorization Tutorial Code

If you want to first see the tutorial code in action, you can skip directly to Running the Authorization Tutorial Code and then go back to the other sections to learn more.

What is JAAS Authorization?

JAAS authorization extends the existing Java security architecture that uses a security policy (see Default Policy Implementation and Policy File Syntax) to specify what access rights are granted to executing code. That architecture is *code-centric*. That is, the permissions are granted based on code characteristics: where the code is coming from and whether it is digitally signed and if so by whom. We saw an example of this in the `jaasacn.policy` file used in the JAAS Authentication tutorial. That file contains the following:

```
grant codebase "file:./JaasAcn.jar" {
   permission javax.security.auth.AuthPermission
                    "createLoginContext.JaasSample";
};
```

This grants the code in the `JaasAcn.jar` file, located in the current directory, the specified permission. (No signer is specified, so it doesn't matter whether the code is signed or not.)

JAAS authorization augments the existing code-centric access controls with new *user-centric* access controls. Permissions can be granted based not just on what code is running but also on *who* is running it.

When an application uses JAAS authentication to authenticate the user (or other entity such as a service), a `Subject` is created as a result. The purpose of the `Subject` is to represent the authenticated user. A `Subject` is comprised of a set of `Principal`s, where each `Principal` represents an identity for that user. For example, a `Subject` could have a name `Principal` ("Susan Smith") and a Social Security Number Principal ("987-65-4321"), thereby distinguishing this `Subject` from other Subjects.

Permissions can be granted in the policy to specific `Principals`. After the user has been authenticated, the application can associate the `Subject` with the current access control context. For each subsequent security-checked operation, (a local file access, for example), the Java runtime will automatically determine whether the policy

grants the required permission only to a specific `Principal` and if so, the operation will be allowed only if the `Subject` associated with the access control context contains the designated `Principal`.

How Is JAAS Authorization Performed?

To make JAAS authorization take place, the following is required:

- The user must be authenticated, as described in the JAAS Authentication tutorial.
- Principal-based entries must be configured in the security policy. See How Do You Make Principal-Based Policy File Statements?
- The Subject that is the result of authentication must be associated with the current access control context. See How Do You Associate a Subject with an Access Control Context?

How Do You Make Principal-Based Policy File Statements?

Policy file `grant` statements can now optionally include one or more Principal fields (see Default Policy Implementation and Policy File Syntax). Inclusion of a Principal field indicates that the user or other entity represented by the specified Principal, executing the specified code, has the designated permissions.

Thus, the basic format of a `grant` statement is now

```
grant <signer(s) field>, <codeBase URL>
  <Principal field(s)> {
    permission perm_class_name "target_name", "action";
    ....
    permission perm_class_name "target_name", "action";
  };
```

where each of the signer, codeBase and Principal fields is optional and the order between the fields doesn't matter.

A Principal field looks like the following:

```
Principal Principal_class "principal_name"
```

That is, it is the word "Principal" (where case doesn't matter) followed by the (fully qualified) name of a Principal class and a principal name.

A `Principal` class is a class that implements the `java.security.Principal` interface. All Principal objects have an associated name that can be obtained by calling their `getName` method. The format used for the name is dependent on each `Principal` implementation.

The type of `Principal` placed in the `Subject` created by the Kerberos authentication mechanism used by this tutorial is `javax.security.auth.kerberos.KerberosPrincipal`, so that is what should be used as the *`Principal_class`* part of our `grant` statement's `Principal` designation. User names for `KerberosPrincipals` are of the form *`name@realm`*. Thus, if the user name is `mjones` and the realm is `KRBNT-OPS.ABC.COM`, the full *`principal_name`* designation to use in the `grant` statement is `mjones@KRBNT-OPS.ABC.COM`.

It is possible to include more than one Principal field in a `grant` statement. If multiple Principal fields are specified, then the permissions in that `grant` statement are granted only if the Subject associated with the current access control context contains *all* of those Principals.

To grant the same set of permissions to different Principals, create multiple `grant` statements where each lists the permissions and contains a single Principal field designating one of the Principals.

The policy file for this tutorial includes one `grant` statement with a Principal field:

```
grant codebase "file:./SampleAction.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm"  {

   permission java.util.PropertyPermission "java.home", "read";
   permission java.util.PropertyPermission "user.home", "read";
   permission java.io.FilePermission "foo.txt", "read";
};
```

where you substitute your Kerberos user name (complete with "@" and realm) for `your_user_name@your_realm`. This specifies that the indicated permissions are granted to the specified principal executing the code in `SampleAction.jar`.

How Do You Associate a Subject with an Access Control Context?

To create and associate a Subject with an access control context, you need the following:

- The user must first be authenticated, as described in JAAS Authentication.
- The static `doAs` method from the `Subject` class must be called, passing it an authenticated `Subject` and a `java.security.PrivilegedAction` or `java.security.PrivilegedExceptionAction`. (See Appendix A: API for Privileged Blocks for a comparison of `PrivilegedAction` and `PrivilegedExceptionAction`.) The `doAs` method associates the provided `Subject` with the current access control context and then invokes the `run` method from the action. The `run` method implementation contains all the code to be executed as the specified `Subject`. The action thus executes as the specified `Subject`.
  The static `doAsPrivileged` method from the `Subject` class may be called instead of the `doAs` method. In addition to the parameters passed to `doAs`, `doAsPrivileged` requires a third parameter: an `AccessControlContext`. Unlike `doAs`, which associates the provided Subject with the current access control context, `doAsPrivileged` associates the `Subject` with the provided access control context. See doAs versus doAsPrivileged in the JAAS Reference Guide for a comparison of those methods.

The Authorization Tutorial Code

The code for this tutorial consists of two files:

- `JaasAzn.java` is exactly the same as the `JaasAcn.java` from the JAAS Authentication tutorial except for the additional code needed to call `Subject.doAsPrivileged`.

- `SampleAction.java` contains the `SampleAction` class. This class implements `PrivilegedAction` and has a `run` method that contains all the code we want to be executed with `Principal`-based authorization checks.

JaasAzn.java

`JaasAzn.java` is exactly the same as the `JaasAcn.java` code used in the previous tutorial except with three statements added at the end of the `main` method, after the authentication is done. These statements result in (1) association of a `Subject` representing the authenticated user with the current access control context and (2) execution of the code in the `run` method of `SampleAction`. Associating the `Subject` with the access control context enables security-sensitive operations in the `SampleAction` `run` method (and any code it invokes directly or indirectly) to be executed if a `Principal` representing the authenticated user is granted the required permissions in the current policy.

Like `JaasAcn.java`, `JaasAzn.java` instantiates a `LoginContext` `lc` and calls its `login` method to perform the authentication. If successful, the authenticated `Subject` (which includes a `Principal` representing the user) is obtained by calling the `LoginContext`'s `getSubject` method:

```
Subject mySubject = lc.getSubject();
```

The `main` method then calls `Subject.doAsPrivileged`, passing it the authenticated `Subject` `mySubject`, a `PrivilegedAction` (`SampleAction`) and a `null` `AccessControlContext`, as described in the following.

The `SampleAction` class is instantiated via the following:

```
PrivilegedAction action = new SampleAction();
```

The call to `Subject.doAsPrivileged` is performed via:

```
Subject.doAsPrivileged(mySubject, action, null);
```

The `doAsPrivileged` method invokes execution of the `run` method in the `PrivilegedAction` `action` (`SampleAction`) to initiate execution of the rest of the code, which is considered to be executed on behalf of the Subject `mySubject`.

Passing `null` as the `AccessControlContext` (third) argument to `doAsPrivileged` indicates that `mySubject` should be associated with a new empty `AccessControlContext`. The result is that security checks occurring during execution of `SampleAction` will only require permissions for the `SampleAction` code itself (or other code it invokes), running as `mySubject`. Note that the caller of `doAsPrivileged` (and the callers on the execution stack at the time `doAsPrivileged` was called) do not require any permissions while the action executes.

SampleAction.java

`SampleAction.java` contains the `SampleAction` class. This class implements `java.security.PrivilegedAction` and has a `run` method that contains all the code we want to be executed as the Subject `mySubject`. For this tutorial, we will perform

three operations, each of which cannot be done unless code has been granted required permissions. We will:

- Read and print the value of the `java.home` system property,
- Read and print the value of the `user.home` system property, and
- Determine whether or not a file named `foo.txt` exists in the current directory.

The Login Configuration File

The login configuration file used for this tutorial can be exactly the same as that used by the JAAS Authentication tutorial. Thus we can use `jaas.conf`, which contains just one entry:

```
JaasSample {
  com.sun.security.auth.module.Krb5LoginModule required;
};
```

This entry is named "JaasSample" and that is the name that both our tutorial applications `JaasAcn` and `JaasAzn` use to refer to it. The entry specifies that the `LoginModule` to be used to do the user authentication is the `Krb5LoginModule` in the `com.sun.security.auth.module` package and that this `Krb5LoginModule` is required to "succeed" in order for authentication to be considered successful. The `Krb5LoginModule` succeeds only if the name and password supplied by the user are successfully used to log the user into the Kerberos KDC.

The Policy File

This authorization tutorial contains two classes, `JaasAzn` and `SampleAction`. The code in each class contains some security-sensitive operations and thus relevant permissions are required in a policy file in order for the operations to be executed.

Permissions Required by JaasAzn

The main method of the `JaasAzn` class does two operations for which permissions are required. It

- creates a `LoginContext`, and
- calls the `doAsPrivileged` static method of the `Subject` class.

The `LoginContext` creation is exactly the same as was done in the authentication tutorial, and it thus needs the same `javax.security.auth.AuthPermission` permission with target `createLoginContext.JaasSample`.

In order to call the `doAsPrivileged` method of the Subject class, you need to have a `javax.security.auth.AuthPermission` with target `doAsPrivileged`.

Assuming the `JaasAzn` class is placed in a JAR file named `JaasAzn.jar`, these permissions can be granted to the `JaasAzn` code via the following `grant` statement in the policy file:

```
grant codebase "file:./JaasAzn.jar" {
   permission javax.security.auth.AuthPermission
                     "createLoginContext.JaasSample";
```

```
   permission javax.security.auth.AuthPermission "doAsPrivileged";
};
```

Permissions Required by SampleAction

The `SampleAction` code does three operations for which permissions are required. It

- reads the value of the `java.home` system property.
- reads the value of the `user.home` system property.
- checks to see whether or not a file named `foo.txt` exists in the current directory.

The permissions required for these operations are the following:

```
permission java.util.PropertyPermission "java.home", "read";
permission java.util.PropertyPermission "user.home", "read";
permission java.io.FilePermission "foo.txt", "read";
```

We need to grant these permissions to the code in `SampleAction.class`, which we will place in a JAR file named `SampleAction.jar`. However, for this particular `grant` statement we want to grant the permissions not just to the *code* but to a specific user executing the code, to demonstrate how to restrict access to a particular user.

Thus, as explained in How Do You Make Principal-Based Policy File Statements?, our `grant` statement looks like the following:

```
grant codebase "file:./SampleAction.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm"  {

   permission java.util.PropertyPermission "java.home", "read";
   permission java.util.PropertyPermission "user.home", "read";
   permission java.io.FilePermission "foo.txt", "read";
};
```

You substitute your Kerberos user name (complete with "@" and realm) for `your_user_name@your_realm`. For example, if your user name is `mjones` and your realm is `KRBNT-OPERATIONS.ABC.COM`, you would use `mjones@KRBNT-OPERATIONS.ABC.COM`.

Running the Authorization Tutorial Code

To execute our JAAS authorization tutorial code, all you have to do is

1. Place the following files into a directory:
   - The `JaasAzn.java` source file.
   - The `SampleAction.java` source file.
   - The `jaas.conf` login configuration file.
   - The `jaasazn.policy` policy file.
2. Replace `your_user_name@your_realm` in `jaasazn.policy` with your user name and realm.

3. Compile `SampleAction.java` and `JaasAzn.java`:

```
javac SampleAction.java JaasAzn.java
```

4. Create a JAR file named `JaasAzn.jar` containing `JaasAzn.class`:

```
jar -cvf JaasAzn.jar JaasAzn.class
```

5. Create a JAR file named `SampleAction.jar` containing `SampleAction.class`:

```
jar -cvf SampleAction.jar SampleAction.class
```

6. Execute the `JaasAzn` application, specifying
   a. by an appropriate `-classpath` clause that classes should be searched for in the `JaasAzn.jar` and `SampleAction.jar` JAR files,
   b. by `-Djava.security.manager` that a security manager should be installed,
   c. by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified.
   d. by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified.
   e. by `-Djava.security.policy=jaasazn.policy` that the policy file to be used is `jaasazn.policy`, and
   f. by `-Djava.security.auth.login.config=jaas.conf` that the login configuration file to be used is `jaas.conf`.

   Below are the full commands to use for Windows, Solaris, Linux, and macOS. The only difference is that on Windows you use semicolons to separate classpath items, while you use colons for that purpose on Solaris, Linux, and macOS.

   > **Note:**
   >
   > **Be sure to replace *<your_realm>* with your Kerberos realm, and *<your_kdc>* with your Kerberos KDC.**

   Here is the full command for Windows:

```
java -classpath JaasAzn.jar;SampleAction.jar
 -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djava.security.policy=jaasazn.policy
 -Djava.security.auth.login.config=jaas.conf JaasAzn
```

   Here is the full command for Solaris, Linux, and macOS:

```
java -classpath JaasAzn.jar:SampleAction.jar
 -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
```

```
-Djava.security.krb5.kdc=<your_kdc>
-Djava.security.policy=jaasazn.policy
-Djava.security.auth.login.config=jaas.conf JaasAzn
```

Type the full command on one line. Multiple lines are used here for legibility. If the command is too long for your system, you may need to place it in a .bat file (for Windows) or a .sh file (for Solaris, Linux, and macOS) and then run that file to execute the command.

You will be prompted for your Kerberos user name and password, and the underlying Kerberos authentication mechanism specified in the login configuration file will log you into Kerberos. If your login is successful, you will see the message "Authentication succeeded!" and if not, you will see "Authentication Failed."

For login troubleshooting suggestions, see Troubleshooting.

Once authentication is successfully completed, the rest of the program (in `SampleAction`) will be executed on behalf of you, the user, requiring you to have been granted appropriate permissions. The `jaasazn.policy` policy file grants you the required permissions, so you will see a display of the values of your `java.home` and `user.home` system properties and a statement as to whether or not you have a file named `foo.txt` in the current directory.

Use of JAAS Login Utility

The previous two tutorials, JAAS Authentication and JAAS Authorization, show how you can use the `LoginContext` and `Subject` classes to write a program to

- authenticate the user to verify his or her identity and
- associate an instance of a particular class representing the user (a `Subject` instance) with an access control context in such a way that subsequent security-sensitive operations will be allowed if the current policy grants the user the required permissions.

This tutorial describes a Login utility that performs the above operations and then executes any specified application as the authenticated user.

Use of the Login utility with a sample application is demonstrated in this tutorial. The next tutorial, Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges, a client/server application using the Java GSS-API, also uses the Login utility.

It is not necessary to read the previous two tutorials on JAAS authentication and authorization prior to reading this one. However, you may want to refer to some sections in those tutorials to obtain further details regarding certain topics, such as JAAS Authorization. You should also read Appendix B: JAAS Login Configuration File for information as to what a login configuration file is, since one is needed for this and all other tutorials in this series.

As with all tutorials in this series of tutorials, the underlying technology used to support authentication is Kerberos. See Kerberos Requirements.

- What You Need to Know About the Login Utility
- Application and Other File Requirements

- The Sample Application Program
- The Login Configuration File
- The Policy File
- Running the Sample Program with the Login Utility

If you want to first see the tutorial code in action, you can skip directly to Running the Sample Program with the Login Utility and then go back to the other sections to learn more.

What You Need to Know About the Login Utility

You do not need to understand the code contained in `Login.java`; you can just use it as is. However, you need to understand some facts about what it does so that your program, policy file, and login configuration file will properly work with it. Below is a summary of these facts, followed by sections with further information and examples.

The `Login` class does the following:

- Assumes it is passed, as arguments, your application's top-level class name, followed by any arguments your application may require.
- Assumes that the class name of your top-level application class is also used as the name of the entry to be looked up in your login configuration file.
- Specifies the `TextCallbackHandler` class (from the `com.sun.security.auth.callback` package) as the class to be used when communicating with the user. This class can prompt the user for a user name and password.
- Uses a `LoginContext` to authenticate the user. The `LoginContext` invokes the appropriate authentication technology, or `LoginModule`, to perform the authentication. `LoginModules` use a `CallbackHandler` (in our case, `TextCallbackHandler`) as needed to communicate with the user.
- Allows the user three attempts to successfully log in.
- Creates an instance of the `MyAction` class (also in `Login.java`), passing it the application arguments, if any.
- Invokes `Subject.doAsPrivileged`, passing it a `Subject` representing the user, the `MyAction` instance, and a null `AccessControlContext`. The result is that the public static main method from your application is invoked and your application code is considered to be executed on behalf of the user.

Application and Other File Requirements

To utilize the Login utility to authenticate the user and execute your application, you may need a small number of additions or modifications to your login configuration file and policy file, as described in the following.

- Application Requirements
- Login Configuration File Requirements
- Policy File Requirements

Application Requirements

In order to utilize the Login utility, your application code does not need anything special. All you need is for the entry point of your application to be the `main` method of a class you write, as usual.

The way to invoke Login such that it will authenticate the user and then instantiate `MyAction` to invoke your application is the following:

```
java <options> Login <AppName> <app arguments>
```

where *<AppName>* is your application's top-level class name and *<app arguments>* are any arguments required by your application. See Running the Sample Program with the Login Utility for the full command used for this tutorial.

Login Configuration File Requirements

Whenever a LoginContext is used to authenticate the user, you need a login configuration file to specify the desired login module. See the The Login Configuration section in the JAAS authentication tutorial for more information as to what a login configuration file is and what it contains.

When you use the Login utility, the name for the login configuration file entry must be exactly the same as your top-level application class name. See The Login Configuration File in this tutorial for an example.

Policy File Requirements

Whenever you run an application with a security manager, you need a policy indicating the permissions granted to specific code, or to specific code being executed by a specific user (or users). One way of specifying the policy is by `grant` statements in a policy file. See The Policy File for more information.

If you use the Login utility to invoke your application, then you will need to grant it various permissions, as described in Permissions Required by the Login and MyAction Classes.

The Sample Application Program

The `Sample.java` application used for this tutorial performs the same actions as the `SampleAction.java` application did in the previous (JAAS Authorization) tutorial. It does the following:

- Reads and prints the value of the `java.home` system property,
- Reads and prints the value of the `user.home` system property, and
- Determines whether or not a file named `foo.txt` exists in the current directory.

The Login Configuration File

The `sample.conf` login configuration file for this tutorial contains a single entry, just like the login configuration file for the previous (JAAS Authorization) tutorial. The entry contents are the same since the class implementing the desired authentication

technology in both cases is the `Krb5LoginModule` in the `com.sun.security.auth.module` package.

The only difference is the name used for the entry. In the previous tutorial we used the name "JaasSample", since that is the name used by the `JaasAzn` class to look up the entry. When you use the Login utility with your application, it expects the name for your login configuration file entry to be the same as the name of your top-level application class. That application class for this tutorial is named "Sample" so that must also be the name of the login configuration file entry. Thus the login configuration file looks like the following:

```
Sample {
   com.sun.security.auth.module.Krb5LoginModule required;
};
```

The "required" indicates that login using the `Krb5LoginModule` is required to "succeed" in order for authentication to be considered successful. The `Krb5LoginModule` succeeds only if the name and password supplied by the user are successfully used to log the user into the Kerberos KDC.

See the `Krb5LoginModule` Javadoc API documentation for information about all the possible options that can be passed to `Krb5LoginModule`.

The Policy File

The `Login`, `MyAction`, and `Sample` classes all perform some security-sensitive operations and thus relevant permissions are required in a policy file, `sample.policy`, in order for the operations to be executed.

Permissions Required by the Login and MyAction Classes

For this tutorial, you will create a `Login.jar` JAR file containing the `Login.class` and `MyAction.class` files. You need to grant `Login.jar` various permissions, specifically the ones required for invoking the security-sensitive methods the `Login.jar` classes call, as well as all the permissions required by your application. Otherwise, access control checks will fail.

The simplest thing to do, and what we recommend, is to grant `Login.jar` `AllPermission`. For this tutorial, the `Login.jar` file is assumed to be in the current directory and the policy file includes the following:

```
grant codebase "file:./Login.jar" {
   permission java.security.AllPermission;
};
```

Permissions Required by Sample

(Note: This section is essentially a modified copy of the Permissions Required by SampleAction section from the previous (JAAS Authorization) tutorial, since `Sample` and `SampleAction` perform the same operations and thus require the same permissions.)

The `Sample` code does three operations for which permissions are required. It

- reads the value of the `java.home` system property.
- reads the value of the `user.home` system property.
- checks to see whether or not a file named `foo.txt` exists in the current directory.

The permissions required for these operations are the following:

```
permission java.util.PropertyPermission "java.home", "read";
permission java.util.PropertyPermission "user.home", "read";
permission java.io.FilePermission "foo.txt", "read";
```

We need to grant these permissions to the code in `Sample.class`, which we will place in a JAR file named `Sample.jar`. However, our `grant` statement will grant the permissions not just to the *code* but to a specific authenticated user executing the code. This illustrates how you can use a Principal designation in a `grant` statement to restrict execution of security-sensitive operations in code to a specific user rather than allowing the permissions to all users executing the code.

Thus, as explained in JAAS Authorization, our `grant` statement looks like the following:

```
grant codebase "file:./Sample.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm"  {

   permission java.util.PropertyPermission "java.home", "read";
   permission java.util.PropertyPermission "user.home", "read";
   permission java.io.FilePermission "foo.txt", "read";
};
```

> **Note:**
>
> **Important: You must substitute your Kerberos user name (complete with "@" and realm) for `your_user_name@your_realm`.**
>
> For example, if your user name is `mjones` and your realm is `KRBNT-OPERATIONS.ABC.COM`, you would use `mjones@KRBNT-OPERATIONS.ABC.COM`.

Running the Sample Program with the Login Utility

To execute the `Sample` application with the Login utility, do the following:

1. Place the following files into a directory:
   - The `Login.java` source file.
   - The `Sample.java` source file.
   - The `sample.conf` login configuration file.
   - The `sample.policy` policy file.
2. Replace `your_user_name@your_realm` in `sample.policy` with your user name and realm.

3. Compile `Login.java` and `Sample.java`:

```
javac Login.java Sample.java
```

   Note that `Login.java` contains two classes and thus compiling `Login.java` creates `Login.class` and `MyAction.class`.

4. Create a JAR file named `Login.jar` containing `Login.class and MyAction.class`:

```
jar -cvf Login.jar Login.class MyAction.class
```

5. Create a JAR file named `Sample.jar` containing `Sample.class`:

```
jar -cvf Sample.jar Sample.class
```

6. Execute the `Login` class, specifying

   - by an appropriate `-classpath` clause that classes should be searched for in the `Login.jar` and `Sample.jar` JAR files,
   - by `-Djava.security.manager` that a security manager should be installed,
   - by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified.
   - by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified.
   - by `-Djava.security.policy=sample.policy` that the policy file to be used is `sample.policy`, and
   - by `-Djava.security.auth.login.config=sample.conf` that the login configuration file to be used is `sample.conf`.

   You pass the name of your application (in this case, `Sample`) as an argument to `Login`. You would then add as arguments any arguments required by your application, but in our case `Sample` does not require any.

   Below are the full commands to use for Windows, Solaris, Linux, and macOS. The only difference is that on Windows you use semicolons to separate classpath items, while you use colons for that purpose on Solaris, Linux, and macOS. **Be sure to replace `<your_realm>` with your Kerberos realm, and `<your_kdc>` with your Kerberos KDC.**

   Here is the full command for Windows:

```
java -classpath Login.jar;Sample.jar
 -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djava.security.policy=sample.policy
 -Djava.security.auth.login.config=sample.conf Login Sample
```

   Here is the full command for Solaris, Linux, and macOS:

```
java -classpath Login.jar:Sample.jar
 -Djava.security.manager
```

```
-Djava.security.krb5.realm=<your_realm>
-Djava.security.krb5.kdc=<your_kdc>
-Djava.security.policy=sample.policy
-Djava.security.auth.login.config=sample.conf Login Sample
```

Type the full command on one line. Multiple lines are used here for legibility. If the command is too long for your system, you may need to place it in a .bat file (for Windows) or a .sh file (for Solaris, Linux, and macOS) and then run that file to execute the command.

You will be prompted for your Kerberos user name and password, and the underlying Kerberos login module specified in the login configuration file will log you into Kerberos. Once authentication is successfully completed, the `Sample` code will be executed on behalf of you, the user. The `sample.policy` policy file grants you the required permissions, so you will see a display of the values of your `java.home` and `user.home` system properties and a statement as to whether or not you have a file named `foo.txt` in the current directory.

For login troubleshooting suggestions, see Troubleshooting.

Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges

This tutorial presents two sample applications to demonstrate the use of the Java GSS-API. This API permits secure exchanges of messages between communicating applications. Here are the sample client and server applications you'll need for this tutorial:

- `SampleClient.java`
- `SampleServer.java`

> **Note:**
>
> This tutorial uses the same client and server applications as the Use of Java GSS-API for Secure Message Exchanges Without JAAS Programming tutorial. In that tutorial, JAAS (Java Authentication and Authorization Service) programming is not required. Instead, you let the underlying mechanism decide how to get credentials.

This tutorial uses policy files and a more complex login configuration file. The programs are run with a security manager; as a result, security-sensitive operations are not allowed unless the required permissions were explicitly granted. This tutorial also demonstrates how JAAS authorization adds *user-centric* access control that applies control based on *who* is running the code – not just on what code is running.

- Before You Start: Recommended Reading
- Overview of the Client and Server Applications
- Kerberos User and Service Principal Names
- The Login Configuration File
- The Policy Files

- Running the SampleClient and SampleServer Programs

Before You Start: Recommended Reading

In this Java GSS-API tutorial, the first step is JAAS authentication. Previous tutorials demonstrated the use of JAAS for user authentication and authorization, and presented examples of policy files and of login configuration files (specifying the underlying authentication technology to be used) that JAAS requires. Applications in the JAAS introductory tutorials, JAAS Authentication and JAAS Authorization, made direct calls to JAAS methods. The Use of JAAS Login Utility tutorial showed the use of a utility program that frees the application from having to do this. The client and server applications in the current tutorial also use the same utility program, so we recommend you **read the login utility tutorial first**.

As with all tutorials in this series, the underlying technology used to support authentication and secure communication for the applications in this tutorial is Kerberos. See Kerberos Requirements.

Overview of the Client and Server Applications

The applications for this tutorial are named `SampleClient` and `SampleServer`.

Each is invoked by executing the Login utility supplied with this tutorial and passing it as arguments the name of the application (`SampleClient` or `SampleServer`), followed by the arguments needed by the application. The Login utility uses a JAAS `LoginContext` to authenticate the user using Kerberos. Finally, the Login utility invokes the `main` method of the application class, in our case either `SampleClient` or `SampleServer`, and passes the application its arguments.

Here is a summary of execution of the `SampleClient` and `SampleServer` applications:

1. Run the `SampleServer` application by running the Login utility and passing it as arguments the name "`SampleServer`" followed by the arguments for the `SampleServer` program. The Login utility prompts you for the password for the principal that `SampleServer` should run as. (See Kerberos User and Service Principal Names.) After authentication is complete, `SampleServer` is run it:
   a. Reads its argument, the port number that it should listen on for client connections.
   b. Creates a `ServerSocket` for listening for client connections on that port.
   c. Listens for a connection.
2. Run the `SampleClient` application (possibly on a different machine), by running the Login utility and passing it as arguments the name "`SampleClient`" followed by the arguments for the `SampleClient` program. The Login utility prompts you for your Kerberos name and password. After authentication is complete, `SampleClient` is run. It
   a. Reads its arguments: (1) The name of the Kerberos principal that represents `SampleServer`. (See Kerberos User and Service Principal Names.), (2) the name of the host (machine) on which `SampleServer` is running, and (3) the port number on which `SampleServer` listens for client connections.
   b. Attempts a socket connection with the `SampleServer`, using the host and port it was passed as arguments.

3. The socket connection is accepted by `SampleServer` and both applications initialize a `DataInputStream` and a `DataOutputStream` from the socket input and output streams, to be used for future data exchanges.
4. `SampleClient` and `SampleServer` each instantiate a `GSSContext` and establish a shared context that will enable subsequent secure data exchanges.
5. `SampleClient` and `SampleServer` can now securely exchange messages.
6. When `SampleClient` and `SampleServer` are done exchanging messages, they perform clean-up operations.

> **Note:**

Refer to the The SampleClient and SampleServer Code section of the Use of Java GSS-API for Secure Message Exchanges Without JAAS Programming tutorial for a full discussion of the code used in this tutorial.

Kerberos User and Service Principal Names

Since the underlying authentication and secure communication technology used by this tutorial is Kerberos V5, we use Kerberos-style principal names wherever a user or service is called for (see Principals).

For example, when you run `SampleClient` you are asked to provide your **user name**. Your Kerberos-style user name is simply the user name you were assigned for Kerberos authentication. It consists of a base user name (like `mjones`) followed by an "`@`" and your realm (like `mjones@KRBNT-OPERATIONS.EXAMPLE.COM`).

A server program like `SampleServer` is typically considered to offer a "service" and to be run on behalf of a particular "**service principal**." A service principal name for `SampleServer` is needed in several places:

- When you run `SampleServer` you must log in as the appropriate service principal. The login configuration file for this tutorial actually specifies the service principal name (as an option to the `Krb5LoginModule`), so the JAAS authentication (done by the `Login` utility) just asks you to specify the password for that service principal. If you specify the correct password, the authentication is successful, a `Subject` is created containing a `Principal` with the service principal name, and that `Subject` is associated with a new access control context. The subsequently-executed code (the `SampleServer` code) is considered to be executed on behalf of the specified principal.
- When you run `SampleClient`, one of the arguments is the service principal name. This is needed so `SampleClient` can initiate establishment of a security context with the appropriate service.
- The client and server policy files each require a `ServicePermission` with name equal to the service principal name and action equal to "initiate" or "accept" (for initiating or accepting establishment of a security context).

Throughout this document, and in the accompanying login configuration file and policy files, `service_principal@your_realm` is used as a placeholder to be replaced by the actual name to be used in your environment. *Any* Kerberos principal can actually be used for the service principal name. So **for the purposes of trying out this tutorial,**

you could use your user name as both the client user name and the service principal name.

In a production environment, system administrators typically like servers to be run as specific principals only and may assign a particular name to be used. Often the Kerberos-style service principal name assigned is of the form

```
service_name/machine_name@realm;
```

For example, an nfs service run on a machine named "raven" in the realm named `KRBNT-OPERATIONS.EXAMPLE.COM` could have the service principal name

```
nfs/raven@KRBNT-OPERATIONS.EXAMPLE.COM
```

Such multi-component names are not required, however. Single-component names, just like those of user principals, can be used. For example, an installation might use the same ftp service principal `ftp@realm` for all ftp servers in that realm, while another installation might have different ftp principals for different ftp servers, such as `ftp/host1@realm` and `ftp/host2@realm` on machines `host1` and `host2`, respectively.

When the Realm is Required in Principal Names

If the realm of a user or service principal name is the default realm (see Kerberos Requirements), you can leave off the realm when you are logging into Kerberos (that is, when you are prompted for your user name). Thus, for example, if your user name is `mjones@KRBNT-OPERATIONS.EXAMPLE.COM`, and you run `SampleClient`, when it requests your user name you could just specify `mjones`, leaving off the realm. The name is interpreted in the context of being a Kerberos principal name and the default realm is appended, as needed.

You can also leave off the realm if a principal name will be converted to a `GSSName` by a `GSSManager` `createName` method. For example, when you run `SampleClient`, one of the arguments is the server service principal name. You can specify the name without including the realm, because `SampleClient` passes the name to such a `createName` method, which appends the default realm as needed.

It is recommended that you always include realms when principal names are used in login configuration files and policy files, because the behavior of the parsers for such files may be implementation-dependent; they may or may not append the default realm before such names are utilized and subsequent actions may fail if there is no realm in the name.

The Login Configuration File

Whenever JAAS is used, a login configuration is required to specify the desired authentication technology. (See Appendix B: JAAS Login Configuration File for more information about what a login configuration file is.) Both `SampleClient` and `SampleServer` can use the same login configuration file, if that file contains two entries, one entry for the client side and one for the server side.

The `csLogin.conf` login configuration file used for this tutorial is the following:

```
SampleClient {
  com.sun.security.auth.module.Krb5LoginModule required;
```

```
};

SampleServer {
  com.sun.security.auth.module.Krb5LoginModule required storeKey=true
    principal="service_principal@your_realm";
};
```

Note that the name for each entry matches the respective class names for our two top-level applications, `SampleClient` and `SampleServer`. Recall that this is also the name that is passed to the Login utility that performs JAAS operations for the application. That utility expects the name of the entry to be looked up in your login configuration file to be the same as the name it is passed.

Both entries specify that Oracle's Kerberos V5 LoginModule must be used to successfully authenticate the user. The Krb5LoginModule succeeds only if the attempt to log in to the Kerberos KDC as a specified entity is successful. In the case of `SampleClient`, the user will be prompted for their name and password. In the case of `SampleServer`, a name is already supplied in this login configuration file (the specified principal, as described below) and the user running `SampleServer` is just asked for the password for the entity specified by that name. They must specify the correct password in order for authentication to succeed.

The `SampleServer` entry **`storeKey=true`** indicates that a secret key should be calculated from the password provided during login and it should be stored in the private credentials of the `Subject` created as a result of login. This key is subsequently utilized during mutual authentication when establishing a security context between `SampleClient` and `SampleServer`.

The `Krb5LoginModule` has a **`principal`** option that can be used to specify that only the specified principal (entity/user) should be logged in for the given program. Here, the `SampleClient` entry does not specify a principal (although it could, if desired), so the user is prompted for a user name and password and anyone with a valid user name and password can run `SampleClient`. `SampleServer`, on the other hand, indicates a particular principal because system administrators usually like servers to be run as specific principals only. In this case, the user running `SampleServer` is prompted for that principal's password and must supply the correct one in order for authentication to succeed.

Note that you must replace `service_principal@your_realm` with the name of the service principal that represents `SampleServer`. (See Kerberos User and Service Principal Names.)

If the server has a keytab file containing secret keys, then use the following JAAS login entry:

```
SampleServer {
  com.sun.security.auth.module.Krb5LoginModule required
  principal="service_principal@your_realm"
  storeKey=true useKeyTab=true keyTab=keytab.file.name
  isInitiator=false;
  };
```

Because the keytab file already provides the keys, you will not be prompted for a password. If the keytab file contains keys for more than one service principal and the

server is designed to act as all these service principals, then you can set the principal entry to the following:

```
principal=*
```

See the `Krb5LoginModule` Javadoc API documentation for information about all the possible options that can be passed to `Krb5LoginModule`.

The Policy Files

The policy file used when running `SampleClient` is `client.policy`, and the policy file used when running `SampleServer` is `server.policy`. Their contents are described below.

The Client Policy File

Permissions Required by the Login Utility Classes

A number of permissions are required by the classes in `Login.java` (`Login` and `MyAction`). As recommended in Use of JAAS Login Utility on the use of the Login utility, we create a `Login.jar` JAR file containing the `Login.class` and `MyAction.class` files and in the `client.policy` policy file we grant `Login.jar` `AllPermission`:

```
grant codebase "file:./Login.jar" {
   permission java.security.AllPermission;
};
```

Permissions Required by SampleClient

The `SampleClient` code does two types of operations for which permissions are required. It

- opens a socket connection with the host machine running the `SampleServer` application.
- initiates establishment of a security context with `SampleServer`.

The permission required to open a socket connection is

```
permission java.net.SocketPermission "*", "connect";
```

You may replace the "*" with the hostname or IP address of the machine that `SampleServer` will be running on.

The permission(s) required to initiate establishment of a security context will depend on the underlying mechanism. This tutorial uses Kerberos as the underlying mechanism, and for that two `javax.security.auth.kerberos.ServicePermissions` are required. A `ServicePermission` contains a service principal name and an action (or list of actions). To initiate establishment of a security context, you need two `ServicePermissions` with action "initiate", whose names specify:

- the service principal name for the ticket granting service for your realm. Granting this permission essentially allows the use of Kerberos as a client.

- the service principal name representing `SampleServer`. (See Kerberos User and Service Principal Names.) Granting this permission allows you to interact with the service, `SampleServer`, using Kerberos.

We want to grant the permissions to a specific authenticated user executing `SampleClient`, so we specify both the `SampleClient` code location (in `SampleClient.jar`) and a Principal designation indicating the user name and realm for the user (you, the person who will run `SampleClient`). (See How Do You Make Principal-Based Policy File Statements? in JAAS Authorization for information on policy file `grant` statements that include `Principal` designations.)

Here is the basic form for the `grant` statement:

```
  grant CodeBase "file:./SampleClient.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm" {

    permission java.net.SocketPermission "*", "connect";

    permission javax.security.auth.kerberos.ServicePermission
        "krbtgt/your_realm@your_realm",
        "initiate";

    permission javax.security.auth.kerberos.ServicePermission
        "service_principal@your_realm",
        "initiate";
};
```

You must substitute your Kerberos user name (complete with "@" and realm) for `your_user_name@your_realm`. For example, if your user name is `mjones` and your realm is `KRBNT-OPERATIONS.EXAMPLE.COM`, you would use `mjones@KRBNT-OPERATIONS.EXAMPLE.COM`.

You must also substitute your realm in `krbtgt/your_realm@your_realm` and the service principal name for the service principal representing the server (see Kerberos User and Service Principal Names for the service principal name for the service principal representing the server for `service_principal@your_realm`. Suppose the former is `krbtgt/KRBNT-OPERATIONS.EXAMPLE.COM@KRBNT-OPERATIONS.EXAMPLE.COM` and the latter is `sample/raven.example.com@KRBNT-OPERATIONS.EXAMPLE.COM`, and your user name is as specified in the previous paragraph. Then the `grant` statement would be

```
grant CodeBase "file:./SampleClient.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "mjones@KRBNT-OPERATIONS.EXAMPLE.COM" {

    permission java.net.SocketPermission "*", "connect";

    permission javax.security.auth.kerberos.ServicePermission
        "krbtgt/KRBNT-OPERATIONS.EXAMPLE.COM@KRBNT-
OPERATIONS.EXAMPLE.COM",
        "initiate";

    permission javax.security.auth.kerberos.ServicePermission
```

```
        "sample/raven.example.com@KRBNT-OPERATIONS.EXAMPLE.COM",
        "initiate";
};
```

The Server Policy File

Permissions Required by the Login Utility Classes

The `grant` statement in the server policy file for the Login classes is exactly the same as the one in the client policy file, as described in Permissions Required by the Login Utility Classes:

```
grant codebase "file:./Login.jar" {
   permission java.security.AllPermission;
};
```

Permissions Required by SampleServer

The `SampleServer` code does two types of operations for which permissions are required. It

- accepts socket connections.
- accepts establishment of a security context, that is, it is the "acceptor" for security context establishment.

The permission required to accept socket connections is

```
permission java.net.SocketPermission "*", "accept";
```

You may replace the "*" with the hostname or IP address of the machine that `SampleClient` will be running on.

The permission required to accept establishment of a security context is

```
permission javax.security.auth.kerberos.ServicePermission
    "service_principal@your_realm",
    "accept";
```

where `service_principal@your_realm` is the Kerberos name of the service principal that represents `SampleServer` (see Kerberos User and Service Principal Names).

We want to grant the permissions to a specific authenticated user executing `SampleServer` (the service principal considered to represent `SampleServer`), so we specify both the `SampleServer` code location (in `SampleServer.jar`) and a Principal designation indicating the service principal. Suppose this name is `sample/raven.example.com@KRBNT-OPERATIONS.EXAMPLE.COM`. Then the `grant` statement would be

```
  grant CodeBase "file:./SampleServer.jar"
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "sample/raven.example.com@KRBNT-OPERATIONS.EXAMPLE.COM" {
```

```
    permission java.net.SocketPermission "*", "accept";

    permission javax.security.auth.kerberos.ServicePermission
        "sample/raven.example.com@KRBNT-OPERATIONS.EXAMPLE.COM", "accept";
};
```

Running the SampleClient and SampleServer Programs

To execute the `SampleClient` and `SampleServer` programs, do the following:

- Prepare SampleServer for Execution
- Prepare SampleClient for Execution
- Execute SampleServer
- Execute SampleClient

Prepare SampleServer for Execution

To prepare `SampleServer` for execution, do the following:

1. Copy the following files into a directory accessible by the machine on which you will run `SampleServer`:
   - The `Login.java` source file.
   - The `SampleServer.java` source file.
   - The `csLogin.conf` login configuration file.
   - The `server.policy` policy file.
2. Replace `service_principal@your_realm` in `csLogin.conf` with the name of the service principal representing `SampleServer` (see Kerberos User and Service Principal Names).
3. In both places it appears, replace `service_principal@your_realm` in `server.policy` with the Kerberos name of the service principal that represents `SampleServer`. (The same name as that used in the login configuration file.)
4. Compile `Login.java` and `SampleServer.java`:

   ```
   javac Login.java SampleServer.java
   ```

   Note that `Login.java` contains two classes and thus compiling `Login.java` creates `Login.class` and `MyAction.class`.
5. Create a JAR file named `Login.jar` containing `Login.class and MyAction.class`:

   ```
   jar -cvf Login.jar Login.class MyAction.class
   ```
6. Create a JAR file named `SampleServer.jar` containing `SampleServer.class`:

   ```
   jar -cvf SampleServer.jar SampleServer.class
   ```

Prepare SampleClient for Execution

To prepare `SampleClient` for execution, do the following:

1. Copy the following files into a directory accessible by the machine on which you will run `SampleClient`:
   - The `Login.java` source file.
   - The `SampleClient.java` source file.
   - The `csLogin.conf` login configuration file.
   - The `client.policy` policy file.
2. Replace parts of `client.policy`:
   - replace `your_user_name@your_realm` with your user name and realm.
   - replace `your_realm` in `krbtgt/your_realm@your_realm` with your realm.
   - replace `service_principal@your_realm` with the Kerberos name of the service principal that represents `SampleServer` (see Kerberos User and Service Principal Names).
3. Compile `Login.java` and `SampleClient.java`:

   ```
   javac Login.java SampleClient.java
   ```

4. Create a JAR file named `Login.jar` containing `Login.class and MyAction.class`:

   ```
   jar -cvf Login.jar Login.class MyAction.class
   ```

5. Create a JAR file named `SampleClient.jar` containing `SampleClient.class`:

   ```
   jar -cvf SampleClient.jar SampleClient.class
   ```

Execute SampleServer

It is important to execute `SampleServer` before `SampleClient` because `SampleClient` will try to make a socket connection to `SampleServer` and that will fail if `SampleServer` is not yet running and accepting socket connections.

To execute `SampleServer`, be sure to run it on the machine it is expected to be run on. This machine name (host name) is specified as an argument to `SampleClient`. The service principal name appears in several places, including the login configuration file and the policy files.

Go to the directory in which you have prepared `SampleServer` for execution. Execute the `Login` class, specifying

- by an appropriate `-classpath` clause that classes should be searched for in the `Login.jar` and `SampleServer.jar` JAR files,
- by `-Djava.security.manager` that a security manager should be installed,

- by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified. For example, if your realm is `KRBNT-OPERATIONS.EXAMPLE.COM` you'd put `-Djava.security.krb5.realm=KRBNT-OPERATIONS.EXAMPLE.COM`.
- by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified. For example, if your KDC is `samplekdc.example.com` you'd put `-Djava.security.krb5.kdc=samplekdc.example.com`.
- by `-Djava.security.policy=server.policy` that the policy file to be used is `server.policy`, and
- by `-Djava.security.auth.login.config=csLogin.conf` that the login configuration file to be used is `csLogin.conf`.

You pass the name of your application (in this case, `SampleServer`) as an argument to Login. You then add as arguments any arguments required by your application, which in the case of `SampleServer` is a single argument specifying the port number to be used for listening for client connections. Choose a high port number unlikely to be used for anything else. An example would be something like 4444.

Below are the full commands to use for Windows and Solaris, Linux, and macOS. The only difference is that Windows you use semicolons to separate class path items, while you use colons for that purpose on Solaris, Linux, and macOS.

> **Note:**
>
> **Important: In these commands, you must replace *<port_number>* with an appropriate port number, *<your_realm>* with your Kerberos realm, and *<your_kdc>* with your Kerberos KDC.**

Here is the command for Windows:

```
java -classpath Login.jar;SampleServer.jar
 -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djava.security.policy=server.policy
 -Djava.security.auth.login.config=csLogin.conf
 Login SampleServer <port_number>
```

Here is the command for Solaris, Linux, and macOS:

```
java -classpath Login.jar:SampleServer.jar
 -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djava.security.policy=server.policy
 -Djava.security.auth.login.config=csLogin.conf
 Login SampleServer <port_number>
```

Type the full command on one line. Multiple lines are used here for legibility. If the command is too long for your system, you may need to place it in a .bat file (for

Windows) or a .sh file (for Solaris, Linux, and macOS) and then run that file to execute the command.

You will be prompted for the Kerberos password for the service principal. The underlying Kerberos authentication mechanism specified in the login configuration file will log the service principal into Kerberos. Once authentication is successfully completed, the `SampleServer` code will be executed on behalf of the service principal. It will listen for socket connections on the specified port.

For login troubleshooting suggestions, see Troubleshooting.

Execute SampleClient

To execute `SampleClient`, go to the directory in which you have prepared `SampleClient` for execution. Then execute the `Login` class, specifying

- by an appropriate `-classpath` clause that classes should be searched for in the `Login.jar` and `SampleClient.jar` JAR files,
- by `-Djava.security.manager` that a security manager should be installed,
- by `-Djava.security.krb5.realm=<your_realm>` that your Kerberos realm is the one specified.
- by `-Djava.security.krb5.kdc=<your_kdc>` that your Kerberos KDC is the one specified.
- by `-Djava.security.policy=client.policy` that the policy file to be used is `client.policy`, and
- by `-Djava.security.auth.login.config=csLogin.conf` that the login configuration file to be used is `csLogin.conf`.

Pass to Login the name of your application (`SampleClient`) followed by the arguments required by `SampleClient`. The `SampleClient` arguments are (1) the Kerberos name of the service principal that represents `SampleServer` (see Kerberos User and Service Principal Names, (2) the name of the host (machine) on which `SampleServer` is running, and (3) the port number on which `SampleServer` is listening for client connections.

Below are the full commands to use for Windows, Solaris, Linux, and macOS.

> **Note:**
>
> **Important: In these commands, you must replace `<service_principal>`, `<host>`, `<port_number>`, `<your_realm>`, and `<your_kdc>` with appropriate values** (and note that the port number must be the same as the port number passed as an argument to `SampleServer`). These values need not be placed in quotes.

Here is the command for Windows:

```
java -classpath Login.jar;SampleClient.jar
 -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_kdc>
 -Djava.security.policy=client.policy
```

```
 -Djava.security.auth.login.config=csLogin.conf
 Login SampleClient <service_principal> <host> <port_number>
```

Here is the command for Solaris, Linux, and macOS:

```
java -classpath Login.jar:SampleClient.jar
 -Djava.security.manager
 -Djava.security.krb5.realm=<your_realm>
 -Djava.security.krb5.kdc=<your_realm>
 -Djava.security.policy=client.policy
 -Djava.security.auth.login.config=csLogin.conf
 Login SampleClient <service_principal> <host> <port_number>
```

Type the full command on one line. Multiple lines are used here for legibility. As with the command for executing `SampleServer`, if the command is too long to type directly into your command window, place it in a .bat file (Microsoft Windows) or a .sh file (Solaris, Linux, and macOS) and then execute that file.

When prompted, type your Kerberos user name and password. The underlying Kerberos authentication mechanism specified in the login configuration file will log you into Kerberos. Once authentication is successfully completed, the `SampleClient` code will be executed on behalf of you. It will request a socket connection with `SampleServer`. Once `SampleServer` accepts the connection, `SampleClient` and `SampleServer` establish a shared context and then exchange messages as described in this tutorial.

For login troubleshooting suggestions, see Troubleshooting.

More Things You Can Do with Java GSS-API and JAAS

The previous tutorial, Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges, demonstrated how two applications, in particular a client and a server, could use the Java GSS-API to establish a secure context between them and then securely exchange messages.

There are additional operations the context acceptor (the server in our client/server example) can perform once the context has been established with the context initiator (the client). Basically, the server can "impersonate" the client. The level of impersonation depends upon whether or not the client has delegated credentials to the server.

- Executing Code on Behalf of the Client User
- Using Credentials Delegated from the Client

Executing Code on Behalf of the Client User

One possible type of client impersonation the server can do is causing code to be executed on behalf of the same entity (user) the client code is executed on behalf of. Normally, a method executed by a thread uses the access control settings for that thread itself. However, when impersonating a client in this tutorial, the server uses the client's access control settings so that the server has access to exactly those resources that the client itself has when it runs.

One major benefit of the approach used in this tutorial is that the JAAS authorization component can be used for access control. Without the JAAS authorization component, the server principal would need access to any resources accessed by the code executed on behalf of the client user, and the server code would need to include access control logic to determine whether the user was authorized to access such resources. By utilizing JAAS authorization, providing principal-based access control, the access control is handled automatically. Permissions for the security-sensitive operations in such code only need to be granted to that user and not also to the server code. See the JAAS Authorization tutorial for more information on JAAS authorization.

- Basic Approach
- Sample Code and Policy File
- Running the Sample Code

Basic Approach

How does the server "impersonate" the client to execute code on behalf of the user running the client code? Essentially the same way the client code is set up to be run on behalf of that user. All the server code needs to know is the user's principal name, which it can obtain from the context established with the client.

Recall that JAAS authentication of the user executing the client code results in creation of a `Subject` containing a `Principal` with the user (principal) name. The Subject is subsequently associated with a new access control context (via a `Subject.doAsPrivileged` call from the Login utility) and the client code is considered to be executed on behalf of the user; subsequent access control decisions are based on whether or not that particular user, executing the client code, is granted the required permissions.

The server code is similarly handled, except in that case the Principal specified for authentication is typically a "service principal", not a user principal. Again, a `Subject` containing a `Principal` with the specified principal name is created, `Subject.doAsPrivileged` is called, and the server code is considered to be executed on behalf of the specified principal; subsequent access control decisions are based on whether or not that particular principal, executing the server code, is granted the required permissions.

Once the client and server have established a mutual context, the context initiator's name (the client's principal name) can be determined by the following:

```
GSSName clientGSSName = context.getSrcName();
```

The context acceptor (the server) can use this name to construct a `Subject` containing a `Principal` that represents the same entity. For example, you can construct such a `Subject` with Oracle's JDK via the following:

```
Subject client =
  com.sun.security.jgss.GSSUtil.createSubject(clientGSSName, null);
```

The `createSubject` method creates a new Subject from the GSSName and GSSCredential specified as arguments. If the server code is just going to execute code on behalf of the user in the local JVM, the user's credentials are not required – and in fact cannot even be obtained unless the client has delegated credentials to the

server, as discussed in Using Credentials Delegated from the Client. Since the credentials are not needed here, we pass a `null` for the GSSCredential argument.

> **Note:**
>
> Note: If you are not using Oracle's JDK, an alternative way to do this is to construct a `KerberosPrincipal` instance as follows:
>
> ```
> KerberosPrincipal principal =
>  new KerberosPrincipal(clientGSSName.toString());
> ```
>
> Then use this principal to construct a new `Subject` or populate this principal in the principal set of an existing `Subject`.
>
> The code that the server would like to execute on behalf of the user should be initiated from the `run` method of a class that implements `java.security.PrivilegedAction` (or `java.security.PrivilegedExceptionAction`). That is, the code can either be in such a `run` method or invoked from such a `run` method.
>
> The server code can pass the `Subject`, along with an instance of the `PrivilegedAction` (or `PrivilegedExceptionAction`), to `Subject.doAsPrivileged` to execute the subsequent code, starting with the `run` method in the `PrivilegedAction`, on behalf of the principal (user) in the specified Subject.
>
> For example, suppose the `PrivilegedAction` class is called `ReadFileAction` and it takes as an argument a `String` with the principal name. You can create an instance of this class by
>
> ```
> String clientName = clientGSSName.toString();
> PrivilegedAction readFile =
>  new ReadFileAction(clientName);
> ```
>
> The call to `doAsPrivileged` is then
>
> ```
> Subject.doAsPrivileged(client, readFile, null);
> ```

Sample Code and Policy File

The following sample code and policy file illustrate the server impersonating the client in order to execute code whose security-sensitive operations are only permitted to be done by the specific user executing the client.

- `SampleServerImp.java`
- `ReadFileAction.java`
- `serverimp.policy`

SampleServerImp.java

The `SampleServerImp.java` file is exactly the same as the `SampleServer.java` file from the previous (Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges) tutorial, except that after exchanging messages with the client, it has the following code to perform a `ReadFileAction` as the client user:

```
System.out.println("Impersonating client.");

/*
 * Extract the KerberosPrincipal from the client GSSName and
 * populate it in the principal set of a new Subject. Pass in a
 * null for credentials since credentials will not be needed.
 */
GSSName clientGSSName = context.getSrcName();
System.out.println("clientGSSName: " + clientGSSName);
Subject client =
   com.sun.security.jgss.GSSUtil.createSubject(clientGSSName,
        null);

/*
* Construct an action that will read a file meant only for the
* client
*/
String clientName = clientGSSName.toString();
PrivilegedAction readFile =
   new ReadFileAction(clientName);

/*
* Invoke the action via a doAsPrivileged. This allows the
* action to be executed as the client subject, and it also
* runs that code as privileged. This means that any permission
* checking that happens beyond this point applies only to
* the code being run as the client.
*/
Subject.doAsPrivileged(client, readFile, null);
```

ReadFileAction.java

The `ReadFileAction.java` file contains the `ReadFileAction` class. Its constructor takes as an argument a `String` for the name of the client user. The client user name is used to construct a file name for a file from which `ReadFileAction` will attempt to read. The file name will be:

```
./data/<name>_info.txt
```

where <*name*> is the client user name without its corresponding realm. For example, if the full user name is `mjones@KRBNT-OPERATIONS.EXAMPLE.COM`, then the file name is

```
./data/mjones_info.txt
```

> **Note:**
>
> On Window, the forward slashes will be backward slashes.

The `ReadFileAction` `run` method reads the specified file and prints its contents.

serverimp.policy

`ReadFileAction` attempts to read a file, which is a security-checked operation. Since `ReadFileAction` is considered to be executed as the client user (`Principal`), the appropriate permission must be granted not only to the `ReadFileAction` code itself, but to the client `Principal` as well.

Assuming the `ReadFileAction` class is placed in a JAR file named `ReadFileAction.jar`, and the user principal name is `mjones@KRBNT-OPERATIONS.EXAMPLE.COM`, this permission can be granted via the following in a policy file:

```
grant CodeBase "file:./ReadFileAction.jar"
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "mjones@KRBNT-OPERATIONS.EXAMPLE.COM" {

    permission java.io.FilePermission "data/mjones_info.txt",
        "read";
};
```

The `serverimp.policy` file is exactly the same as the `server.policy` file from the previous (Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges) tutorial, except that it grants the `SampleServer` code the `javax.security.auth.AuthPermission "doAsPrivileged"` permission it needs in order to call the `doAsPrivileged` method, and it has the following placeholder for granting the `FilePermission` shown above:

```
grant CodeBase "file:./ReadFileAction.jar"
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm" {

    permission java.io.FilePermission "data/your_user_name_info.txt",
        "read";
};
```

You must substitute your Kerberos realm for `your_realm`, and your user name for `your_user_name` in both `your_user_name@your_realm` and `data/your_user_name_info.txt`. If you are working on Windows, you also replace the "/" in `data/your_user_name_info.txt` with a "\".

Running the Sample Code

To run the sample code illustrating the server impersonating the client, do everything listed in Running the SampleClient and SampleServer Programs in the previous tutorial, except for the following:

- In the "Prepare `SampleServer` for Execution" step:
  - Use `SampleServerImp.java` instead of `SampleServer.java`. Compile it and create a JAR file named `SampleServerImp.jar` containing `SampleServerImp.class` via the following:

    ```
    javac SampleServerImp.java
    jar -cvf SampleServerImp.jar SampleServerImp.class
    ```

  - Use the `serverimp.policy` policy file instead of `server.policy`.
  - Use the `csImpLogin.conf` login configuration file instead of `cs.conf`.
  - Copy `ReadFileAction.java` to the same directory as the other files. Compile it and place it in a JAR file via the following:

    ```
    javac ReadFileAction.java
    jar -cvf ReadFileAction.jar ReadFileAction.class
    ```

  - In `csImpLogin.conf`, replace `service_principal@your_realm` with the Kerberos name of the service principal that represents `SampleServer` (see Kerberos User and Service Principal Names).
  - In `serverimp.policy`, replace `service_principal@your_realm` in both places it appears with the Kerberos name of the service principal that represents `SampleServer`. (The same name as that used in the login configuration file.) In addition, substitute your Kerberos realm for `your_realm`, and your user name for `your_user_name` in both `your_user_name@your_realm` and `data/your_user_name_info.txt`. If you are running on Windows, then replace the "/" in `data/your_user_name_info.txt` with a "\".
  - Create a `data` subdirectory of your current directory and create a short text file of the specified name in that directory. For example, if your user name is `mjones`, the file to be placed in the data subdirectory should be named `mjones_info.txt`.
- In the "Execute `SampleServer`" step:
  - Use the following commands instead of those specified in that section so that `SampleServerImp` is executed, `serverimp.policy` and `csImpLogin.conf` are used, and `ReadFileAction.jar` is included.

    > **Note:**
    >
    > **Important: In these commands, you must replace <*port_number*> with an appropriate port number (a high port number such as 4444), <*your_realm*> with your Kerberos realm, and <*your_kdc*> with your Kerberos KDC.**

    Here is the command for Windows:

    ```
    java -classpath Login.jar;SampleServerImp.jar;ReadFileAction.jar
      -Djava.security.manager
      -Djava.security.krb5.realm=<your_realm>
      -Djava.security.krb5.kdc=<your_kdc>
    ```

```
  -Djava.security.policy=serverimp.policy
  -Djava.security.auth.login.config=csImpLogin.conf
  Login SampleServerImp <port_number>
```

Here is the command for Solaris, Linux, and macOS:

```
java -classpath Login.jar:SampleServerImp.jar:ReadFileAction.jar
  -Djava.security.manager
  -Djava.security.krb5.realm=<your_realm>
  -Djava.security.krb5.kdc=<your_kdc>
  -Djava.security.policy=serverimp.policy
  -Djava.security.auth.login.config=csImpLogin.conf
  Login SampleServerImp <port_number>
```

As usual, type the full command on one line. Multiple lines are used here for legibility. If the command is too long for your system, you may need to place it in a .bat file (for Windows) or a .sh file (for Solaris, Linux, and macOS) and then run that file to execute the command.

As when running `SampleServer`, you will be prompted for the Kerberos password for the service principal under which `SampleServerImp` is expected to be run. The Kerberos login module specified in the login configuration file will log the service principal into Kerberos. Once authentication is successfully completed, the `SampleServerImp` code will be executed on behalf of the service principal. It will listen for socket connections on the specified port.

After you follow the "Prepare `SampleClient` for Execution" and "Execute `SampleClient`" instructions as usual and perform the user login, the client code will request a socket connection with `SampleServerImp`. Once `SampleServerImp` accepts the connection, `SampleClient` and `SampleServerImp` establish a shared context and then exchange messages as described in the previous tutorial.

After the message exchange, `SampleServerImp` determines the principal name of the user executing the client code, creates a new `Subject` containing a `Principal` with that name, and calls `Subject.doAsPrivileged` to execute the code in `ReadFileAction` on behalf of the specified user. `ReadFileAction` reads the file named `your_user_name_info.txt` (where `your_user_name` represents the actual user name) in the `data` subdirectory of the current directory, and prints out its contents.

For login troubleshooting suggestions, see Troubleshooting.

Using Credentials Delegated from the Client

The most complete type of client impersonation is possible if the client delegates its credentials to the server.

Recall that prior to context establishment with the context acceptor (the server in our previous tutorial), the context initiator (the client) sets various context options. If the initiator calls the `requestCredDeleg` method on the `context` object with a `true` argument, as in

```
context.requestCredDeleg(true);
```

then this requests that the initiator's credentials be delegated to the acceptor during context establishment.

Delegation of credentials from the initiator to the acceptor enables the acceptor to authenticate itself as an agent or delegate of the initiator.

First, after context establishment, the acceptor must determine whether or not credential delegation actually took place. It does so by calling the `getCredDelegState` method:

```
boolean delegated = context.getCredDelegState();
```

If credentials were delegated, the acceptor can obtain those credentials by calling the `getDelegCr` method:

```
GSSCredential clientCr = context.getDelegCred();
```

The resulting `GSSCredential` object can then be used to initiate subsequent GSS-API contexts as a "delegate" of the initiator. For example, the server could authenticate as the client to a backend server that cares more about who the original client was than who the intermediate server is.

Acting as the client, the server can establish a connection with the backend server, establish a joint security context, and exchange messages in basically the same manner that the client and server did.

One way it could be done is that when the server calls the `createContext` method of a `GSSManager`, it could pass `createContext` the delegated credentials instead of passing a `null`.

Alternatively, the server code could first call the `com.sun.security.jgss.GSSUtil` `createSubject` method and pass it the delegated credentials. That method returns a `Subject` containing those credentials as the default credentials. The server could then associate this `Subject` with the current `AccessControlContext`, as described in How Do You Associate a Subject with an Access Control Context? in the JAAS Authorization tutorial. Then, when the server code calls the `GSSManager` `createContext` method, it can pass a null (indicating the credentials for the "current" Subject should be used). In other words, the server would effectively become the client. Subsequent connections to backend servers using GSS could be made exactly as described in the previous tutorials. This approach is useful if you want the code that will use the delegated credentials to be identical to the code that uses the default local credentials.

Constrained Delegation

If constrained delegation is configured in a KDC server, then, on the server side, the `getCredDelegState()` call might still return true and `getDelegCred()` would return delegated credentials, depending on the KDC settings, even if the client has not called `requestCredDeleg(true)`.

Permission Required In Order to Delegate Credentials

In order to delegate credentials, the context initiator (`SampleClient` in our previous tutorial) must have a `javax.security.auth.kerberos.DelegationPermission`. An example using placeholders in italics for actual values is the following:

```
permission javax.security.auth.kerberos.DelegationPermission
  "\"service_principal@your_realm\"
     \"krbtgt/your_realm@your_realm\"";
```

Note that `DelegationPermission` has a single target in quotes that contains two items, both of which are quoted. Each inner quote is escaped by a "\". Thus the first item is

```
"service_principal@your_realm"
```

and the second is

```
"krbtgt/your_realm@your_realm"
```

This basically gives the code executing on behalf of the client the permission to forward a Kerberos ticket to the specified peer (*`service_principal`*), where the Kerberos ticket is meant to avail service from `krbtgt/`*`your_realm@your_realm`*.

Substitute your realm for all places *`your_realm`* appears. Also substitute the service principal name for the service principal representing the server for *`service_principal@your_realm`*. (See Kerberos User and Service Principal Names in the previous tutorial.) Suppose your realm is `KRBNT-OPERATIONS.EXAMPLE.COM` and the service principal is `sample/raven.example.com@KRBNT-OPERATIONS.EXAMPLE.COM`. Then the permission could appear in a policy file as

```
permission javax.security.auth.kerberos.DelegationPermission
  "\"sample/raven.example.com@KRBNT-OPERATIONS.EXAMPLE.COM\"
     \"krbtgt/KRBNT-OPERATIONS.EXAMPLE.COM@KRBNT-OPERATIONS.EXAMPLE.COM\"";
```

Kerberos Requirements

Kerberos Version 5 is used for both the authentication and secure communication aspects of the client and server applications developed in this tutorial. The reader is assumed to already be familiar with Kerberos. See the Kerberos reference documentation.

The JAAS framework, and the Kerberos mechanism required by the Java GSS-API methods, are built into JDKs from all vendors. The Kerberos LoginModule required for the JAAS authentication in this tutorial may not be available in all vendors' JDKs. We will be using the LoginModule for Kerberos provided by Oracle's JDK.

In order to run the sample programs, you will need access to a Kerberos installation. As described in the following sections, you may also need a `krb5.conf` Kerberos configuration file and an indication as to where that file is located.

As with all Kerberos installations, a Kerberos Key Distribution Center (KDC) is required. It needs to contain the user name and password you will use to be authenticated to Kerberos.

> **Note:**
>
> A KDC implementation is part of a Kerberos installation and not a part of the JDK.

As with most Kerberos installations, a Kerberos configuration file `krb5.conf` is consulted to determine such things as the default realm and KDC. If you are using a Kerberos implementation that does not include a `krb5.conf` file (such as one from Windows), you will either need to create one or use system properties as described in Setting Properties to Indicate the Default Realm and KDC.

Setting Properties to Indicate the Default Realm and KDC

Typically, the default realm and the KDC for that realm are indicated in the Kerberos `krb5.conf` configuration file. However, if you like, you can instead specify these values by setting the following system properties to indicate the realm and KDC, respectively:

```
java.security.krb5.realm
java.security.krb5.kdc
```

If you set one of these properties you must set them both.

Also note that if you set these properties, then no cross-realm authentication is possible unless a `krb5.conf` file is also provided from which the additional information required for cross-realm authentication may be obtained.

If you set values for these properties, then they override the default realm and KDC values specified in `krb5.conf` (if such a file is found). The `krb5.conf` file is still consulted if values for items other than the default realm and KDC are needed. If no `krb5.conf` file is found, then the default values used for these items are implementation-specific.

Locating the krb5.conf Configuration File

The essential Kerberos configuration information is the default realm and the default KDC. As shown in Setting Properties to Indicate the Default Realm and KDC, if you set properties to indicate these values, they are not obtained from a `krb5.conf` configuration file.

If these properties do not have values set, or if other Kerberos configuration information is needed, an attempt is made to find the required information in a `krb5.conf` file. The algorithm to locate the `krb5.conf` file is the following:

- If the system property `java.security.krb5.conf` is set, its value is assumed to specify the path and file name.
- If that system property value is not set, then the configuration file is looked for in the directory

  - `<java-home>\conf\security` (Windows)
  - `<java-home>/conf/security` (Solaris, Linux, and macOS)

  Here `<java-home>` refers to the directory where the JDK is installed.

- If the file is still not found, then an attempt is made to locate it as follows:
  - `/etc/krb5/krb5.conf` (Solaris)
  - `C:\Windows\krb5.ini` (Windows)
  - `/etc/krb5.conf` (Linux)
  - `~/Library/Preferences/edu.mit.Kerberos`, `/Library/Preferences/edu.mit.Kerberos`, or `/etc/krb5.conf` (macOS)
- If the file is still not found, and the configuration information being searched for is *not* the default realm and KDC, then implementation-specific defaults are used. If, on the other hand, the configuration information being searched for is the default realm and KDC because they weren't specified in system properties, and the `krb5.conf` file is not found either, then an exception is thrown.
- On Windows, if a `krb5.conf` file cannot be found or it does not contain settings for the default realm and its KDC, then the environment variables `USERDNSDOMAIN` and `LOGONSERVER` are used as the default realm and its KDC.

Naming Conventions for Realm Names and Hostnames

By convention, all Kerberos realm names are uppercase and all DNS hostname and domain names are lowercase. On Windows, domains are also Kerberos realms; however, the realm name is always the uppercase version of the domain name.

Hostnames are case insensitive and by convention they are all lowercase. They must resolve to the same hostname on the client and server by their respective naming services.

However, in the Kerberos database hostnames are case sensitive. In all host-based Kerberos service principals in the KDC, hostnames are case-sensitive. The hostnames used in the Kerberos service principal names must exactly match the hostnames returned by the naming service. For example, if the naming service returns a fully qualified lowercased DNS hostname, such as `raven.example.com`, then the administrator must use the same fully qualified lowercased DNS hostname when creating host-based principal names in the KDC: `host/raven.example.com`.

Cross-Realm Authentication

In cross-realm authentication, a principal in one realm can authenticate to principals in another realm.

In Kerberos, cross-realm authentication is implemented by sharing an encryption key between two realms. The KDCs in two different realms share a special cross-realm secret; this secret is used to prove identity when crossing the boundary between realms.

The key that is shared is the Ticket Granting Service principal's key. Here's a typical Ticket Granting Service principal for a single realm:

```
ktbtgt/EXAMPLE.COM@EXAMPLE.COM
```

In cross realm authentication, two principals are created on each participating realm. For two realms, `ENG.EAST.EXAMPLE.COM` and `SALES.WEST.EXAMPLE.COM`, these principals would be:

```
krbtgt/ENG.EAST.EXAMPLE.COM@SALES.WEST.EXAMPLE.COM
krbtgt/SALES.WEST.EXAMPLE.COM@ENG.EAST.EXAMPLE.COM
```

These principals, known as remote Ticket Granting Server principals, must be created on both realms.

For a Windows KDC, the `krbtgt` account is created automatically when a Windows domain is created. This account cannot be deleted and renamed.

Troubleshooting

Below are listed some problems that may occur when attempting a login, and suggestions for solving them.

- **Configurable Kerberos Settings:** The Kerberos Key Distribution Center (KDC) name and realm settings are provided in the Kerberos configuration file or via the system properties `java.security.krb5.kdc` and `java.security.krb5.realm.` A boolean option `refreshKrb5Config` can be specified in the entry for `Krb5LoginModule` in the JAAS configuration file. If this option is set to `true`, then the configuration values will be refreshed before the login method of the `Krb5LoginModule` is called.

  > **Note:**
  >
  > When switching Kerberos configurations, it is REQUIRED that `refreshKrb5Config` should be set to `true.` Failure to set this value can lead to unexpected results.

- **java.lang.SecurityException at javax.security.auth.login.Configuration.getConfiguration**

  Cause: There was a problem processing the JAAS login configuration file, possibly due to a syntax error in the file.

  Solution: Check the configuration file carefully for errors. See Appendix B: JAAS Login Configuration File for information about the syntax required in the login configuration file.

- **javax.security.auth.login.LoginException: KrbException: Pre-authentication information was invalid (24) - Preauthentication failed**

  Cause 1: The password entered is incorrect.

  Solution 1: Verify the password.

  Cause 2: If you are using the keytab to get the key (e.g., by setting the `useKeyTab` option to `true` in the Krb5LoginModule entry in the JAAS login configuration file), then the key might have changed since you updated the keytab.

  Solution 2: Consult your Kerberos documentation to generate a new keytab and use that keytab.

Cause 3: Clock skew - If the time on the KDC and on the client differ significantly (typically 5 minutes), this error can be returned.

Solution 3: Synchronize the clocks (or have a system administrator do so).*Cause 4:* The Kerberos realm name is not all uppercase.

Solution 4: Make the Kerberos realm name all uppercase. **Note**: It is recommended to have all uppercase realm names. See Naming Conventions for Realm Names and Hostnames.

- **GSSException: No valid credentials provided (Mechanism level: Attempt to obtain new INITIATE credentials failed! (null)) . . . Caused by: javax.security.auth.login.LoginException: Clock skew too great**

  Cause: Kerberos requires the time on the KDC and on the client to be loosely synchronized. (The default is within 5 minutes.) If that's not the case, you will get this error.

  Solution: Synchronize the clocks (or have a system administrator do so).

- **javax.security.auth.login.LoginException: KrbException: Null realm name (601) - default realm not specified**

  Cause: The default realm is not specified in the Kerberos configuration file `krb5.conf` (if used), provided as a part of the user name, or specified via the `java.security.krb5.realm` system property.

  Solution: Verify that your Kerberos configuration file (if used) contains an entry specifying the default realm, or directly specify it by setting the value of the `java.security.krb5.realm` system property and/or including it in your user name when authenticating using Kerberos.

- **javax.security.auth.login.LoginException: java.net.SocketTimeoutException: Receive timed out**

  Solution: Verify that the Kerberos KDC is up and running.

- **GSSException: No valid credentials provided (Mechanism level: Failed to find any Kerberos Ticket)**

  Cause: This may occur if no valid Kerberos credentials are obtained. In particular, this occurs if you want the underlying mechanism to obtain credentials but you forgot to indicate this by setting the `javax.security.auth.useSubjectCredsOnly` system property value to `false` (for example via `-Djavax.security.auth.useSubjectCredsOnly=false` in your execution command).

  Solution: Be sure to set the `javax.security.auth.useSubjectCredsOnly` system property value to `false` if you want the underlying mechanism to obtain credentials, rather than your application or a wrapper program (such as the Login utility used by some of the tutorials) performing authentication using JAAS.

- **javax.security.auth.login.LoginException: Could not load configuration file <krb5.conf> (No such file or directory)**

  Cause: The tutorials' sample execution commands specify the default Kerberos realm and KDC by setting values for the `java.security.krb5.realm` and `java.security.krb5.kdc` system properties. If you like, you can instead have a `krb5.conf` Kerberos configuration file used. Such a file includes information about what the default realm and KDC are. To use a `krb5.conf` file, you either set the system property `java.security.krb5.conf` (instead of the `realm` and `kdc` properties) to specify the location of the file or you don't set any of these properties

and therefore an attempt is made to locate the `krb5.conf` file in a default location. You will get the error "Could not load configuration file <krb5.conf> (No such file or directory)" if the file could not be found.

Solution: Verify that the Kerberos configuration file `krb5.conf` is available and readable. Check Kerberos Requirements for information about how to specify the location of the `krb5.conf` file and where such a file is searched for by default if you don't explicitly indicate the location.

- **javax.security.auth.login.LoginException: KrbException: KDC has no support for encryption type (14) - KDC has no support for encryption type**

  Cause 1: Your KDC does not support the encryption type requested.

  Solution 1: Oracle's implementation of Kerberos supports the following encryption types: `aes256-cts-hmac-sha1-96`, `aes128-cts-hmac-sha1-96`, `des3-cbc-sha1`, `arcfour-hmac-md5`, `des-cbc-crc`, and `des-cbc-md5`.

  Applications can select the desired encryption type by specifying following tags in the Kerberos Configuration file `krb5.conf`:

  ```
  [libdefaults]
  default_tkt_enctypes = des-cbc-md5 des-cbc-crc des3-cbc-sha1
  default_tgs_enctypes = des-cbc-md5 des-cbc-crc des3-cbc-sha1
  permitted_enctypes = des-cbc-md5 des-cbc-crc des3-cbc-sha1
  ```

  If not specified, then the default value is:

  ```
  aes256-cts-hmac-sha1-96 aes128-cts-hmac-sha1-96 des3-cbc-sha1 arcfour-
  hmac-md5
  ```

  If `allow_weak_crypto` in `krb5.conf` is set to true, then `des-cbc-crc` and `des-cbc-md5` are also supported.

  Cause 2: This exception is thrown when using native ticket cache on some Windows platforms. Microsoft has added a new feature in which they no longer export the session keys for Ticket-Granting Tickets (TGTs). As a result, the native TGT obtained on Windows has an "empty" session key and null EType.

  Solution 2: You need to update the Windows registry to disable this new feature. The registry key **`allowtgtsessionkey`** should be added – and set correctly – to allow session keys to be sent in the Kerberos Ticket-Granting Ticket.

  See Registry Key to Allow Session Keys to Be Sent in Kerberos Ticket-Granting-Ticket from Microsoft Support for more information. Usually, the following is the required registry setting:

  ```
  HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Kerberos
  \Parameters
  Value Name: allowtgtsessionkey
  Value Type: REG_DWORD
  Value: 0x01  ( default is 0 )
  ```

  By default, the value is 0; setting it to "0x01" allows a session key to be included in the TGT.

- **KDC reply did not match expectations**

Cause: The KDC sent a response that cannot be understood by the client.

Solution: Verify that you have set correctly all the `krb5.conf` file configuration parameters and consult your KDC vendor's guide.

> **Note:**
>
> A debugging mode can be enabled by setting the system property `sun.security.krb5.debug` to "true". This setting allows you to follow the program's execution of the Kerberos V5 protocol.

Source Code for JAAS and Java GSS-API Tutorials

SampleServer.java

```
import org.ietf.jgss.*;
import java.io.*;
import java.net.Socket;
import java.net.ServerSocket;

/**
 * A sample server application that uses JGSS to do mutual authentication
 * with a client using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the client.
 *
 * Every message exchanged with the client includes a 4-byte application-
 * level header that contains the big-endian integer value for the number
 * of bytes that will follow as part of the JGSS token.
 *
 * The protocol is:
 *    1.  Context establishment loop:
 *         a. client sends init sec context token to server
 *         b. server sends accept sec context token to client
 *         ....
 *    2. client sends a wrap token to the server.
 *    3. server sends a mic token to the client for the application
 *       message that was contained in the wrap token.
 */

public class SampleServer  {

    public static void main(String[] args)
        throws IOException, GSSException {

        // Obtain the command-line arguments and parse the port number

        if (args.length != 1) {
            System.err.println("Usage: java <options> Login SampleServer
<localPort>");
            System.exit(-1);
        }
```

```
        int localPort = Integer.parseInt(args[0]);

        ServerSocket ss = new ServerSocket(localPort);

        GSSManager manager = GSSManager.getInstance();

        while (true) {

            System.out.println("Waiting for incoming connection...");

            Socket socket = ss.accept();
            DataInputStream inStream =
                new DataInputStream(socket.getInputStream());
            DataOutputStream outStream =
                new DataOutputStream(socket.getOutputStream());

            System.out.println("Got connection from client "
                               + socket.getInetAddress());

            /*
             * Create a GSSContext to receive the incoming request
             * from the client. Use null for the server credentials
             * passed in. This tells the underlying mechanism
             * to use whatever credentials it has available that
             * can be used to accept this connection.
             */
            GSSContext context =
manager.createContext((GSSCredential)null);

            // Do the context eastablishment loop

            byte[] token = null;

            while (!context.isEstablished()) {

                token = new byte[inStream.readInt()];
                System.out.println("Will read input token of size "
                                   + token.length
                                   + " for processing by
acceptSecContext");
                inStream.readFully(token);

                token = context.acceptSecContext(token, 0, token.length);

                // Send a token to the peer if one was generated by
                // acceptSecContext
                if (token != null) {
                    System.out.println("Will send token of size "
                                       + token.length
                                       + " from acceptSecContext.");
                    outStream.writeInt(token.length);
                    outStream.write(token);
                    outStream.flush();
                }
            }
```

```
System.out.print("Context Established! ");
System.out.println("Client is " + context.getSrcName());
System.out.println("Server is " + context.getTargName());
/*
 * If mutual authentication did not take place, then
 * only the client was authenticated to the
 * server. Otherwise, both client and server were
 * authenticated to each other.
 */
if (context.getMutualAuthState())
    System.out.println("Mutual authentication took place!");

/*
 * Create a MessageProp which unwrap will use to return
 * information such as the Quality-of-Protection that was
 * applied to the wrapped token, whether or not it was
 * encrypted, etc. Since the initial MessageProp values
 * are ignored, just set them to the defaults of 0 and false.
 */
MessageProp prop = new MessageProp(0, false);

/*
 * Read the token. This uses the same token byte array
 * as that used during context establishment.
 */
token = new byte[inStream.readInt()];
System.out.println("Will read token of size "
                   + token.length);
inStream.readFully(token);

byte[] bytes = context.unwrap(token, 0, token.length, prop);
String str = new String(bytes);
System.out.println("Received data \""
                   + str + "\" of length " + str.length());

System.out.println("Confidentiality applied: "
                   + prop.getPrivacy());

/*
 * Now generate a MIC and send it to the client. This is
 * just for illustration purposes. The integrity of the
 * incoming wrapped message is guaranteed irrespective of
 * the confidentiality (encryption) that was used.
 */

/*
 * First reset the QOP of the MessageProp to 0
 * to ensure the default Quality-of-Protection
 * is applied.
 */
prop.setQOP(0);

token = context.getMIC(bytes, 0, bytes.length, prop);
```

```
            System.out.println("Will send MIC token of size "
                               + token.length);
            outStream.writeInt(token.length);
            outStream.write(token);
            outStream.flush();

            System.out.println("Closing connection with client "
                               + socket.getInetAddress());
            context.dispose();
            socket.close();
        }
    }
}
```

bcsLogin.conf

```
/**
 * Login Configuration for JAAS.
 */

com.sun.security.jgss.initiate {
  com.sun.security.auth.module.Krb5LoginModule required;
};

com.sun.security.jgss.accept {
  com.sun.security.auth.module.Krb5LoginModule required storeKey=true;
};
```

SampleClient.java

```
import org.ietf.jgss.*;
import java.net.Socket;
import java.io.IOException;
import java.io.DataInputStream;
import java.io.DataOutputStream;

/**
 * A sample client application that uses JGSS to do mutual authentication
 * with a server using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the server.
 *
 * Every message sent to the server includes a 4-byte application-level
 * header that contains the big-endian integer value for the number
 * of bytes that will follow as part of the JGSS token.
 *
 * The protocol is:
 *    1.  Context establishment loop:
 *         a. client sends init sec context token to server
 *         b. server sends accept sec context token to client
 *         ....
 *    2. client sends a wrap token to the server.
 *    3. server sends a MIC token to the client for the application
 *       message that was contained in the wrap token.
 */
```

```
public class SampleClient {

    public static void main(String[] args)
       throws IOException, GSSException  {

        // Obtain the command-line arguments and parse the port number

        if (args.length < 3) {
            System.err.println("Usage: java <options> Login SampleClient "
                               + " <server> <hostName> <port>");
            System.exit(-1);
        }

        String server = args[0];
        String hostName = args[1];
        int port = Integer.parseInt(args[2]);

        Socket socket = new Socket(hostName, port);
        DataInputStream inStream =
          new DataInputStream(socket.getInputStream());
        DataOutputStream outStream =
          new DataOutputStream(socket.getOutputStream());

        System.out.println("Connected to server "
                           + socket.getInetAddress());

        /*
         * This Oid is used to represent the Kerberos version 5 GSS-API
         * mechanism. It is defined in RFC 1964. We will use this Oid
         * whenever we need to indicate to the GSS-API that it must
         * use Kerberos for some purpose.
         */
        Oid krb5Oid = new Oid("1.2.840.113554.1.2.2");

        GSSManager manager = GSSManager.getInstance();

        /*
         * Create a GSSName out of the server's name. The null
         * indicates that this application does not wish to make
         * any claims about the syntax of this name and that the
         * underlying mechanism should try to parse it as per whatever
         * default syntax it chooses.
         */
        GSSName serverName = manager.createName(server, null);

        /*
         * Create a GSSContext for mutual authentication with the
         * server.
         *    - serverName is the GSSName that represents the server.
         *    - krb5Oid is the Oid that represents the mechanism to
         *      use. The client chooses the mechanism to use.
         *    - null is passed in for client credentials
         *    - DEFAULT_LIFETIME lets the mechanism decide how long the
         *      context can remain valid.
```

```
 * Note: Passing in null for the credentials asks GSS-API to
 * use the default credentials. This means that the mechanism
 * will look among the credentials stored in the current Subject
 * to find the right kind of credentials that it needs.
 */
GSSContext context = manager.createContext(serverName,
                                krb5Oid,
                                null,
                                GSSContext.DEFAULT_LIFETIME);

// Set the desired optional features on the context. The client
// chooses these options.

context.requestMutualAuth(true);  // Mutual authentication
context.requestConf(true);  // Will use confidentiality later
context.requestInteg(true); // Will use integrity later

// Do the context eastablishment loop

byte[] token = new byte[0];

while (!context.isEstablished()) {

    // token is ignored on the first call
    token = context.initSecContext(token, 0, token.length);

    // Send a token to the server if one was generated by
    // initSecContext
    if (token != null) {
        System.out.println("Will send token of size "
                           + token.length
                           + " from initSecContext.");
        outStream.writeInt(token.length);
        outStream.write(token);
        outStream.flush();
    }

    // If the client is done with context establishment
    // then there will be no more tokens to read in this loop
    if (!context.isEstablished()) {
        token = new byte[inStream.readInt()];
        System.out.println("Will read input token of size "
                           + token.length
                           + " for processing by initSecContext");
        inStream.readFully(token);
    }
}

System.out.println("Context Established! ");
System.out.println("Client is " + context.getSrcName());
System.out.println("Server is " + context.getTargName());

/*
 * If mutual authentication did not take place, then only the
 * client was authenticated to the server. Otherwise, both
```

```
         * client and server were authenticated to each other.
         */
        if (context.getMutualAuthState())
            System.out.println("Mutual authentication took place!");

        byte[] messageBytes = "Hello There!\0".getBytes();

        /*
         * The first MessageProp argument is 0 to request
         * the default Quality-of-Protection.
         * The second argument is true to request
         * privacy (encryption of the message).
         */
        MessageProp prop =  new MessageProp(0, true);

        /*
         * Encrypt the data and send it across. Integrity protection
         * is always applied, irrespective of confidentiality
         * (i.e., encryption).
         * You can use the same token (byte array) as that used when
         * establishing the context.
         */

        token = context.wrap(messageBytes, 0, messageBytes.length, prop);
        System.out.println("Will send wrap token of size " + token.length);
        outStream.writeInt(token.length);
        outStream.write(token);
        outStream.flush();

        /*
         * Now we will allow the server to decrypt the message,
         * calculate a MIC on the decrypted message and send it back
         * to us for verification. This is unnecessary, but done here
         * for illustration.
         */

        token = new byte[inStream.readInt()];
        System.out.println("Will read token of size " + token.length);
        inStream.readFully(token);
        context.verifyMIC(token, 0, token.length,
                          messageBytes, 0, messageBytes.length,
                          prop);

        System.out.println("Verified received MIC for message.");

        System.out.println("Exiting...");
        context.dispose();
        socket.close();
    }
}
```

JaasAcn.java

```
import javax.security.auth.*;
import javax.security.auth.callback.*;
import javax.security.auth.login.*;
import com.sun.security.auth.callback.TextCallbackHandler;

/**
 * This JaasAcn application attempts to authenticate a user
 * and reports whether or not the authentication was successful.
 */
public class JaasAcn {

    public static void main(String[] args) {

        // Obtain a LoginContext, needed for authentication. Tell it
        // to use the LoginModule implementation specified by the
        // entry named "JaasSample" in the JAAS login configuration
        // file and to also use the specified CallbackHandler.
        LoginContext lc = null;
        try {
            lc = new LoginContext("JaasSample", new TextCallbackHandler());
        } catch (LoginException le) {
            System.err.println("Cannot create LoginContext. "
                + le.getMessage());
            System.exit(-1);
        } catch (SecurityException se) {
            System.err.println("Cannot create LoginContext. "
                + se.getMessage());
            System.exit(-1);
        }

        try {

            // attempt authentication
            lc.login();

        } catch (LoginException le) {

            System.err.println("Authentication failed:");
            System.err.println("  " + le.getMessage());
            System.exit(-1);

        }

        System.out.println("Authentication succeeded!");

    }
}
```

jass.conf

```
/** Login Configuration for the JaasAcn and
 ** JaasAzn Applications
 **/

JaasSample {
   com.sun.security.auth.module.Krb5LoginModule required;
};
```

jassacn.policy

```
/** Java Access Control Policy for the JaasAcn Application **/


grant codebase "file:./JaasAcn.jar" {

   permission javax.security.auth.AuthPermission
"createLoginContext.JaasSample";
};
```

JaasAzn.java

```
import javax.security.auth.Subject;
import javax.security.auth.callback.*;
import javax.security.auth.login.*;
import com.sun.security.auth.callback.TextCallbackHandler;
import java.security.PrivilegedAction;

/**
 * This JaasAzn application attempts to authenticate a user
 * and reports whether or not the authentication was successful.
 * If successful, it then sets up subsequent execution of
 * code in the run method of the SampleAction class such that
 * access control checks for security-sensitive operations will be
 * based on the user running the code.
 */
public class JaasAzn {

    public static void main(String[] args) {

        // Obtain a LoginContext, needed for authentication. Tell it
        // to use the LoginModule implementation specified by the
        // entry named "JaasSample" in the JAAS login configuration
        // file and to also use the specified CallbackHandler.
        LoginContext lc = null;
        try {
            lc = new LoginContext("JaasSample", new TextCallbackHandler());
        } catch (LoginException le) {
            System.err.println("Cannot create LoginContext. "
                + le.getMessage());
            System.exit(-1);
        } catch (SecurityException se) {
```

```
            System.err.println("Cannot create LoginContext. "
                + se.getMessage());
            System.exit(-1);
        }

        try {

            // attempt authentication
            lc.login();

        } catch (LoginException le) {

            System.err.println("Authentication failed:");
            System.err.println("  " + le.getMessage());
            System.exit(-1);

        }

        System.out.println("Authentication succeeded!");

        // now try to execute the SampleAction as the authenticated Subject
        Subject mySubject = lc.getSubject();
        PrivilegedAction action = new SampleAction();
        Subject.doAsPrivileged(mySubject, action, null);

    }
}
```

SampleAction.java

```
import java.io.File;
import java.security.PrivilegedAction;

/**
 * This is a sample PrivilegedAction implementation, designed to be
 * used with the JaasAzn class.
 */
public class SampleAction implements PrivilegedAction {

    /**
     * This sample PrivilegedAction performs the following operations:
     * <ul>
     * <li> Access the System property <i>java.home</i>
     * <li> Access the System property <i>user.home</i>
     * <li> Access the file <i>foo.txt</i>
     * </ul>
     *
     * @return <code>null</code> in all cases.
     *
     * @exception SecurityException if the caller does not have permission
     *          to perform any of the operations listed above.
     */
    public Object run() {
        System.out.println("\nYour java.home property value is: "
```

```
                            +System.getProperty("java.home"));

        System.out.println("\nYour user.home property value is: "
                            +System.getProperty("user.home"));

        File f = new File("foo.txt");
        System.out.print("\nfoo.txt does ");
        if (!f.exists())
            System.out.print("not ");
        System.out.println("exist in the current working directory.");
        return null;
    }
}
```

jassazn.policy

```
/** Java Access Control Policy for the JaasAzn Application **/

/** Code-Based Access Control Policy for JaasAzn **/

grant codebase "file:./JaasAzn.jar" {

   permission javax.security.auth.AuthPermission
                    "createLoginContext.JaasSample";
   permission javax.security.auth.AuthPermission "doAsPrivileged";
};


/** User-Based Access Control Policy for the SampleAction class
 ** instantiated by JaasAzn
 **/

grant    codebase "file:./SampleAction.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm"  {

   permission java.util.PropertyPermission "java.home", "read";
   permission java.util.PropertyPermission "user.home", "read";
   permission java.io.FilePermission "foo.txt", "read";
};
```

Login.java

```
import java.io.*;
import java.lang.reflect.*;
import java.util.Arrays;

import javax.security.auth.callback.*;
import javax.security.auth.login.*;
import javax.security.auth.Subject;
import com.sun.security.auth.callback.TextCallbackHandler;

/**
 * <p> This class authenticates a <code>Subject</code> and then
```

```
 * executes a specified application as that <code>Subject</code>.
 * To use this class, the java interpreter would typically be invoked as:
 *
 * <pre>
 *    % java -Djava.security.manager \
 *        Login \
 *        <applicationclass> <applicationClass_args>
 * </pre>
 *
 * <p> <i>applicationClass</i> represents the application to be executed
 * as the authenticated <code>Subject</code>,
 * and <i>applicationClass_args</i> are passed as arguments to
 * <i>applicationClass</i>.
 *
 * <p> To perform the authentication, <code>Login</code> uses a
 * <code>LoginContext</code>.  A <code>LoginContext</code> relies on a
 * <code>Configuration</code> to determine the modules that should be used
 * to perform the actual authentication.  The location of the Configuration
 * is dependent upon each Configuration implementation.
 * The default Configuration implementation
 * (<code>com.sun.security.auth.login.ConfigFile</code>)
 * allows the Configuration location to be specified (among other ways)
 * via the <code>java.security.auth.login.config</code> system property.
 * Therefore, the <code>Login</code> class can also be invoked as:
 *
 * <pre>
 *    % java -Djava.security.manager \
 *        -Djava.security.auth.login.config=<configuration_url> \
 *        Login \
 *        <your_application_class> <your_application_class_args>
 * </pre>
 */

public class Login {

    /**
     * <p> Instantate a <code>LoginContext</code> using the
     * provided application classname as the index for the login
     * <code>Configuration</code>.  Authenticate the <code>Subject</code>
     * (three retries are allowed) and invoke
     * <code>Subject.doAsPrivileged</code>
     * with the authenticated <code>Subject</code> and a
     * <code>PrivilegedExceptionAction</code>.
     * The <code>PrivilegedExceptionAction</code>
     * loads the provided application class, and then invokes
     * its public static <code>main</code> method, passing it
     * the application arguments.
     *
     * <p>
     *
     * @param args the arguments for <code>Login</code>.  The first
     *          argument must be the class name of the application to be
     *          invoked once authentication has completed, and the
     *          subsequent arguments are the arguments to be passed
     *          to that application's public static <code>main</code>
```

```
method.
     */
    public static void main(String[] args) {

        // check for the application's main class
        if (args == null || args.length == 0) {
            System.err.println("Invalid arguments: " +
                "Did not provide name of application class.");
            System.exit(-1);
        }

        LoginContext lc = null;
        try {
            lc = new LoginContext(args[0], new TextCallbackHandler());
        } catch (LoginException le) {
            System.err.println("Cannot create LoginContext. "
                + le.getMessage());
            System.exit(-1);
        } catch (SecurityException se) {
            System.err .println("Cannot create LoginContext. "
                + se.getMessage());
            System.exit(-1);
        }

        // the user has 3 attempts to authenticate successfully
        int i;
        for (i = 0; i < 3; i++) {
            try {

                // attempt authentication
                lc.login();

                // if we return with no exception, authentication succeeded
                break;

            } catch (AccountExpiredException aee) {

                System.err.println("Your account has expired.  " +
                                "Please notify your administrator.");
                System.exit(-1);

            } catch (CredentialExpiredException cee) {

                System.err.println("Your credentials have expired.");
                System.exit(-1);

            } catch (FailedLoginException fle) {

                System.err.println("Authentication Failed");
                try {
                      Thread.currentThread().sleep(3000);
                } catch (Exception e) {
                      // ignore
                }
```

```
            } catch (Exception e) {

                System.err.println("Unexpected Exception - unable to
continue");
                e.printStackTrace();
                System.exit(-1);
            }
        }

        // did they fail three times?
        if (i == 3) {
            System.err.println("Sorry");
            System.exit(-1);
        }

        // push the subject into the current ACC
        try {
            Subject.doAsPrivileged(lc.getSubject(),
                                   new  MyAction(args),
                                   null);
        } catch (java.security.PrivilegedActionException pae) {
            pae.printStackTrace();
            System.exit(-1);
        }

        System.exit(0);
    }
}

class MyAction implements java.security.PrivilegedExceptionAction {

    String[] origArgs;

    public MyAction(String[] origArgs) {
        this.origArgs = (String[])origArgs.clone();
    }

    public Object run() throws Exception {

        // get the ContextClassLoader
        ClassLoader cl = Thread.currentThread().getContextClassLoader();

        try {
            // get the application class's main method
            Class c = Class.forName(origArgs[0], true, cl);
            Class[] PARAMS = { origArgs.getClass() };
            java.lang.reflect.Method mainMethod = c.getMethod("main",
PARAMS);

            // invoke the main method with the remaining args
            String[] appArgs = new String[origArgs.length - 1];
            System.arraycopy(origArgs, 1, appArgs, 0, origArgs.length - 1);
            Object[] args = { appArgs };
            mainMethod.invoke(null /*ignored*/, args);
        } catch (Exception e) {
```

```
            throw new java.security.PrivilegedActionException(e);
        }

        // successful completion
        return null;
    }
}
```

Sample.java

```
import java.io.File;

public class Sample {

    /**
     * This sample class performs the following operations:
     * <ul>
     * <li> Access the System property <i>java.home</i>
     * <li> Access the System property <i>user.home</i>
     * <li> Access the file <i>foo.txt</i>
     * </ul>
     *
     * @exception SecurityException if the caller does not have permission
     *          to perform any of the operations listed above.
     */
    public static void main (String[] args) throws SecurityException {

        // If there were any arguments to read, we'd do it here.

        System.out.println("\nYour java.home property value is: "
                           +System.getProperty("java.home"));

        System.out.println("\nYour user.home property value is: "
                           +System.getProperty("user.home"));

        File f = new File("foo.txt");
        System.out.print("\nfoo.txt does ");
        if (!f.exists())
            System.out.print("not ");
        System.out.println("exist in the current working directory.");
    }
}
```

sample.conf

```
/** Login Configuration for the Sample Application **/

Sample {
   com.sun.security.auth.module.Krb5LoginModule required;
};
```

sample.policy

```
/** Access Control Policy for the Sample Application **/


grant codebase "file:./Login.jar" {
   permission java.security.AllPermission;
};


grant    codebase "file:./Sample.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm"  {

   permission java.util.PropertyPermission "java.home", "read";
   permission java.util.PropertyPermission "user.home", "read";
   permission java.io.FilePermission "foo.txt", "read";
};
```

csLogin.conf

```
/**
 * Login Configuration for JAAS.
 */

SampleClient {
  com.sun.security.auth.module.Krb5LoginModule required;
};

SampleServer {
  com.sun.security.auth.module.Krb5LoginModule required storeKey=true
principal="service_principal@your_realm";
};
```

client.policy

```
grant CodeBase "file:./Login.jar" {
        permission java.security.AllPermission;
};

grant CodeBase "file:./SampleClient.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm" {

    permission java.net.SocketPermission "*", "connect";

    permission javax.security.auth.kerberos.ServicePermission
        "krbtgt/your_realm@your_realm",
        "initiate";

    permission javax.security.auth.kerberos.ServicePermission
        "service_principal@your_realm",
```

```
        "initiate";
};
```

server.policy

```
grant CodeBase "file:./Login.jar" {
        permission java.security.AllPermission;
};

grant CodeBase "file:./SampleServer.jar"
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "service_principal@your_realm" {

    permission java.net.SocketPermission "*", "accept";

    permission javax.security.auth.kerberos.ServicePermission
        "service_principal@your_realm", "accept";
};
```

SampleServerImp.java

```
import org.ietf.jgss.*;
import java.io.*;
import java.net.Socket;
import java.net.ServerSocket;
import javax.security.auth.Subject;
import java.security.PrivilegedAction;

/**
 * A sample server application that uses JGSS to do mutual authentication
 * with a client using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the client.
 *
 * Every message exchanged with the client includes a 4-byte application-
 * level header that contains the big-endian integer value for the number
 * of bytes that will follow as part of the JGSS token.
 *
 * The protocol is:
 *    1.  Context establishment loop:
 *         a. client sends init sec context token to server
 *         b. server sends accept sec context token to client
 *         ....
 *    2. client sends a wrap token to the server.
 *    3. server sends a mic token to the client for the application
 *       message that was contained in the wrap token.
 */

public class SampleServerImp  {

    public static void main(String[] args)
        throws IOException, GSSException {

        // Obtain the command-line arguments and parse the port number
```

```
        if (args.length != 1) {
            System.err.println("Usage: java <options> Login SampleServer
<localPort>");
            System.exit(-1);
        }

        int localPort = Integer.parseInt(args[0]);

        ServerSocket ss = new ServerSocket(localPort);

        GSSManager manager = GSSManager.getInstance();

        while (true) {

            System.out.println("Waiting for incoming connection...");

            Socket socket = ss.accept();
            DataInputStream inStream =
                new DataInputStream(socket.getInputStream());
            DataOutputStream outStream =
                new DataOutputStream(socket.getOutputStream());

            System.out.println("Got connection from client "
                               + socket.getInetAddress());

            /*
             * Create a GSSContext to receive the incoming request
             * from the client. Use null for the server credentials
             * passed in. This tells the underlying mechanism
             * to use whatever credentials it has available that
             * can be used to accept this connection.
             */
            GSSContext context =
manager.createContext((GSSCredential)null);

            // Do the context eastablishment loop

            byte[] token = null;

            while (!context.isEstablished()) {

                token = new byte[inStream.readInt()];
                System.out.println("Will read input token of size "
                                   + token.length
                                   + " for processing by
acceptSecContext");
                inStream.readFully(token);

                token = context.acceptSecContext(token, 0, token.length);

                // Send a token to the peer if one was generated by
                // acceptSecContext
                if (token != null) {
                    System.out.println("Will send token of size "
                                       + token.length
```

```
                            + " from acceptSecContext.");
        outStream.writeInt(token.length);
        outStream.write(token);
        outStream.flush();
    }
}

System.out.println("Context Established! ");
System.out.println("Client is " + context.getSrcName());
System.out.println("Server is " + context.getTargName());
/*
 * If mutual authentication did not take place, then
 * only the client was authenticated to the
 * server. Otherwise, both client and server were
 * authenticated to each other.
 */
if (context.getMutualAuthState())
    System.out.println("Mutual authentication took place!");

/*
 * Create a MessageProp which unwrap will use to return
 * information such as the Quality-of-Protection that was
 * applied to the wrapped token, whether or not it was
 * encrypted, etc. Since the initial MessageProp values
 * are ignored, just set them to the defaults of 0 and false.
 */
MessageProp prop = new MessageProp(0, false);

/*
 * Read the token. This uses the same token byte array
 * as that used during context establishment.
 */
token = new byte[inStream.readInt()];
System.out.println("Will read token of size "
                   + token.length);
inStream.readFully(token);

byte[] bytes = context.unwrap(token, 0, token.length, prop);
String str = new String(bytes);
System.out.println("Received data \""
                   + str + "\" of length " + str.length());

System.out.println("Confidentiality applied: "
                   + prop.getPrivacy());

/*
 * Now generate a MIC and send it to the client. This is
 * just for illustration purposes. The integrity of the
 * incoming wrapped message is guaranteed irrespective of
 * the confidentiality (encryption) that was used.
 */

/*
 * First reset the QOP of the MessageProp to 0
 * to ensure the default Quality-of-Protection
```

```
             * is applied.
             */
            prop.setQOP(0);

            token = context.getMIC(bytes, 0, bytes.length, prop);

            System.out.println("Will send MIC token of size "
                               + token.length);
            outStream.writeInt(token.length);
            outStream.write(token);
            outStream.flush();


            /*
             * Impersonate client
             */

            System.out.println("Impersonating client.");

            /*
             * Extract the KerberosPrincipal from the client GSSName and
populate
             * it in the principal set of a new Subject. Pass in a null for
             * credentials. If we were to pass in the delegated
GSSCredential
             * instead of null, then the resulting Subject's private
credential
             * set would also be populated.
             */
            GSSName clientGSSName = context.getSrcName();
            System.out.println("clientGSSName: " + clientGSSName);
            Subject client =
                com.sun.security.jgss.GSSUtil.createSubject(clientGSSName,
                                                            null);

            /*
             * Construct an action that will read a file meant only for the
             * client
             */
            PrivilegedAction readFile =
                new ReadFileAction(clientGSSName.toString());

            /*
             * Invoke the action via a doAsPrivileged. This allows the
             * action to be executed as the client subject, and it also
runs
             * that code as privileged. This means that any permission
checking
             * that happens beyond this point applies only to the code
being
             * run as the client.
             */
            Subject.doAsPrivileged(client, readFile, null);
```

```
            /*
             * Clean up
             */


            System.out.println("Closing connection with client "
                               + socket.getInetAddress());
            context.dispose();
            socket.close();
        }
    }
}
```

ReadFileAction.java

```
import java.security.PrivilegedAction;
import java.io.*;

/**
 * This class implements the PrivilegedAction interface to demonstrate the
 * reading of a file that belongs to the client. This code will be
 * executed by the server while impersonating the client principal.
 */
public class ReadFileAction implements PrivilegedAction {

    private String fileName;

    /**
     * Contructs a ReadFileAction instance.
     *
     * @param kerberosPrincipalName the name of the Kerberos principal
     * who owns the file that will be read. The filename is constructed
     * from the name of the principal.
     */
    public ReadFileAction(String kerberosPrincipalName) {
        /*
         * Separate the realm component from the name and use the rest of
         * it for constructing the filename. If the principal name is
         * "joe@REALM" then the file that will be read is
         * "data/joe_info.txt". The path separator "/" might be "\" in the
         * case of Windows.
         */
        int realmSeparatorPos = kerberosPrincipalName.lastIndexOf('@');
        fileName = "data" + File.separatorChar
            + kerberosPrincipalName.substring(0, realmSeparatorPos)
            + "_info.txt";
    }

    /**
     * Does the actual reading of the file. It displays the text contained
     * in the file.
     */
    public Object run() {
```

```
System.out.println("==============================================");
        System.out.println("Reading file: " + fileName);
        try {
            BufferedReader reader = new BufferedReader(new
FileReader(fileName));
            String str = reader.readLine();
            while (str != null) {
                System.out.println(str);
                str = reader.readLine();
            }
        } catch (IOException e) {
            System.err.println(e);
        }

System.out.println("==============================================");
        return null;
    }
}
```

serverimp.policy

```
grant CodeBase "file:./Login.jar" {
        permission java.security.AllPermission;
};

grant CodeBase "file:./SampleServerImp.jar"
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "service_principal@your_realm" {

    permission java.net.SocketPermission "*", "accept";

    permission javax.security.auth.kerberos.ServicePermission
        "service_principal@your_realm", "accept";

    permission javax.security.auth.AuthPermission "doAsPrivileged";
};

grant CodeBase "file:./ReadFileAction.jar"
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "your_user_name@your_realm" {

    permission java.io.FilePermission
        "data/your_user_name_info.txt", "read";
};
```

csImpLogin.conf

```
/**
 * Login Configuration for JAAS.
 */

SampleClient {
  com.sun.security.auth.module.Krb5LoginModule required;
```

```
};

SampleServerImp {
  com.sun.security.auth.module.Krb5LoginModule required storeKey=true
principal="service_principal@your_realm";
};
```

Related Documentation

- API specifications
  - `com.sun.security.jgss` package
  - `com.sun.security.auth` package
  - `com.sun.security.auth.callback` package
  - `com.sun.security.auth.login` package
  - `com.sun.security.auth.module` package
- User guides and tutorials
  - Java Authentication and Authorization Service (JAAS) Reference Guide
  - Java Security Tutorial
- Other Java Security Documentation
  - Default Policy Implementation and Policy File Syntax
  - Permissions in the JDK
  - Single Sign-on Using Kerberos in Java
  - Java SE Platform Security Architecture
- Reference document
  - Generic Security Service API Version 2: Java Bindings Update

Accessing Native GSS-API

To help Java platform applications achieve seamless integration with native applications, the JDK enhances the Java GSS-API to use native GSS-API instead of its own implementation of cryptographic mechanisms when configured to do so. When using the native GSS-API and its underlying native cryptographic mechanisms, the native credentials and settings in users' environment will be picked up automatically. This is different from the default case in which the Java GSS-API uses its own implementation of cryptographic mechanisms. When using Kerberos, Java applications have to supply Kerberos configuration information using the designated Kerberos system properties for the Java GSS-API to function. Introduction to JAAS and Java GSS-API Tutorials covers the default case in great detail, so this section will focus on how to enable or configure Java GSS-API to use native GSS-API.

Before you enable Java GSS-API to use native GSS-API, ensure that native GSS-API and its underlying cryptographic mechanism are available and functioning with user settings. For example, ensure that native GSS libraries are installed at the appropriate directories with proper configurations, and the same applies to the Kerberos library and configurations. Note that native GSS-API assumes that before an application calls

its APIs, it has already obtained and stored the mechanism-specific credentials in a location that the native mechanism implementation is aware of. Thus, when an application uses native GSS-API with Kerberos, it must already have obtained the appropriate native credentials, such as Kerberos tickets and keys by using the `kinit` tool on the initiator side or a keytab file on the acceptor side.

To make the Java GSS-API use native GSS-API, Java applications must explicitly enable this behavior by setting one or more of the following system properties:

- `sun.security.jgss.native` (required): Set this to true to enable the Java GSS-API to use native GSS-API.
- `sun.security.jgss.lib` (optional, but required for Windows): Set to the full path of the native GSS library. If this is not set, then the Java GSS-API will look for the native GSS library using the default Java library path. If this is not set, then the Java GSS-API will try to search for some well-known native GSS libraries, for example, `libgss.so` on Solaris, `libgssapi.so` or `libgssapi_krb5.so` on Linux, and `libgssapi_krb5.dylib` on macOS. There is no well-known native GSS-API library on Windows, and you must specify its full path name with this system property.

As mentioned previously, native GSS-API requires that the application had obtained these credentials and that they are accessible. Java applications can access these native credentials through the Java GSS-API and use them for establishing GSS-API security contexts with peers. Note that when a Subject is present, for example,

```
javax.security.auth.Subject.getSubject(AccessController.getContext()) != null
```

then the Java GSS-API mandates that the credentials be obtained from the private or public credential sets of the current Subject and that the Java GSS-API call must fail if the desired credential cannot be found. Thus, Java platform applications that execute the Java GSS-API calls inside a `Subject.doAs/doAsPrivileged(...)` call should either populate the Subject's credential sets with the appropriate Java `GSSCredential` objects that encapsulate the native credentials or explicitly set the system property `javax.security.auth.useSubjectCredsOnly` to false so that the Java GSS-API can obtain credentials from other locations, for example, from native credential caches, in addition to the Subject's credential sets.

When delegated to establish a GSS-API security context on behalf of others, Java applications can either specify the delegated credential, as returned by `GSSContext.getDelegCred()`, explicitly in Java GSS-API calls, or create a `Subject` object with this delegated credential and execute the Java GSS-API calls inside the `Subject.doAs/doAsPrivileged(...)` calls.

Once the native GSS-API is enabled, Java platform applications that indirectly call Java GSS-API through mechanisms or protocols such as Simple Authentication and Security Layer (SASL) (see Java SASL API Programming and Deployment Guide) will also use user's native settings and credentials.

Here is some sample code that helps demonstrate how to use Java GSS-API to establish GSS-API security contexts and securely exchange data between three parties: `SampleClient` contacts `FooServer`, which in turn contacts `FooServer2` on behalf of `SampleClient`. Note:

- The sample code should be invoked with native GSS-API enabled. The Principal names `host@foo.sample.com` and `host@foo2.sample.com` are placeholders and should be replaced with actual principal names in your Kerberos database.

- When a security manager is installed, some Java GSS-API calls require that permissions be granted. Check the Java documentation of the following classes for more details:
  - `javax.security.auth.kerberos.ServicePermission`
  - `javax.security.auth.kerberos.DelegationPermission`
- To simplify the example, token exchanges between peers are represented by two pseudo-methods: `SEND_TOKEN(byte[])` and `READ_TOKEN()`. Their actual implementation are application-specific and thus not shown here.
- To reduce code duplication, context establishment code is referred by a pseudo-method, `ESTABLISH_CONTEXT(GSSContext)`, in the code segments for `SampleClient`, `FooServer`, and `FooServer2`.

The following is the implementation for `ESTABLISH_CONTEXT(GSSContext)` using Java GSS-API.

```
/**
 * ESTABLISH_CONTEXT(GSSContext ctxt): establishes a context
 * with data confidentiality and mutual authentication.
 */
ctxt.requestConf(true);
ctxt.requestMutualAuth(true);

byte[] inToken = new byte[0];
byte[] outToken = null;

if (ctxt.isInitiator()) {
    while (!ctxt.isEstablished()) {
        // Note: initSecContext(...) always ignores the arguments
        // for the first call because there is no incoming token.
        outToken = ctxt.initSecContext(inToken, 0, inToken.length);

        // Send the output token if generated.
        if (outToken != null) SEND_TOKEN(outToken); // to acceptor

        // Check whether more incoming tokens are expected.
        if (!ctxt.isEstablished()) {
            inToken = READ_TOKEN(); // from acceptor
        }
    }
} else {
    while (!ctxt.isEstablished()) {
        inToken = READ_TOKEN(); // from initiator
        outToken =
            ctxt.acceptSecContext(inToken, 0, inToken.length);

        // Send the output token if generated.
        if (outToken != null) SEND_TOKEN(outToken); // to initiator
    }
}
```

Following are the code segments for `SampleClient`, `FooServer`, and `FooServer2`:

`SampleClient`: It contacts `FooServer` and delegates the server to act on its behalf. If all goes well, it should get back a personalized hello message produced by `FooServer2`.

```
GSSManager gssMgr = GSSManager.getInstance();
GSSName serverName = gssMgr.createName(
    "host@foo.sample.com", GSSName.NT_HOSTBASED_SERVICE);
```

```
GSSContext context = gssMgr.createContext(
    serverName, null /* default mechanism, which is Kerberos*/,
    null /* default initiator cred */,
    GSSContext.DEFAULT_LIFETIME);
context.requestCredDelegState(true);

ESTABLISH_CONTEXT(context);

// Make sure credential delegation is available.
if (!context.getCredDeleg()) {
    context.dispose();
    throw new Exception("credential delegation is denied");
}

byte[] token = READ_TOKEN(); // from "FooServer"
byte[] data =
   context.unwrap(token, 0, token.length, new MessageProp(true));
context.dispose();

// Should print "Hello from FooServer2 to <client name>" where
// <client name> is the name of the default initiator.
System.out.println(new String(data));
```

FooServer: It contacts `FooServer2` as `SampleClient` and forwards the received reply to `SampleClient`.

```
GSSManager gssMgr = GSSManager.getInstance();
GSSName myName = gssMgr.createName(
    "host@foo.sample.com", GSSName.NT_HOSTBASED_SERVICE);
GSSCredential myCred = gssMgr.createCredential(
    acceptorName, GSSCredential.INDEFINITE_LIFETIME,
    (Oid[]) null /* default set of mechanisms */,
    GSSCredential.ACCEPT_ONLY);
GSSContext acontext = gssMgr.createContext(myCred);

ESTABLISH_ACC_CONTEXT(acontext);

GSSCredential delegCred = acontext.getDelegCred();
if (delegCred != null) {
    byte[] data, token;
    // Establish a context on client's behalf using the delegated
    // credential.
    GSSName serverName = gssMgr.createName(
        "host@foo2.sample.com", GSSName.NT_HOSTBASED_SERVICE);
    GSSContext icontext = gssMgr.createContext(
        serverName, null /* default mechanism Kerberos */,
        delegCred /* act on SampleClient's behalf */,
        GSSContext.DEFAULT_LIFETIME);

    ESTABLISH_CONTEXT(icontext);

    token = READ_TOKEN(); // from "FooServer2"

    MessageProp msgProp = new MessageProp(true);

    // Forward the reply from FooServer2 to SampleClient.
    data = icontext.unwrap(token, 0, token.length, msgProp);
    token = acontext.wrap(data, 0, data.length, msgProp);
    SEND_TOKEN(token); // to "SampleClient"
    icontext.dispose();
```

```
}
acontext.dispose();
```

FooServer2: It always replies with a hello message personalized to the name of the initiator of the established context.

```
GSSManager gssMgr = GSSManager.getInstance();
GSSName myName = gssMgr.createName(
    "host@foo2.sample.com", GSSName.NT_HOSTBASED_SERVICE);
GSSCredential myCred = gssMgr.createCredential(
    myName, GSSCredential.INDEFINITE_LIFETIME,
    (Oid[]) null /* default set of mechanisms */,
    GSSCredential.ACCEPT_ONLY);
GSSContext context = gssMgr.createContext(myCred);

ESTABLISH_CONTEXT(context);

byte[] data = new String("Hello from FooServer2 to " +
    context.getSrcName()).getBytes();
byte[] token =
    context.wrap(data, 0, data.length, new MessageProp(true));

SEND_TOKEN(token); // to "FooServer"

context.dispose();
```

Single Sign-on Using Kerberos in Java

Abstract

A significant enhancement to the Java SE security architecture is the capability to achieve single sign-on using Kerberos Version 5. A single sign-on solution lets users authenticate themselves just once to access information on any of several systems. This is done using JAAS for authentication and authorization and Java GSS-API to establish a secure context for communication with a peer application. Our focus is on Kerberos V5 as the underlying security mechanism for single sign-on, although other security mechanisms may be added in the future.

Introduction

With the increasing use of distributed systems users need to access resources that are often remote. Traditionally users have had to sign-on to multiple systems, each of which may involve different user names and authentication techniques. In contrast, with single sign-on, the user needs to authenticate only once and the authenticated identity is securely carried across the network to access resources on behalf of the user.

In this paper we discuss how to use single sign-on based on the Kerberos V5 protocol. We use the Java Authentication and Authorization Service (JAAS) to authenticate a principal to Kerberos and obtain credentials that prove its identity. We show how Oracle's implementation of a Kerberos login module can be made to read credentials from an existing cache on platforms that contain native Kerberos support. We then use the Java Generic Security Service API (Java GSS-API) to authenticate to a remote peer using the previously obtained Kerberos credentials. We also show how to delegate Kerberos credentials for single sign-on in a multi-tier environment.

Kerberos V5

Kerberos V5 is a trusted third party network authentication protocol designed to provide strong authentication using secret key cryptography. When using Kerberos V5, the user's password is never sent across the network, not even in encrypted form, except during Kerberos V5 administration. Kerberos was developed in the mid-1980's as part of MIT's Project Athena. A full description of the Kerberos V5 protocol is beyond the scope of this paper. For more information on the Kerberos V5 protocol please refer to [1] and [2].

Kerberos V5 is a mature protocol and has been widely deployed. A free reference implementation in C is available from MIT. For these reasons we have selected Kerberos V5 as the underlying technology for single sign-on in Java SE.

Java Authentication and Authorization Service (JAAS)

The Java SE security architecture used to solely determine privileges by the origin of the code and the public key certificates matching the code signers. However, in a multi-user environment it is desirable to further specify privileges based on the authenticated identity of the user running the code.

JAAS supplies such a capability. JAAS is a pluggable framework and programming interface specifically targeted for authentication and access control based on the authenticated identities.

Pluggable and Stackable Framework

JAAS authentication framework is based on Pluggable Authentication Module (PAM); see [3] and [4]. JAAS authentication is performed in a pluggable fashion allowing system administrators to add appropriate authentication modules. This permits Java applications to remain independent of underlying authentication technologies, and new or updated authentication technologies can be seamlessly configured without requiring modifications to the application itself.

JAAS authentication framework also supports the stacking of authentication modules. Multiple modules can be specified and they are invoked by the JAAS framework in the order they were specified. The success of the overall authentication depends on the results of the individual authentication modules.

Authentication and Authorization

The JAAS framework can be divided into two components: an authentication component and an authorization component.

The JAAS authentication component provides the ability to reliably and securely determine who is currently executing Java code, regardless of whether the code is running as an application, an applet, a bean, or a servlet.

The JAAS authorization component supplements the existing Java security framework by providing the means to restrict the executing Java code from performing sensitive tasks, depending on its codesource and depending on who is executing the code.

Subject

JAAS uses the term Subject to refer to any entity that is the source of a request to access resources. A Subject may be a user or a service. Since an entity may have many names or principals JAAS uses Subject as an extra layer of abstraction that handles multiple names per entity. Thus a Subject is comprised of a set of principals. There are no restrictions on principal names.

A Subject is only populated with authenticated principals. Authentication typically involves the user providing proof of identity, such as a password.

A Subject may also have security related attributes, which are referred to as credentials. The credentials can be public or private. Sensitive credentials such as private cryptographic keys are stored in the private credentials set of the Subject.

The `Subject` class has methods to retrieve the principals, public credentials and private credentials associated with it.

Please note that different permissions may be required for operations on these classes. For example `AuthPermission("modifyPrincipals")` may be required to modify the principal set of the Subject. Similar permissions are required to modify the public credentials, private credentials and to get the current Subject.

doAs and doAsPrivileged

Java SE enforces runtime access controls via `java.lang.SecurityManager`. The `SecurityManager` is consulted anytime sensitive operations are attempted. The `SecurityManager` delegates this responsibility to `java.security.AccessController`. The `AccessController` obtains a current image of the `AccessControlContext` and verifies that it has sufficient permission to do the operation requested.

JAAS provides two methods, `doAs` and `doAsPrivileged`, that can be used to associate an authenticated Subject with the `AccessControlContext` dynamically.

The `doAs` method associates the Subject with the current thread's access control context and subsequent access control checks are made on the basis of the code being executed and the Subject executing it.

```
public static Object doAs(final Subject subject,
                          final PrivilegedAction action)

public static Object doAs(final Subject subject,
                          final PrivilegedExceptionAction action)
    throws  PrivilegedActionException;
```

Both forms of the `doAs` method first associate the specified subject with the current Thread's `AccessControlContext`, and then execute the action. This achieves the effect of having the action run as the Subject. The first method can throw runtime exceptions but normal execution has it returning an `Object` from the `run()` method of its action argument. The second method behaves similarly except that it can throw a checked `PrivilegedActionException` from its `run()` method. An `AuthPermission("doAs")` is required to call the `doAs` methods.

The following methods also execute code as a particular Subject:

```
public static Object doAsPrivileged(final Subject subject,
                                    final PrivilegedAction action,
                                    final AccessControlContext  acc);

public static Object doAsPrivileged(final Subject subject,
                                    final PrivilegedExceptionAction action,
                                    final AccessControlContext acc)
    throws PrivilegedActionException;
```

The `doAsPrivileged` method behaves exactly as `doAs`, except that it allows the caller to specify an access control context. Thus it effectively throws away the current `AccessControlContext` and authorization decisions will be based on the `AccessControlContext` passed in.

Since the `AccessControlContext` is set on a per thread basis, different threads within the JVM can assume different identities. The Subject associated with a specific `AccessControlContext` can be retrieved by using the following method:

```
public static Subject getSubject(final AccessControlContext acc);
```

LoginContext

The `LoginContext` class provides the basic methods used to authenticate Subjects. It also allows an application to be independent of the underlying authentication technologies. The `LoginContext` consults a configuration that determines the authentication services or `LoginModules` configured for a particular application. If the application does not have a specific entry, it defaults to the entry identified as "other".

To support the stackable nature of `LoginModules`, `LoginContext` performs authentication in two phases. In the first phase or login phase, it invokes each configured LoginModule to attempt the authentication. If all the necessary `LoginModules` succeed, then `LoginContext` enters the second phase where it invokes each LoginModule again to formally commit the authentication process. During this phase the Subject is populated with the authenticated principals and their credentials. If either of the phase fails, then the `LoginContext` invokes each configured module to abort the entire authentication attempt. Each `LoginModule` then cleans up any relevant state associated with the authentication attempt.

`LoginContext` has four constructors that can be used to instantiate it. All of them require the configuration entry name to be passed. In addition the Subject and/or a `CallbackHandler` can also be passed to the constructors.

Callbacks

The login modules invoked by JAAS must be able to garner information from the caller for authentication. For example the Kerberos login module may require users to enter their Kerberos password for authentication.

The `LoginContext` allows the application to specify a callback handler that the underlying login modules use to interact with users. There are two callback handlers - one based on the command line and another based on a GUI.

LoginModules

Oracle provides an implementation of the `UnixLoginModule`, `NTLoginModule`, `JNDILoginModule`, `KeyStoreLoginModule` and `Krb5LoginModule`.

The Kerberos Login Module

The class `com.sun.security.auth.module.Krb5LoginModule` is Oracle's implementation of a login module for the Kerberos version 5 protocol. Upon successful authentication the Ticket Granting Ticket (TGT) is stored in the Subject's private credentials set and the Kerberos principal is stored in the Subject's principal set.

Based on certain configurable options, `Krb5LoginModule` can also use an existing credentials cache, such as a native cache in the operating system, to acquire the TGT and/or use a `keytab` file containing the secret key to implicitly authenticate a principal. Both Solaris and Windows contain a credentials cache that `Krb5LoginModule` can use for fetching the TGT. Solaris also contains a system-wide `keytab` file that `Krb5LoginModule` can use for fetching the secret key. On all platforms, `Krb5LoginModule` supports options to set the file path to a ticket cache or `keytab` file of choice. This is useful when third-party Kerberos support is installed and Java integration is desired. Please consult the documentation for `Krb5LoginModule` to learn about these options. In the absence of a native cache or `keytab`, the user will be prompted for the password and the TGT obtained from the key distribution center (KDC).

The following is a sample JAAS login configuration entry for a client application. In this example, `Krb5LoginModule` will use the native ticket cache to get the TGT available in it. The authenticated identity will be the identity of the Kerberos principal that the TGT belongs to.

```
// Sample client configuration entry

SampleClient {
    com.sun.security.auth.module.Krb5LoginModule required
useTicketCache=true
};
```

The following is a sample login configuration entry for a server application. With this configuration, the secret key from the keytab is used to authenticate the principal `nfs/bar.example.com` and both the TGT obtained from the Kerberos KDC and the secret key are stored in the Subject's private credentials set. The stored key may be used later to validate a service ticket sent by a client (See the section on Java GSS-API.)

```
// Sample server configuration entry

SampleServer {
     com.sun.security.auth.module.Krb5LoginModule
         required useKeyTab=true storeKey=true principal="nfs/
bar.example.com"
};
```

In the following client code example, the configuration entry `SampleClient` will be used by the `LoginContext`. The `TextCallbackHandler` class will be used to prompt the user for the Kerberos password. Once the user has logged in, the Subject will be populated with the Kerberos Principal name and the TGT. Thereafter the user can execute code using `Subject.doAs` passing in the Subject obtained from the `LoginContext`.

```
// Sample client code

LoginContext lc = null;

try {
        lc = new LoginContext("SampleClient", new TextCallbackHandler());
        // attempt authentication
        lc.login();
} catch (LoginException le) {
    ...
}

// Now try to execute ClientAction as the authenticated Subject

Subject mySubject = lc.getSubject();
PrivilegedAction action = new ClientAction();
Subject.doAs(mySubject, action);
```

`ClientAction` could be an action that is allowed only for authenticated Kerberos client Principals with a specific value.

The following shows server side sample code. It is similar to the sample client code except for the application entry name and the `PrivilegedAction`.

```
// Sample server code

LoginContext lc = null;

try {
        lc = new LoginContext("SampleServer", new TextCallbackHandler());
        // attempt authentication
        lc.login();
} catch (LoginException le) {
   ...
}

// Now try to execute ServerAction as the authenticated Subject

Subject mySubject = lc.getSubject();
PrivilegedAction action = new ServerAction();
Subject.doAs(mySubject, action);
```

Kerberos Classes

To enable other vendors to provide their own Kerberos login module implementation that can be used with Java GSS-API, three standard Kerberos classes have been introduced in the `javax.security.auth.kerberos` package. These are

`KerberosPrincipal` for Kerberos principals, `KerberosKey` for the long-term Kerberos secret key and `KerberosTicket` for Kerberos tickets. All implementations of the Kerberos login module must use these classes to store principals, keys and tickets in the Subject.

Authorization

Upon successful authentication of a Subject, access controls can be enforced based upon the principals associated with the authenticated Subject. The JAAS principal based access controls augment the `CodeSource` access controls of Java SE. Permissions granted to a Subject are configured in `Policy`, which is an abstract class for representing the system wide access control policy. Oracle provides a file based implementation of the `Policy` class. The `Policy` class is provider based so that others can provide their own policy implementation.

Java Generic Security Service Application Program Interface (Java GSS-API)

Generic Security Service API (GSS-API)

Enterprise applications often have varying security requirements and deploy a range of underlying technologies to achieve this. In such a scenario how do we develop a client-server application so that it can easily migrate from one technology to another? The GSS-API was designed in the Common Authentication Technology working group of the IETF to solve this problem by providing a uniform application programming interface for peer to peer authentication and secure communication that insulates the caller from the details of the underlying technology.

The API, described in a language independent form in RFC 2743 [6], accommodates the following security services: authentication, message confidentiality and integrity, sequencing of protected messages, replay detection, and credential delegation. The underlying security technology or "security mechanism" being used, has a choice of supporting one or more of these features beyond the essential one way authentication.[1]

There are mainly two standard security mechanisms that the IETF has defined: Kerberos V5 [6] and the Simple Public Key Mechanism (SPKM) [8].

The API is designed such that an implementation may support multiple mechanisms simultaneously, giving the application the ability to choose one at runtime. However, a client application and a server application that communicate with each other must use the same security mechanism. Figure 7-1 illustrates this. It shows a client-server application that uses the GSS-API for secure communication. The GSS framework enables this application to support multiple security mechanisms (in this example, Kerberos V5 and SPKM). Once the GSS-API negotiates a security mechanism for the client or server application (in this example, Kerberos V5) the other must use the same.

[1] The GSS-API Kerberos mechanism performs client authentication at the minimum.

Figure 7-1 A Multi-Mechanism GSS-API Implementation

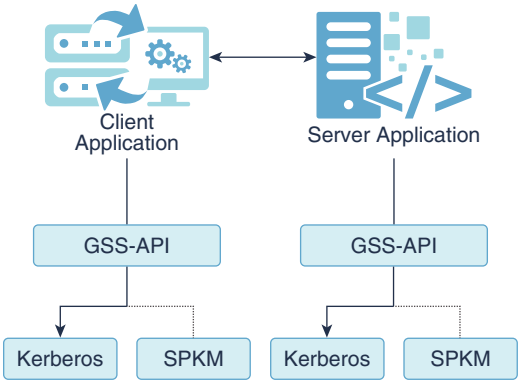

Mechanisms are identified by means of unique object identifier's (OID's) that are registered with the IANA. For instance, the Kerberos V5 mechanism is identified by the OID `{iso(1) member-body(2) United States(840) mit(113554) infosys(1) gssapi(2) krb5(2)}`

Another important feature of the API is that it is token based. i.e., Calls to the API generate opaque octets that the application must transport to its peer. This enables the API to be transport independent.

Java GSS-API

The Java API for the Generic Security Service was also defined at the IETF and is documented in RFC 2853 [10]. Oracle is pursuing the standardization of this API under the Java Community Process (JCP) [11] and plans to deliver a reference implementation with Merlin. Because the JCP is merely endorsing this externally defined API, the IETF assigned package namespace `org.ietf.jgss` will be retained in Merlin.

Oracle's implementation of Java GSS-API, will initially ship with support for the Kerberos V5 mechanism only. Kerberos V5 mechanism support is mandatory for all Java GSS-API implementations in Java SE, although they are free to support additional mechanisms. In a future release, a Service Provider Interface (SPI) will be added so that new mechanisms can be configured statically or even at runtime. Even now the reference implementation in Merlin will be modular and support a private provider SPI that will be converted to public when standardized.

The Java GSS-API framework itself is quite thin, and all security related functionality is delegated to components obtained from the underlying mechanisms. The `GSSManager` class is aware of all mechanism providers installed and is responsible for invoking them to obtain these components.

The implementation of the default `GSSManager` that will ship with Java SE is obtained as follows:

```
GSSManager manager = GSSManager.getInstance();
```

The `GSSManager` can be used to configure new providers and to list all mechanisms already present. The `GSSManager` also serves as a factory class for three important interfaces: `GSSName`, `GSSCredential`, and `GSSContext`. These interfaces are

described below with the methods to instantiate their implementations. For a complete API specification, readers are referred to [9] and [11].

Most calls to Java GSS-API throw a `GSSException` that encapsulate problems that occur both within the GSS-API framework, and within the mechanism providers.

The GSSName Interface

This interface represents an entity for the purposes of Java GSS-API. An implementation of this interface is instantiated as follows:

```
GSSName GSSManager.createName(String name, Oid nameType)
    throws GSSException
```

For example:

```
GSSName clientName = manager.createName("duke", GSSName.NT_USER_NAME);
```

This call returns a `GSSName` that represents the user principal `duke` at a mechanism independent level. Internally, it is assumed that each supported mechanism will map the generic representation of the user to a more mechanism specific form. For instance a Kerberos V5 mechanism provider might map this name to `duke@EXAMPLE.COM` where `EXAMPLE.COM` is the local Kerberos realm. Similarly, a public key based mechanism provider might map this name to an X.509 Distinguished Name.

If we were referring to a principal that was not a user, but some sort of service, we would indicate that to the Java GSS-API call so that the mechanism knows to interpret it differently.

Example:

```
GSSName serverName = manager.createName("nfs@bar.example.com",
                                        GSSName.NT_HOSTBASED_SERVICE);
```

The Kerberos V5 mechanism would map this name to the Kerberos specific form `nfs/bar.example.com@EXAMPLE.COM` where `EXAMPLE.COM` is the realm of the principal. This principal represents the service `nfs` running on the host machine `bar.example.com`.

Oracle's implementation of the `GSSName` interface is a container class. The container class lazily asks the individual providers to perform their mapping when their mechanism is used and then stores each mapped element in a set of principals. In this respect an implementation of `GSSName` is similar to the principal set stored in a Subject. It may even contain the same elements that are in a Subject's principal set, but its use is restricted to the context of Java GSS-API.

The name element stored by the Oracle Kerberos V5 provider is an instance of a subclass of `javax.security.auth.kerberos.KerberosPrincipal`.

The GSSCredential Interface

This interface encapsulates the credentials owned by one entity. Like the `GSSName`, this interface too is a multi-mechanism container.

Its implementation is instantiated as follows:

```
GSSCredential createCredential(GSSName name,
                               int lifetime,
                               Oid[] desiredMechs,
                               int usage)
    throws GSSException
```

Here is an example of this call on the client side:

```
GSSCredential clientCreds =
    manager.createCredential(clientName,
                             8*3600,
                             desiredMechs,
                             GSSCredential.INITIATE_ONLY);
```

The `GSSManager` invokes the providers of the mechanisms listed in the `desiredMechs` for credentials that belong to the `GSSName` `clientName`. Additionally, it imposes the restriction that the credential must be the kind that can initiate outbound requests (i.e., a client credential), and requests a lifetime of 8 hours for it. The returned object contains elements from a subset of `desiredMechs` that had some credential available to satisfy this criteria. The element stored by the Kerberos V5 mechanism is an instance of a subclass of `javax.security.auth.kerberos.KerberosTicket` containing a TGT that belongs to the user.

Credential acquisition on the server side occurs as follows:

```
GSSCredential serverCreds =
    manager.createCredential(serverName,
                             GSSCredential.INDEFINITE_LIFETIME,
                             desiredMechs,
                             GSSCredential.ACCEPT_ONLY);
```

The behavior is similar to the client case, except that the kind of credential requested is one that can accept incoming requests (i.e., a server credential). Moreover, servers are typically long lived and like to request a longer lifetime for the credentials such as the `INDEFINITE_LIFETIME` shown here. The Kerberos V5 mechanism element stored is an instance of a subclass of `javax.security.auth.kerberos.KerberosKey` containing the secret key of the server.

This step can be an expensive one, and applications generally acquire a reference at initialization time to all the credentials they expect to use during their lifetime.

The GSSContext Interface

The `GSSContext` is an interface whose implementation provides security services to the two peers.

On the client side a `GSSContext` implementation is obtained with the following API call:

```
GSSContext GSSManager.createContext(GSSName peer,
                                    Oid mech,
                                    GSSCredential clientCreds,
                                    int lifetime)
    throws GSSException
```

This returns an initialized security context that is aware of the peer that it must communicate with and the mechanism that it must use to do so. The client's credentials are necessary to authenticate to the peer.

On the server side the `GSSContext` is obtained as follows:

```
GSSContext GSSManager.createContext(GSSCredential serverCreds)
    throws GSSException
```

This returns an initialized security context on the acceptor's side. At this point it does not know the name of the peer (client) that will send a context establishment request or even the underlying mechanism that will be used. However, if the incoming request is not for service principal represented by the credentials `serverCreds`, or the underlying mechanism requested by the client side does not have a credential element in `serverCreds`, then the request will fail.

Before the `GSSContext` can be used for its security services it has to be established with an exchange of tokens between the two peers. Each call to the context establishment methods will generate an opaque token that the application must somehow send to its peer using a communication channel of its choice.

The client uses the following API call to establish the context:

```
byte[] GSSContext.initSecContext(byte[] inToken,
                                 int offset,
                                 int len)

   throws GSSException
```

The server uses the following call:

```
byte[] acceptSecContext(byte[] inToken,
                        int offset,
                        int len)

   throws GSSException
```

These two methods are complementary and the input accepted by one is the output generated by the other. The first token is generated when the client calls `initSecContext` for the first time. The arguments to this method are ignored during that call. The last token generated depends on the particulars of the security mechanism being used and the properties of the context being established.

The number of round trips of GSS-API tokens required to authenticate the peers varies from mechanism to mechanism and also varies with characteristics such as whether mutual authentication or one-way authentication is desired. Thus each side of the application must continue to call the context establishment methods in a loop until the process is complete.

In the case of the Kerberos V5 mechanism, there is no more than one round trip of tokens during context establishment. The client first sends a token generated by its `initSecContext()` containing the Kerberos AP-REQ message [2]. In order to generate the AP-REQ message, the Kerberos provider obtains a service ticket for the target server using the client's TGT. The service ticket is encrypted with the server's long-term secret key and is encapsulated as part of the AP-REQ message. After the server receives this token, it is passed to the `acceptSecContext()` method which decrypts the service ticket and authenticates the client. If mutual authentication was not requested, both the client and server side contexts would be established, and the server side `acceptSecContext()` would generate no output.

However, if mutual authentication were enabled, then the server's `acceptSecContext()` would generate an output token containing the Kerberos AP-REP [2] message. This token would need to be sent back to the client for processing by its `initSecContext()`, before the client side context is established.

Note that when a `GSSContext` is initialized on the client side, it is clear what underlying mechanism needs to be used. The Java GSS-API framework can obtain a context implementation from the appropriate mechanism provider. Thereafter, all calls made to the `GSSContext` object are delegated to the mechanism's context implementation. On the server side, the mechanism to use is not decided until the first token from the client side arrives.

Here is a class showing how the client side of an application would be coded. This is the `ClientAction` class that was executed using the `doAs` method in the sample client code in the section The Kerberos Login Module:

```
// Sample client using Java GSS-API

class ClientAction implements PrivilegedAction {

    public Object run() {
        ...
        ...
        try {
            GSSManager manager = GSSManager.getInstance();
            GSSName clientName =
                manager.createName("duke", GSSName.NT_USER_NAME);
            GSSCredential clientCreds =
                manager.createCredential(clientName,
                                         8*3600,
                                         desiredMechs,
                                         GSSCredential.INITIATE_ONLY);
            GSSName peerName =
                manager.createName("nfs@bar.example.com",
                                   GSSName.NT_HOSTBASED_SERVICE);
            GSSContext secContext =
                manager.createContext(peerName,
                                      krb5Oid,
                                      clientCreds,
```

```
                                GSSContext.DEFAULT_LIFETIME);
            secContext.requestMutualAuth(true);

            // The first input token is ignored
            byte[] inToken = new byte[0];

            byte[] outToken = null;

            boolean established = false;

           // Loop while the context is still not established
           while (!established) {
               outToken =
                   secContext.initSecContext(inToken, 0, inToken.length);

               // Send a token to the peer if one was generated
               if (outToken != null)
                  sendToken(outToken);

               if (!secContext.isEstablished()) {
                  inToken = readToken();
               else
                  established = true;
            }
        } catch (GSSException e) {
             ....
        }
        ...
        ...
    }
}
```

The corresponding section of code on the server side running the `ServerAction` class from the sample server code in the section The Kerberos Login Module:

```
// Sample server using Java GSS-API

class ServerAction implelemts PrivilegedAction {

    public Object run() {
        ...
        ...
        try {
            GSSManager manager = GSSManager.getInstance();
            GSSName serverName =
                manager.createName("nfs@bar.example.com",
                                   GSSName.NT_HOSTBASED_SERVICE);
            GSSCredential serverCreds =
             manager.createCredential(serverName,
                                      GSSCredential.INDEFINITE_LIFETIME,
                                      desiredMechs,
                                      GSSCredential.ACCEPT_ONLY);
            GSSContext secContext = manager.createContext(serverCreds);
```

```
            byte[] inToken = null;
            byte[] outToken = null;

            // Loop while the context is still not established
            while (!secContext.isEstablished()) {
                 inToken = readToken();
                 outToken =
                     secContext.acceptSecContext(inToken, 0,
inToken.length);

                  // Send a token to the peer if one was generated
                  if (outToken != null)
                      sendToken(outToken);
            }
        } catch (GSSException e) {
          ...
        }
        ...
        ...
   }
}
```

Message Protection

Once the security context is established, it can be used for message protection. Java GSS-API provides both message integrity and message confidentiality. The two calls that enable this are as follows:

```
byte[] GSSContext.wrap(byte[] clearText,
                       int offset,
                       int len,
                       MessageProp properties)
    throws GSSException
```

and

```
byte[] unwrap(byte[] inToken,
              int offset,
              int len,
              MessageProp properties)
    throws GSSException
```

The `wrap` method is used to encapsulate a cleartext message in a token such that it is integrity protected. Optionally, the message can also be encrypted by requesting this through a `MessageProp` object. The `wrap` method returns an opaque token that the caller sends to its peer. The original cleartext is returned by the peer's `unwrap` method when the token is passed to it. The `MessageProp` object on the `unwrap` side returns information about whether the message was simply integrity protected or whether it was encrypted as well. It also contains sequencing and duplicate token warnings.

Credential Delegation

Java GSS-API allows the client to securely delegate its credentials to the server, such that the server can initiate other security contexts on behalf of the client. This feature is useful for single sign-on in a multi-tier environment. Figure 7-2 illustrates this.

Figure 7-2 Traditional Credential Delegation

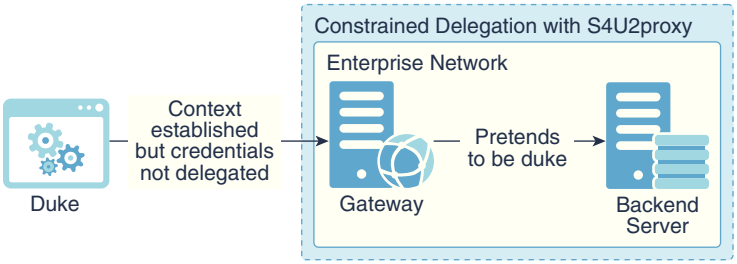

The client requests credential delegation prior to making the first call to `initSecContext()`:

```
void GSSContext.requestCredDeleg(boolean state)
    throws GSSException
```

by setting state to true.

The server receives the delegated credential after context establishment:

```
GSSCredential GSSContext.getDelegCred() throws GSSException
```

The server can then pass this `GSSCredential` to `GSSManager.createContext()` pretending to be the client.

In the case of the Kerberos V5 mechanism, the delegated credential is a forwarded TGT that is encapsulated as part of the first token sent from the client to the server. Using this TGT, the server can obtain a service ticket on behalf of the client for any other service.

MS-SFU Kerberos Extensions

MS-SFU refers to Microsoft Kerberos 5 extensions that allow a service to obtain a Kerberos service ticket on behalf of a client. Microsoft calls this feature constrained delegation. This is useful when the authentication between the client and this service is not established through Kerberos (thus the standard Kerberos delegation cannot be used) but the service needs to access a Kerberos-secured back-end server in the name of the client.

MS-SFU adds two extensions to that protocol: Service for User to Self (S4U2self), which allows a front-end service to obtain a Kerberos service ticket to itself on behalf of a user, and Service for User to Proxy (S4U2proxy), which enables it to obtain a service ticket on behalf of the user to a second, back-end service. Together, these two extensions enable the front-end service to obtain a Kerberos service ticket on behalf of a user. The resulting service ticket can be used to access other services on the local

or remote machines. The public method `ExtendedGSSCredential::impersonate` in the `com.sun.security.jgss` package implements these extensions.

This feature is very useful in secure enterprise deployments. For example, in a typical network service, the front end (such as a web server) often needs to access the back end (such as a database server) on behalf of a client. Normal Kerberos 5 supports delegation, but demands that all layers in this chain explicitly use Kerberos authentication at each step, which is not always possible or desirable.

For example, if a client logs in to a web server using digest authentication, then there are no Kerberos credentials to be delegated, and normal step-by-step Kerberos 5 authentication cannot occur. However, because MS-SFU defines the Service for User (S4U2self) extension so that the front end can access the back end on behalf of the client without presenting the client's Kerberos credentials, MS-SFU could provide authentication in this situation. Figure 7-3 illustrates this.

Figure 7-3 Constrained Delegation with S4U2self

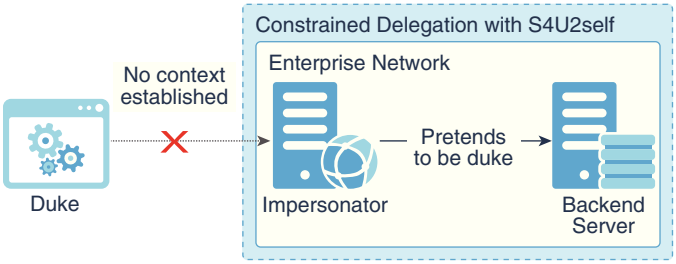

In addition, there are potential security gaps in the standard Kerberos 5 delegation mechanism (which Microsoft calls open delegation). In this mechanism, once the service account has the client's delegated credentials, it has access to any service. Thus, great care is needed with open delegation.

In contrast, with MS-SFU delegation (implemented in S4U2proxy), the administrator can precisely control the services to which a particular service can access on behalf a client.Figure 7-4 illustrates this.

Figure 7-4 Constrained Delegation with S4U2proxy

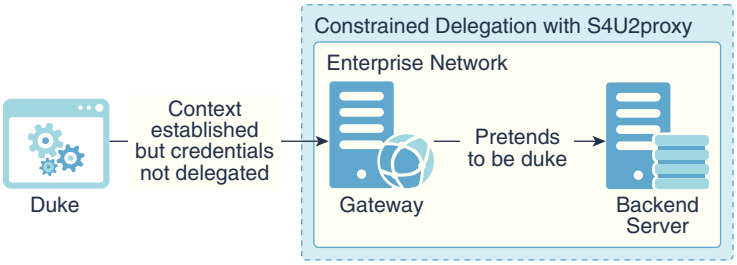

> **Note:**
>
> To delegate credentials as specified in the RFCs in this document, you must use traditional delegation. With constrained delegation, the client is unable to determine if its own credentials can be delegated because this is determined by the KDC.

Default Credential Acquisition Model

Previously we discussed how an application uses the `GSSManager.createCredential()` method to populate a `GSSCredential` object with mechanism specific credentials. The next two sub-sections will focus on how Java GSS-API mechanisms obtain these credentials. The mechanisms do not themselves perform a user login. Instead, the login is performed prior to using Java GSS-API and the credentials are assumed to be stored in some cache that the mechanism provider is aware of. The `GSSManager.createCredential()` method merely obtains references to those credentials and returns them in a GSS-centric container, the `GSSCredential`.

In Java SE, we impose the restriction that the credentials cache that Java GSS-API mechanism providers use to obtain these elements must exclusively be the public and private credential sets in the Subject that is on the current access control context.

This model has the advantage that credential management is simple and predictable from the application's point of view. An application, given the right permissions, can purge the credentials in the Subject or renew them using standard Java API's. If it purged the credentials, it would be sure that the Java GSS-API mechanism would fail, or if it renewed a time based credential it would be sure that the mechanism would succeed.

Here is the sequence of events relevant to credential acquisition when the Kerberos V5 mechanism is used by the client application in Figures 3 and 6:

1. The application invokes a JAAS login, which in turn invokes the configured `Krb5LoginModule`
2. `Krb5LoginModule` obtains a TGT (`KerberosTicket`) for the user either from the KDC or from an existing ticket cache, and stores this TGT in the private credentials set of a Subject
3. The application retrieves the populated Subject, then calls `Subject.doAs`/`doAsPrivileged` which places this Subject on the access control context of the thread executing `ClientAction`
4. `ClientAction` calls the `GSSManager.createCredential` method, passing it the Kerberos V5 OID in `desiredMechs`.
5. GSSManager.createCredential invokes the Kerberos V5 GSS-API provider, asking for a Kerberos credential for initiating security contexts.
6. The Kerberos provider obtains the Subject from the current access control context, and searches through its private credential set for a valid KerberosTicket that represents the TGT for the user.
7. The KerberosTicket is returned to the GSSManager which stores it in a GSSCredential container instance to be returned to the caller.

On the server side, when the Kerberos login is successful in step 2, Krb5LoginModule stores the KerberosKey for the server in the Subject in addition to the KerberosTicket. Later on the KerberosKey is retrieved in steps 5 through 7 and used to decrypt the service ticket that the client sends.

Exceptions to the Model

The default credential acquisition model for Java GSS-API requires credentials to be present in the current Subject. Typically, the credentials are placed there after a JAAS login by the application.

There might be cases where an application wishes to use Kerberos credentials from outside the Subject. It is recommended that such credentials be read as part of the initial JAAS login, either by configuring `Krb5LoginModule` to read them, or by writing a custom login module that reads them. However, some applications might have constrains that either prevent them from using JAAS prior to calling Java GSS-API, or force them to use some Kerberos mechanism provider that does not retrieve credentials from the current Subject.

The system property `javax.security.auth.useSubjectCredsOnly` accommodates such cases while still retaining the standard model for others. This system property serves as a boolean where a value of true requires that the standard credential acquisition model be followed, and a value of false permits the provider to use any cache of it choice. The default value of this property (when it is not set) will be assumed to be true.

If there is no valid Kerberos credential in the current Subject, and this property is true, then the Kerberos mechanism throws a `GSSException`. Setting this property to false does not necessarily mean that the provider has to use a cache other than the current Subject, it only gives the provider the latitude to do so if it wishes.

The Oracle provider for the Kerberos V5 GSS-API mechanism always obtains credentials from a Subject. If there are no valid credentials in the current Subject, and this property is set to false, then the provider attempts to obtain new credentials from a temporary Subject by invoking a JAAS login itself. It uses the text callback handler for input/output with the user, and the JAAS configuration entry identified by "other" for the list of modules and options to use.[2]

The Oracle provider for the Kerberos V5 GSS-API mechanism assumes that one of these modules will be a Kerberos login module. It is possible to configure the modules listed under "other" to read a pre-existing cache so that the user is not unexpectedly prompted for a password in the middle of a Java GSS-API call. The new Subject that is populated by this login is discarded by the Kerberos GSS-API mechanism just as soon as the required credentials are retrieved from it.

Security Risks

The convenience of single sign-on also introduces new risks. What happens if a malicious user gains access to your unattended desktop from where he or she can start applets as you? What happens if malicious applets sign on as you to services that they are not supposed to?

[2] Actually it first tries to use the JAAS configuration entry `com.sun.security.jgss.initiate` for the client and `com.sun.security.jgss.accept` for the server and falls back on the entry for "other", if these entries are missing. This gives system administrators some additional control over its behavior.

For the former, we have no solution but to caution you against leaving your workstation unlocked! For the latter, we have many authorizations checks in place.

To illustrate some details of the permissions model consider an example where your browser has performed a JAAS login at startup time and associated a Subject with all applets that run in it.

The Subject is protected from rogue applets by means of the `javax.security.auth.AuthPermission` class. This permission is checked whenever code tries to obtain a reference to the Subject associated with any access control context.

Even if an applet were given access to a Subject, it needs a `javax.security.auth.PrivateCredentialPermission` to actually read the sensitive private credentials stored in it.

Other kinds of checks are to be done by Java GSS-API mechanism providers as they read credentials and establish security contexts on behalf of the credential's owner. In order to support the Kerberos V5 mechanism, two new permission classes have been added with the package `javax.security.auth.kerberos`:

```
ServicePermission(String servicePrinicipal, String action)
DelegationPermission(String principals)
```

As new GSS-API mechanisms are standardized for Java SE, more packages will be added that contain relevant permission classes for providers of those mechanisms.

The Kerberos GSS-API mechanism permission checks take place at the following points in the program's execution:

Credential Acquisition

The `GSSManager.createCredential()` method obtains mechanism specific credential elements from a cache such as the current Subject and stores them in a `GSSCredential` container. Allowing applets to acquire `GSSCredential` freely, even if they cannot use them to do much, is undesirable. Doing so leaks information about the existence of user and service principals. Thus, before an application can acquire a `GSSCredential` with any Kerberos credential elements in it, a `ServicePermission` check is made.

On the client side, a successful `GSSCredential` acquisition implies that a TGT has been accessed from a cache. Thus the following `ServicePermission` is checked:

```
ServicePermission("krbtgt/EXAMPLE.COM@EXAMPLE.COM", "initiate");
```

The service principal `krbtgt/EXAMPLE.COM@EXAMPLE.COM` represents the ticket granting service (TGS) in the Kerberos realm `EXAMPLE.COM`, and the action "initiate" suggests that a ticket to this service is being accessed. The TGS service principal will always be used in this permission check at the time of client side credential acquisition.

On the server side, a successful `GSSCredential` acquisition implies that a secret key has been accessed from a cache. Thus the following `ServicePermission` is checked:

```
ServicePermission("nfs/bar.example.com@EXAMPLE.COM", "accept");
```

Here the service principal `nfs/bar.example.com` represents the Kerberos service principal and the action "accept" suggests that the secret key for this service is being requested.

Context Establishment

An applet that has permissions to contact a particular server, say the LDAP server, must not instead contact a different server such as the FTP server. Of course, the applet might be restricted from doing so with the help of `SocketPermission`. However, it is possible to use `ServicePermission` to restrict it from authenticating using your identity, even if the network connection was permitted.

When the Kerberos mechanism provider is about to initiate context establishment it checks the `ServicePermission`:

```
ServicePermission("ftp@EXAMPLE.COM", "initiate");
```

This check prevents unauthorized code from obtaining and using a Kerberos service ticket for the principal `ftp@EXAMPLE.COM`.

Providing limited access to specific service principals using this permission is still dangerous. Downloaded code is allowed to communicate back with the host it originated from. A malicious applet could send back the initial GSS-API output token that contains a `KerberosTicket` encrypted in the target service principal's long-term secret key, thus exposing it to an offline dictionary attack. For this reason it is not advisable to grant any "initiate" `ServicePermission` to code downloaded from untrusted sites.

On the server side, the permission to use the secret key to accept incoming security context establishment requests is already checked during credential acquisition. Hence, no checks are made in the context establishment stage.

Credential Delegation

An applet that has permission to establish a security context with a server on your behalf also has the ability to request that your credentials be delegated to that server. But not all servers are trusted to the extent that your credentials can be delegated to them. Thus, before a Kerberos provider obtains a delegated credential to send to the peer, it checks the following permission:

```
DelegationPermission(" \"ftp@EXAMPLE.COM\" \"krbtgt/EXAMPLE.COM@EXAMPLE.COM
\" ");
```

This permission allows the Kerberos service principal ftp@EXAMPLE.COM to receive a forwarded TGT (represented by the ticket granting service krbtgt/EXAMPLE.COM@EXAMPLE.COM).[3]

[3] The use of two principal names in this permission allows for finer grained delegation such as proxy tickets for specific services unlike a carte blanche forwarded TGT. Even though the GSS-API does not allow for proxy tickets, another API such as JSSE might support this idea at some point in the future.

Conclusions

In this paper we have presented a framework to enable single sign-on in Java. This requires sharing of credentials between JAAS which does the initial authentication to obtain credentials, and Java GSS-API which uses those credentials to communicate securely over the wire. We have focused on Kerberos V5 as the underlying security mechanism, but JAAS's stackable architecture and Java GSS-API's multi-mechanism nature allow us to use any number of different mechanisms simultaneously.

The Kerberos login module for JAAS is capable of reading native caches so that users do not have to authenticate themselves beyond desktop login on platforms that support Kerberos. Moreover, the Kerberos V5 mechanism for Java GSS-API allows credentials to be delegated which enables single sign-on in multi-tier environments.

Finally, a number of permissions checks are shown to prevent the unauthorized use of the single-sign on features provided by Kerberos.

Acknowledgements

We thank Gary Ellison, Charlie Lai, and Jeff Nisewanger for their contribution at each stage of the Kerberos single sign-on project. JAAS 1.0 was implemented by Charlie as an optional package for Kestrel (J2SE 1.3). Gary has been instrumental in designing the permissions model for the Kerberos Java GSS-API mechanism. We are grateful to Bob Scheifler for his feedback on integrating JAAS 1.0 into Merlin and to Tim Blackman for the KeyStoreLoginModule and CallbackHandler implementations. We also thank Bruce Rich, Tony Nadalin, Thomas Owusu and Yanni Zhang for their comments and suggestions. We thank Mary Dageforde for the documentation and tutorials. Sriramulu Lakkaraju, Stuart Ke and Shital Shisode contributed tests for the projects. Maxine Erlund provided management support for the project.

References

1. Neuman, Clifford and Tso, Theodore (1994). Kerberos: An Authentication Service for Computer Networks, IEEE Communications, volume 39 pages 33-38
2. J.Kohl and C.Neuman. The Kerberos Network Authentication Service (V5) Internet Engineering Task Force, September 1993 Request for Comments 1510
3. V. Samar and C. Lai. Making Login Services Independent from Authentication Technologies. In Proceedings of the SunSoft Developer's Conference, March 1996.
4. X/Open Single Sign-On Service (XSSO) - Pluggable Authentication. Preliminary Specification P702, The Open Group, June 1997. https://www.opengroup.org
5. A Smart Card Login Module for Java Authentication and Authorization Service. http://www.gemplus.com/techno/smartjaas/index.html
6. J. Linn. Generic Security Service Application Program Interface,Version 2. Internet Engineering Task Force, January 2000 Request for Comments 2743
7. J. Linn. The Kerberos Version 5 GSS-API Mechanism. Internet Engineering Task Force, June 1996 Request for Comments 1964
8. C.Adams. The Simple Public-Key GSS-API Mechanism (SPKM). Internet Engineering Task Force, October 1996 Request for Comments 2025

9. J. Kabat and M.Upadhyay. Generic Security Service API Version 2: Java Bindings. Internet Engineering Task Force, January 1997 Request for Comments 2853
10. JSR 000072 Generic Security Services API
11. Java Platform, Standard Edition API Specification

Advanced Security Programming in Java SE Authentication, Secure Communication and Single Sign-On

Java SE offers a rich set of APIs and features for developing secure Java applications and services. The exercise sessions listed here can help you to use the Java SE GSS APIs to build applications that authenticate their users, to communicate securely with other applications and services, and help you to configure your applications in a Kerberos environment to achieve Single Sign-On. In addition, you will also learn how to use stronger encryption algorithms in a Kerberos environment, and how to use Java GSS mechanisms such as SPNEGO to secure the association.

Setting up your Development Environment

Set up your development environment as follows before proceeding to the first exercise:

1. Install and set up a Solaris machine (required for Exercise 7: Configuring to Use Stronger Encryption Algorithms in a Kerberos Environment, to Secure the Communication).
2. Configure a Kerberos server on a Solaris machine with accounts used by the exercises. See Appendix A: Setting up Kerberos Accounts.
3. Set up the Key Distribution Center (KDC) on your Solaris machine and start the Kerberos server.
4. Set up the Kerberos configuration on your client machine.
5. Set up the JDK environment:
   - Set up the `JAVA_HOME` environment variable to point to the JDK installation directory
   - Place `%JAVA_HOME%\bin` (Windows) or `$JAVA_HOME/bin` (Solaris/Linux) in the `PATH` environment variable.

Exercises

This session includes six lessons. Each part contains one or more coding exercises. Work through the exercises in sequence:

- Part I : Secure Authentication using the Java Authentication and Authorization Service (JAAS)
  - Exercise 1: Using the JAAS API
  - Exercise 2: Configuring JAAS for Kerberos Authentication
- Part II : Secure Communications using the Java SE Security API
  - Exercise 3: Using the Java Generic Security Service (GSS) API
  - Exercise 4: Using the Java SASL API

  - Exercise 5: Using the Java Secure Socket Extension with Kerberos
- Part III : Deploying for Single Sign-On in a Kerberos Environment
  - Exercise 6: Deploying for Single Sign-On
- Part IV : Secure Communications Using Stronger Encryption Algorithms
  - Exercise 7: Configuring to Use Stronger Encryption Algorithms in a Kerberos Environment, to Secure the Communication
- Part V : Secure Authentication Using SPNEGO Java GSS Mechanism
  - Exercise 8: Using the Java Generic Security Services (GSS) API with SPNEGO
- Part VI: HTTP/SPNEGO Authentication
  - Exercise 9: Using HTTP/SPNEGO Authentication

Part I : Secure Authentication using the Java Authentication and Authorization Service (JAAS)

Exercise 1: Using the JAAS API

Goal of This Exercise

The goal of this exercise is to learn how to use the Java Authentication and Authorization (JAAS) API to perform authentication.

Background for This Exercise

JAAS provides a standard pluggable authentication framework (PAM) for the Java platform. An application uses the JAAS API to perform *authentication* - the process of verifying the identity of the user who is using the application and gathering his identity information into a container called a *subject*. The application can then use the identity information in the subject along with the JAAS API to make *authorization* decisions, to decide whether the authenticated user is allowed to access protected resources or perform restricted actions. This exercise demonstrates JAAS Authentication. It does not demonstrate JAAS Authorization.

Resources for This Exercise

- Java Authentication and Authorization Service (JAAS) Reference Guide
- JAAS Tutorials
- JAAS Javadoc API documentation
  - `javax.security.auth`
  - `javax.security.auth.callback`
  - `javax.security.auth.kerberos`
  - `javax.security.auth.login`
  - `javax.security.auth.spi`
  - `javax.security.auth.x500`

Steps to Follow

- Read the `Jass.java` sample code. The code performs the following tasks:
  1. Define a callback handler or use a predefined one.
  2. Create a `LoginContext` with a name that identifies which JAAS configuration entry to use.
  3. Perform the authentication.
  4. Define the task that the authenticated user is to perform.
  5. Perform the action as the authenticated user.
  6. Log out.

  `Subject.doAs` will run the code defined in `MyAction` as the authenticated user [lines 14-15]. This serves two purposes. First, code in `MyAction` that requires identity information for authentication to a service could get it from the subject. This exercise demonstrates this use. Second, if `MyAction` accesses any protected resources/operations, the identity information in the current subject would be used to make the corresponding access control decision. This second aspect is not covered in this exercise.
- Make sure that the `%JAVA_HOME%/bin` is in the PATH environment variable.
- Compile the modified sample code. You will run this code in subsequent exercises after doing some set up. That ends this exercise.

Summary

This exercise introduced the main classes of the JAAS APIs: `LoginContext` and `Subject`. You learned how to use `LoginContext` to authenticate a user and collect its identity information in a Subject. You then learned how to use the Subject to perform an action as the authenticated user.

Next Steps

Proceed to Exercise 2: Configuring JAAS for Kerberos Authentication to learn how to configure the sample application to use Kerberos for authentication.

Exercise 2: Configuring JAAS for Kerberos Authentication

Goal of This Exercise

The goal of this exercise is to learn how to configure a JAAS application to use Kerberos for authentication.

Kerberos Background for This Exercise

Kerberos is an Internet standard protocol for trusted-third party authentication defined in RFC 4120. It is available on most modern computing platforms today, including Solaris, Windows, and Linux.

The Kerberos architecture is centered around a trusted authentication service called the key distribution center, or KDC. Users and services in a Kerberos environment are referred to as principals; each principal shares a secret (such as a password) with the KDC. A principal authenticates to Kerberos by proving to the KDC that it knows the

shared secret. If the authentication is successful, the KDC issues a ticket-granting-ticket (TGT) to the principal. When the principal subsequently wants to authenticate to a service on the network, such as a directory service or a file service, (thereby, acting as a "client" of the service), it gives the TGT to the KDC to obtain a service ticket to communicate with the service. Not only does the service ticket indicate the identities of the client and service principals, it also contains a session key that can be used by the client and service to subsequently establish secure communication. To authenticate to the service, the client sends the service ticket to the service. When the service receives the ticket, it decodes it using the secret it shares with the KDC.

In this architecture, a principal only authenticates directly (once) to the KDC. It authenticates indirectly to all other services via the use of service tickets. Service tickets are how the KDC vouches for the identity of a principal. The ability of a principal to access multiple secure services by performing explicit authentication only once is called single sign-on.

JAAS Background for This exercise

In JAAS, for a client principal, "logging into Kerberos" means acquiring the TGT and placing it in the Subject, so that it can be used for authentication with services that the client will access. For a service principal, "logging into Kerberos" means obtaining the secret keys that the service needs to decode incoming client authentication requests.

Resources for This Exercise

- Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide
- The Kerberos Network Authentication Service (v5)
- Appendix B: JAAS Login Configuration File
- Login module package Javadoc API documentation:`com.sun.security.auth.module`
- Introduction to JAAS and Java GSS-API Tutorials

Steps to Follow

1. Examine the `jaas-krb5.conf` configuration file.

   This file contains two entries, one named *client* and one named *server*. The client entry indicates that the `LoginContext` must use the **`com.sun.security.auth.module.Krb5LoginModule`**. The server entry indicates that the `LoginContext` must use the same login module, and use keys from the `sample.keytab` file for the principal host/machineName.

2. Determine the hostname of your machine by executing the hostname command.
3. Edit this file and change the entry for server principal to use the name of your machine. For example, if your machine name is j1hol-001, this line in the configuration file should look like this:

   ```
   principal="host/j1hol-001"
   ```

4. Perform client authentication by typing the following command:

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf Jaas client
   ```

You will be prompted for a password. You should see the following output. Replace *password* with a password that is secure.

```
Kerberos password for test: password
Authenticated principal: [test@J1LABS.EXAMPLE.COM]
Performing secure action...
```

5. Perform server authentication by typing the following command:

```
% java -Djava.security.auth.login.config=jaas-krb5.conf Jaas server
```

You should see the following output:

```
Authenticated principal: [host/j1hol-001@J1LABS.EXAMPLE.COM] Performing
secure action...
```

Summary

In this exercise, you learned how to configure a JAAS application to use a Kerberos login module, both as a client principal who enters his/her username/password interactively, and as a service principal who gets its keys from a `keytab` file.

Next Steps

Proceed to Part II : Secure Communications using the Java SE Security API to learn how to establish secure communication channels using Java security APIs.

Part II : Secure Communications using the Java SE Security API

This part shows you how to build applications that perform secure communications. The Java SE platform provides three standard APIs that allow applications to perform secure communications: The Java Generic Security Service (GSS), the Java SASL API, and the Java Secure Socket Extension (JSSE). When building an application, which of these APIs should you use? The answer depends on many factors, including requirements of the protocol or service, deployment infrastructure, and integration with other security services. For example, if you are building an LDAP client library, you would need to use the Java SASL API because use of SASL is part of LDAP's protocol definition. As an other example, if the service supports SSL, then the client application attempting to access the service would need to use JSSE.

Exercise 3: Using the Java Generic Security Service (GSS) API

Goal of This Exercise

The goal of this exercise is to learn how to use the Java GSS API to perform secure authentication and communication.

Background for This Exercise

The Generic Security Service API provides a uniform C-language interface to access various security services, such as authentication, message integrity, and message confidentiality. The Java GSS API provides the corresponding interface for Java applications. It allows applications to perform authentication and establish secure

communication with the peer. One of the most common security service accessed via the GSS-API and Java GSS-API is Kerberos.

Resources for This Exercise

- Introduction to JAAS and Java GSS-API Tutorials
- Generic Security Service API Version 2: Java Bindings (RFC 2853)
- Java GSS Javadoc API documentation: `org.ietf.jgss`.

Overview of This Exercise

This exercise is a client-server application that demonstrates how to communicate securely using the Java GSS API. The client and server parts first authenticate to Kerberos, as shown in Exercise 1: Using the JAAS API. This stores the credentials in the subject. The application then executes an action that performs Java GSS operations (with Kerberos as the underlying GSS mechanism) inside of a `Subject.doAs` using the subject. The Java GSS Kerberos mechanism, because it is executing inside the `doAs`, obtains the Kerberos credentials from the subject, and uses them to authenticate with the peer and to exchange messages securely.

Steps to Follow

1. Read the `GssServer.java` code.

   This code fragment defines the action to execute after the service principal has authenticated to the KDC. It replaces the `MyAction` in Exercise 1: Using the JAAS API. The code first creates an instance of **`GSSManager`**, which it uses to obtain its own credentials and to create an instance of `GSSContext`. It uses this context to perform authentication. Upon completing authentication, it accepts encrypted input from the client and uses the established security context to decrypt the data. It then uses the security context to encrypt a reply containing the original input and the date, and then sends it back to the client.

2. Compile the sample code.

3. Read the `GssClient.java` code.

   This code fragment defines the action to execute after the client principal has authenticated to the KDC. It replaces the `MyAction` in Exercise 1: Using the JAAS API. The code first creates an instance of **`GSSManager`**, which it uses to obtain a principal name for the service that it is going to communicate with. It then creates an instance of `GSSContext` to perform authentication with the service. Upon completing authentication, it uses the established security context to encrypt a message, and sends it to the server. It then reads an encrypted message from the server and decodes it using the established security context.

4. Compile the sample code.

5. Launch a new window and start the server:

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf GssServer
   ```

6. Run the client application. `GssClient` takes two parameters: the service name and the name of the server that the service is running on. For example, if the service is host running on the machine j1hol-001, you would enter the following:

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf GssClient host
   j1hol-001
   ```

   When prompted for the password, enter `change_it`.

7. Observe the following output in the respective client and server applications' windows.

   Output for running `GssServer` example:

   ```
   Authenticated principal: [host/j1hol-001@J1LABS.EXAMPLE.COM]
   Waiting for incoming connections..
   Got connection from client /192.0.2.102
   Context Established!
   Client principal is test@J1LABS.EXAMPLE.COM
   Server principal is host/j1hol-001@J1LABS.EXAMPLE.COM
   Mutual authentication took place!
   Received data "Hello There!" of length 12
   Confidentiality applied: true
   Sending: Hello There! Thu May 06 12:11:15 PDT 2005
   ```

   Output for running `GssClient` example:

   ```
   Kerberos password for test: change_it
   Authenticated principal: [test@J1LABS.EXAMPLE.COM]
   Connected to address j1hol-001/192.0.2.102
   Context Established!
   Client principal is test@J1LABS.EXAMPLE.COM
   Server principal is host@j1hol-001
   Mutual authentication took place!
   Sending message: Hello There!
   Will read token of size 93
   Received message: Hello There! Thu May 06 12:11:15 PDT 2005
   ```

Summary

In this exercise, you learned how to write a client-server application that uses the Java GSS API to authenticate and communicate securely with each other.

Next Steps

1. Proceed to Exercise 4: Using the Java SASL API to learn how to write a client/server application that uses the Java SASL API to authenticate and communicate securely with each other.
2. Proceed to Exercise 5: Using the Java Secure Socket Extension with Kerberos to learn how to write a client/server application that uses the JSSE to authenticate and communicate securely with each other.
3. Proceed to Exercise 6: Deploying for Single Sign-On to learn how to configure the sample programs that you have just used to achieve single sign-on in a Kerberos environment.

Exercise 4: Using the Java SASL API

Goal of This Exercise

The goal of this exercise is to learn how to use the Java SASL API to perform secure authentication and communication.

Background for This Exercise

Simple Authentication and Security Layer (SASL) specifies a challenge-response protocol in which data is exchanged between the client and the server for the purposes of authentication and (optional) establishment of a security layer on which to carry on subsequent communications. SASL allows different mechanisms to be used; each such mechanism is identified by a profile that defines the data to be exchanged and a name. SASL is used with connection-based protocols such as LDAPv3 and IMAPv4. SASL is described in RFC 4422.

The Java SASL API defines an API for applications to use SASL in a mechanism-independent way. For example, if you are writing a library for a networking protocol that uses SASL, you can use the Java SASL API to generate the data to be exchanged with the peer. When the library is deployed, you can dynamically configure the mechanisms to use with the library.

In addition to authentication, you can use SASL to negotiate a security layer to be used after authentication. But unlike the GSS-API, the properties of the security layer (such as whether you want integrity or confidentiality) is decided at negotiation time. (the GSS-API allows confidentiality to be turned on or off per message).

Resources for This Exercise

- Java SASL API Programming and Deployment Guide
- `javax.security.sasl`
- Simple Authentication and Security Layer (SASL) (RFC 4422)

Overview of This Exercise

This exercise is a client-server application that demonstrates how to communicate securely using the Java SASL API. The client and server parts first authenticate to Kerberos using Exercise 1: Using the JAAS API. This stores the credentials in the subject. The application then executes an action that performs Java SASL API operations (with Kerberos as the underlying SASL mechanism) inside of a `Subject.doAs` using the subject. The SASL/Kerberos mechanism, because it is executing inside the `doAs`, obtains the Kerberos credentials from the subject, and uses them to authenticate with the peer and to exchange messages securely.

This example uses a simple protocol implemented by the `AppConnection` class. This protocol exchanges authentication commands and data commands. Each command consists of a type (e.g., `AppConnection.AUTH_CMD`), the length of the data to follow, and the data itself. The data is a SASL buffer if it is for authentication or encrypted/integrity-protected application data; it is plain application data otherwise.

Steps to Follow

1. Read the `SaslTestServer.java` sample code.

This code fragment defines the action to execute after the service principal has authenticated to the KDC. It replaces the `MyAction` in Exercise 1: Using the JAAS API. The server specifies the quality of protections (QOP) that it will support and then creates an instance of `SaslServer` to perform the authentication. The challenge-response protocol of SASL is performed in the while loop, with the server sending challenges to the client and processing the responses from the client. After authentication, the identity of the authenticated client can be obtained via a call to the `getAuthorizedID()` method. If a security layer was negotiated, the server can exchange data securely with the client.

2. Compile the sample code.
3. Read the `SaslTestClient.java` sample code.

   This code fragment defines the action to execute after the client principal has authenticated to the KDC. It replaces the MyAction in Exercise 1: Using the JAAS API. The program first specifies the quality of protections that it wants (in this case, confidentiality) and then creates an instance of `SaslClient` to use for authentication. It then checks whether the mechanism has an initial response and if so, gets the response by invoking the `evaluateChallenge()` method with an empty byte array. It then sends the response to the server to begin the authentication. The challenge-response protocol of SASL is performed in the while loop, with the client evaluating the challenges that it gets from the server and sending the server the corresponding responses to the challenges. After authentication, the client can proceed to communicate with the server using the negotiated security layer.
4. Compile the sample code.
5. Launch a new window and start the server. `SaslTestServer` takes two parameters: the service name and the name of the server that the service is running on. For example, if the service is host running on the machine j1hol-001, you would enter the following:

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf SaslTestServer
   host j1hol-001
   ```

6. Run the client application. `SaslTestClient` takes two parameters: the service name and the name of the server that the service is running on. For example, if the service is host running on the machine j1hol-001, you would enter the following:

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf SaslTestClient
   host j1hol-001
   ```

   Provide a secure password.
7. Observe the following output in the respective client and server applications' windows.

   Output for running the `SaslTestServer` example:

   ```
   Authenticated principal: [host/j1hol-001@J1LABS.EXAMPLE.COM]
   Waiting for incoming connections...
   Got connection from client /192.0.2.102
   Client authenticated; authorized client is: test@J1LABS.EXAMPLE.COM
   Negotiated QOP: auth-conf
   ```

```
Received: Hello There!
Sending: Hello There! Fri May 07 15:32:37 PDT 2005
Received data "Hello There!" of length 12
```

Output for running the `SaslTestClient` example (*password* will be replaced by the password that you provided):

```
Kerberos password for test: password
Authenticated principal: [test@J1LABS.EXAMPLE.COM]
Connected to address j1hol-001/192.0.2.102
Client authenticated.
Negotiated QOP: auth-conf
Sending: Hello There!
Received: Hello There! Fri May 07 15:32:37 PDT 2005
```

Summary

In this exercise, you learned how to write a client-server application that uses the Java SASL API to authenticate and communicate securely with each other.

Next Steps

1. Proceed to Exercise 5: Using the Java Secure Socket Extension with Kerberos to learn how to write a client/server application that uses the JSSE to authenticate and communicate securely with each other.
2. Proceed to Exercise 6: Deploying for Single Sign-On to learn how to configure the sample programs that you have just used to achieve single sign-on in a Kerberos environment.

Exercise 5: Using the Java Secure Socket Extension with Kerberos

Goal of This Exercise

The goal of this exercise is to learn how to use the JSSE API to perform secure authentication and communication using Kerberos cipher suites.

Background for This Exercise

Secure Socket Layer (SSL) and Transport Layer Security (TLS) are the most widely used protocols for implementing cryptography on the Internet. TLS is the Internet standard evolved from SSL. SSL/TLS provides application-level protocols (such as HTTP and LDAP) with secure authentication and communication. For example, HTTPS is the resulting protocol of using HTTP over SSL/TLS. SSL/TLS is used not only for standard protocols such as HTTP, it is also widely used when building custom applications using custom protocols that need to communicate securely.

SSL/TLS traditionally used certificate-based authentication and is commonly used for server-authentication. For example, when a Web client such as a browser accesses a secure Web site (server) on behalf of a user, the server sends its certificate to the browser so that the browser can verify the identity of the server. This ensures that the user does not divulge confidential information (such as credit card information) to a bogus server. Recently, a new standard allows the use of Kerberos with TLS. This means instead of using certificate-based authentication, an application can use Kerberos credentials and take advantage of the Kerberos infrastructure in the

deployment environment. Using Kerberos cipher suites also provides automatic support for mutual authentication in which the client is also authenticated in addition to the server.

The decision of whether to use Java GSS, Java SASL, or JSSE for a particular application often depends upon several factors, including (the protocols being used by) the services with which the application interacts, the deployment environment (PKI or Kerberos-based), and the application's security requirements. JSSE provides a secure end-to-end channel that takes care of the I/O and transport, while Java GSS and Java SASL provide encryption and integrity-protection on the data, but the application is responsible for transporting the secured data to its peer. Some details about factors for deciding when to use JSSE versus Java GSS are presented in the document, When to Use Java GSS-API Versus JSSE.

Resources for This Exercise

- Java Secure Socket Extension (JSSE) Reference Guide
- The JSSE Javadoc API: `javax.net` and `javax.net.ssl`
- The SSL Protocol version 3.0
- The TLS Protocol Version 1.0 (RFC 2246)
- Addition of Kerberos Cipher Suites to Transport Layer Security TLS (RFC 2712)
- When to Use Java GSS-API Versus JSSE

Overview of This Exercise

This exercise is a client-server application that demonstrates how to communicate securely using the JSSE and Kerberos cipher suites. The client and server parts first authenticate to Kerberos using Exercise 1: Using the JAAS API. This stores the credentials in the subject. The application then executes an action that performs JSSE operations (using a Kerberos cipher suite) inside of a `Subject.doAs` using the subject. The Kerberos cipher suite implementation, because it is executing inside the `doAs`, obtains the Kerberos credentials from the subject, and uses them to authenticate with the peer and to exchange messages securely. This example sends newline-terminated messages, encrypted using the negotiated cipher suite and integrity-protected, back and forth between client and server.

According to the standard (RFC 2712) all Kerberos-enabled TLS applications use the same service name (*host*). That is why in this exercise, you do not need to explicitly supply the Kerberos service name.

Steps to Follow

1. Read the `JsseServer.java` sample code.

   This code fragment defines the action to execute after the service principal has authenticated to the KDC. It replaces the `MyAction` in Exercise 1: Using the JAAS API. The server first creates an `SSLServerSocket`. This is analogous to an application creating a plain `ServerSocket` except an `SSLServerSocket` will provide automatic authentication, encryption and decryption, as needed. The server then sets the cipher suites that it wants to use. The server then runs in a loop, accepting connections from SSL clients, and reads and writes from the SSL socket. The server can find out the identities of the owners of socket by invoking the `getLocalPrincipal()` and `getPeerPrincipal()` methods.

2. Compile the sample code.

3. Read the `JsseClient.java` sample code.

   This code fragment defines the action to execute after the client principal has authenticated to the KDC. It replaces the `MyAction` in Exercise 1: Using the JAAS API. The client first creates an `SSLSocket`. The client then sets the cipher suites that it wants to use. The client then exchanges messages with the server using the `SSLSocket` by reading and writing to the socket's input/output streams. The client can find out the identities of the owners of socket by invoking the `getLocalPrincipal()` and `getPeerPrincipal()` methods.

4. Compile the sample code.

5. Launch a new window and start the server. `JsseServer` takes one parameter: the name of the server that the JSSE service is running on. For example, if it is running on the machine j1hol-001, you would enter the following:

   ```
   % xterm &
   % java -Djava.security.auth.login.config=jaas-krb5.conf JsseServer
   j1hol-001
   ```

6. Run the client application. `JsseClient` takes one parameter: the name of the server that the JSSE service is running on. For example, if the service is running on the machine j1hol-001, you would enter the following.

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf JsseClient
   j1hol-001
   ```

   Provide a secure password.

7. Observe the following output in the respective client and server applications' windows.

   Output for running the `JsseServer` example:

   ```
   Authenticated principal: [host/j1hol-001@J1LABS.EXAMPLE.COM]
   Waiting for incoming connections...
   Got connection from client /192.0.2.102
   Received: Hello There!
   Sending: Hello There! Fri May 07 15:32:37 PDT 2005
   Cipher suite in use: TLS_KRB5_WITH_3DES_EDE_CBC_SHA
   I am: host/j1hol-001@J1LABS.EXAMPLE.COM
   Client is: test@J1LABS.EXAMPLE.COM
   ```

   Output for running the `JsseClient` example (*password* will be replaced by the password that you provided):

   ```
   Kerberos password for test: password
   Authenticated principal: [test@J1LABS.EXAMPLE.COM]
   Sending: Hello There!
   Received: Hello There! Fri May 07 15:32:37 PDT 2005
   Cipher suite in use: TLS_KRB5_WITH_3DES_EDE_CBC_SHA
   I am: test@J1LABS.EXAMPLE.COM
   Server is: host/j1hol-001@J1LABS.EXAMPLE.COM
   ```

Summary

In this exercise, you learned how to write a client-server application that uses JSSE to authenticate and communicate securely with each other, using Kerberos as the underlying authentication system.

Next Steps

Proceed to Exercise 6: Deploying for Single Sign-On to learn how to configure the sample programs in Exercises 3, 4, and 5 to achieve single sign-on in a Kerberos environment.

Part III : Deploying for Single Sign-On in a Kerberos Environment

Exercise 6: Deploying for Single Sign-On

Goal of This Exercise

The goal of this exercise is to learn how to configure a JAAS application that uses Kerberos for authentication to achieve single sign-on. Single sign-on means that the user needs only authenticate once to a system or a collection of services. After the initial authentication, the user can access other services in the system using the same identity as he used for the initial authentication.

Single sign-on can be used to describe different types of authentication. There are HTTP-based network single sign-on protocols. There is Kerberos-based single sign-on for network services. In this particular exercise, we show how to achieve single sign-on in Kerberos-based systems by showing how to import already-acquired Kerberos credentials from the underlying native operating system.

Background and Resources for This Exercise

See Single Sign-on Using Kerberos in Java. In addition, see the information provided in Exercise 2: Configuring JAAS for Kerberos Authentication and Exercise 4: Using the Java SASL API for background information about Kerberos and Java GSS.

Steps to Follow

1. Edit the `jaas-krb5.conf` configuration file.

   This file contains two entries: one named *client* and one named *server*. Add the line `useTicketCache=true` to the client entry.

2. Perform Kerberos login to the native operating system. To login to Kerberos, use kinit command as follows:

   ```
   % kinit test
   ```

   Provide a secure password.

3. Run the client and server programs in Exercises 1 through 5 and you will note that the client applications no longer ask you to enter a password.

Part IV : Secure Communications Using Stronger Encryption Algorithms

Exercise 7: Configuring to Use Stronger Encryption Algorithms in a Kerberos Environment, to Secure the Communication

Goal of This Exercise

The goal of this exercise is to learn how to use various Kerberos encryption algorithms to secure the communication. Java GSS/Kerberos provides a wide range of encryption algorithms, including AES256, AES128, 3DES, RC4-HMAC, and DES.

> **Note:**
>
> DES-based encryption types are disabled by default.

The following is a list of all the encryption types supported by the Java GSS/Kerberos provider in Java SE:

- AES256-CTS
- AES128-CTS
- AES256-SHA2
- AES128-SHA2
- RC4-HMAC
- DES3-CBC-SHA1
- DES-CBC-MD5
- DES-CBC-CRC

Steps to Follow

1. Configure the Key Distribution Center (KDC) and update the Kerberos database.

   First, you need to update to use the KDC that supports the required Kerberos encryption types, such as the latest version of Solaris or the latest version of Kerberos from the MIT distribution. If you are using Active Directory on a Windows platform, the latest version also supports RC4-HMAC and AES encryption types.

   You need to update the Kerberos database to generate the new keys with stronger encryption algorithms. By default, Solaris 10 KDC will generate the keys for all the above encryption types listed. You can now create a keytab that will include all the keys for all the above encryption types.

> **Note:**
>
> If you want to use Windows 2000 KDC, you can configure to use the DES and RC4-HMAC encryption types that are available on Windows.

2. Edit the Kerberos configuration file (`krb5.conf`).

   You will need to edit the Kerberos configuration file in order to select the desired encryption types used. The required parameters that you will need to insert under the `libdefaults` section of the Kerberos configuration file are listed below. For the purpose of this exercise, all the required entries have been added to a sample Kerberos configuration file included with the exercise, and the entries have been commented out. To enable the desired encryption type, you only need to uncomment the required entries.

   - To only enable AES256-CTS encryption type, add the following:

     ```
     [libdefaults]
     default_tkt_enctypes = aes256-cts
     default_tgs_enctypes = aes256-cts
     permitted_enctypes = aes256-cts
     ```

     > **Note:**
     >
     > Solaris 10 11/06 and earlier does not support AES256 by default. You will need to install the following packages: SUNWcry, SUNWcryr, SUNWcryptoint.

   - To only enable AES128-CTS encryption type, add the following:

     ```
     [libdefaults]
     default_tkt_enctypes = aes128-cts
     default_tgs_enctypes = aes128-cts
     permitted_enctypes = aes128-cts
     ```

   - To only enable RC4-HMAC encryption type, add the following:

     ```
     [libdefaults]
     default_tkt_enctypes = rc4-hmac
     default_tgs_enctypes = rc4-hmac
     permitted_enctypes = rc4-hmac
     ```

   - To only enable DES3-CBC-SHA1 encryption type, add the following:

     ```
     [libdefaults]
     default_tkt_enctypes = des3-cbc-sha1
     default_tgs_enctypes = des3-cbc-sha1
     permitted_enctypes = des3-cbc-sha1
     ```

> **Note:**
>
> - If the above parameters are not added to the Kerberos configuration file, Solaris 10 11/06 and earlier will default to use AES128 enctype. If the exportable crypto packages have been installed, it will default to use AES256 enctype.
> - Destroy any pre-existing Kerberos TGT in the ticket cache from the previous exercise as follows:
>
> ```
> % kdestroy
> ```

3. Launch a new window and start the server using the updated `krb5.conf` as follows:

```
% java -Djava.security.auth.login.config=jaas-krb5.conf \
-Djava.security.krb5.conf=krb5.conf GSSServer
```

4. Run the client application using the updated `krb5.conf`. The `GSSClient` class takes two parameters: the service name and the name of the server that the service is running on. For example, if the service is host running on the machine j1hol-001, use the following (provide a secure password when prompted):

```
% java -Djava.security.auth.login.config=jaas-krb5.conf \
-Djava.security.krb5.conf=krb5.conf \
GSSClient host j1hol-001
```

Summary

In this exercise, you learned how to write a client-server application that uses Java GSS API to authenticate and communicate securely using stronger Kerberos encryption algorithms. You can enable Kerberos debugging (`-Dsun.security.krb5.debug=true`), to obtain information about the Kerberos encryption type used.

Part V : Secure Authentication Using SPNEGO Java GSS Mechanism

Exercise 8: Using the Java Generic Security Services (GSS) API with SPNEGO

Java GSS is a framework that can support multiple security mechanisms; a way to negotiate a security mechanism underneath GSS-API is needed. This is available via SPNEGO.

SPNEGO is standardized at IETF in RFC 4178. It is a pseudo-security mechanism used to negotiate an underlying security mechanism. It provides the flexibility for client and server to securely negotiate a common GSS security mechanism.

Microsoft makes heavy use of SPNEGO. SPNEGO can be used to inter-operate with Microsoft Server over HTTP, to support HTTP-based cross-platform authentication via the Negotiate Protocol.

Currently, when using Java GSS with Kerberos, we specify the Kerberos OID as follows:

```
Oid krb5Oid = new Oid("1.2.840.113554.1.2.2");
```

In order to use SPNEGO, you only need to specify the SPNEGO OID as follows:

```
Oid spnegoOid = new Oid("1.3.6.1.5.5.2");
```

Then you can use the SPNEGO OID when creating a `GSSCredential`, `GSSContext`, etc.

Goal of This Exercise

Currently the only security mechanism available with Java GSS is Kerberos. The goal of this exercise is to learn how to use other Java GSS mechanisms, such as the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO), to secure the association.

Steps to Follow

1. Read the `GssSpNegoClient.java` code.
2. Compile the sample code:

   ```
   % javac GssSpNegoClient.java
   ```

3. Read the `GssSpNegoServer.java` code.
4. Compile the sample code:

   ```
   % javac GssSpNegoServer.java
   ```

5. Launch a new window and start the server:

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf GssSpNegoServer
   ```

6. Run the client application. `GssSpNegoClient` takes two parameters: the service name and the name of the server that the service is running on. For example, if the service is host running on the machine j1hol-001, use the following (provide a secure password when prompted):

   ```
   % java -Djava.security.auth.login.config=jaas-krb5.conf \
   GssSpNegoClient host j1hol-001
   ```

   Sample output for running `GssSpNegoServer`:

   ```
   Authenticated principal: [host/j1hol-001@J1LABS.EXAMPLE.COM]
   Waiting for incoming connections...
   Got connection from client /129.145.128.102
   SPNEGO Negotiated Mechanism = 1.2.840.113554.1.2.2 Kerberos V5
   Context Established!
   Client principal is test@J1LABS.EXAMPLE.COM
   Server principal is
   ```

```
host/j1hol-001@J1LABS.EXAMPLE.COM
Mutual authentication took place!
Received data "Hello There!" of length 12
Confidentiality applied: true
Sending: Hello There! Thu May 06 12:11:15 PDT 2005
```

Sample output for running `GssSpNegoClient` (*`password`* is replaced with the password you provided before):

```
Kerberos password for test: password
Authenticated principal: [test@J1LABS.EXAMPLE.COM]
Connected to address j1hol-001/129.145.128.102
SPNEGO Negotiated Mechanism = 1.2.840.113554.1.2.2 Kerberos V5
Context Established!
Client principal is test@J1LABS.EXAMPLE.COM
Server principal is host@j1hol-001
Mutual authentication took place!
Sending message: Hello There!
Will read token of size 93
Received message: Hello There! Thu May 06 12:11:15 PDT 2005
```

Summary

In this exercise, you learned how to write a client-server application that uses the Java GSS API with SPNEGO to negotiate an underlying security mechanism, such as Kerberos, and communicate securely using Kerberos as the underlying authentication system.

> **Note:**
>
> Microsoft has implemented certain variations of the SPNEGO protocol, hence to inter-operate with Microsoft, we have added a separate mode via a new system property `sun.security.spnego.msinterop`. This property is enabled to true by default. To disable it, you need to explicitly set this property to false. To enable SPNEGO debugging, you can set the system property `sun.security.spnego.debug=true`.

Part VI: HTTP/SPNEGO Authentication

Exercise 9: Using HTTP/SPNEGO Authentication

What is HTTP SPNEGO

HTTP SPNEGO supports the Negotiate authentication scheme in an HTTP communication. SPNEGO supports types of authentication:

Web Authentication

The Web Server responds with

```
HTTP/1.1 401 Unauthorized
WWW-Authenticate: Negotiate
```

the client will need to send a header like

```
Authorization: Negotiate YY.....
```

to authenticate itself to the server

Proxy Authentication

The Web Server responses with

```
HTTP/1.1 407 Proxy Authentication Required
Proxy-Authenticate: Negotiate
```

the client will need to send a header like

```
Proxy-Authorization: Negotiate YY.....
```

to authenticate itself to the proxy server.

This feature supports both types of authentication.

How to use HTTP/SPNEGO Authentication

There is no new public API function involved in the new feature, but several configurations are needed to perform a success communication:

Kerberos 5 Configuration

Since the SPNEGO mechanism will call JGSS, which in turns calls the Kerberos V5 login module to do real works. Kerberos 5 configurations are needed. This includes the following:

- Some way to provide Kerberos configurations. This can be achieved with the Java system property `java.security.krb5.conf`. For example:

  ```
  java -Djava.security.krb5.conf=krb5.conf \
       -Djavax.security.auth.useSubjectCredsOnly=false \
       ClassName
  ```

  A JAAS config file denoting what login module to use. HTTP SPNEGO codes will look for the standard entry named `com.sun.security.jgss.krb5.initiate`.

For example, you can provide a file `spnegoLogin.conf`:

```
com.sun.security.jgss.krb5.initiate {
    com.sun.security.auth.module.Krb5LoginModule
        required useTicketCache=true;
};
```

and run java with:

```
java -Djava.security.krb5.conf=krb5.conf \
     -Djava.security.auth.login.config=spnegoLogin.conf \
     -Djavax.security.auth.useSubjectCredsOnly=false \
     ClassName
```

User Name and Password Retrieval

Just like any other HTTP authentication scheme, the client can provide a customized `java.net.Authenticator` to feed user name and password to the HTTP SPNEGO module **if** they are needed (i.e. there is no credential cache available). The only authentication information needed to be checked in your `Authenticator` is the scheme which can be retrieved with `getRequestingScheme()`. The value should be "Negotiate".

This means your `Authenticator` implementation will look like:

```
class MyAuthenticator extends Authenticator {

        public PasswordAuthentication getPasswordAuthentication () {
            if (getRequestingScheme().equalsIgnoreCase("negotiate")) {
                String krb5user;
                char[] krb5pass;
                // get krb5user and krb5pass in your own way
                ....
                return (new PasswordAuthentication (krb5user,
                            krb5pass));
            } else {
                ....
            }
        }
    }
```

> **Note:**
>
> According to the specification of `java.net.Authenticator`, it's designed to get the user name and password at the same time, so do not specify `principal=`*xxx* in the JAAS config file.

Scheme Preference

The client can still provide system property `http.auth.preference` to denote that a certain scheme should always be used as long as the server request for it. You can use "SPNEGO" or "Kerberos" for this system property. "SPNEGO" means you prefer

to response the Negotiate scheme using the GSS/SPNEGO mechanism; "Kerberos" means you prefer to response the Negotiate scheme using the GSS/Kerberos mechanism. Normally, when authenticating against a Microsoft product, you can use "SPNEGO". The value "Kerberos" also works for Microsoft servers. It's only needed when you encounter a server which knows Negotiate but doesn't know about SPNEGO.

If `http.auth.preference` is not set, the internal order chosen is:

- GSS/SPNEGO -> Digest -> NTLM -> Basic

Notice that Kerberos does not appear in this list, since whenever Negotiate is supported, GSS/SPNEGO is always chosen.

Fallback

If the server has provided more than one authentication scheme (including Negotiate), according to the processing order mentioned in the last section, Java will try to challenge the Negotiate scheme. However, if the protocol cannot be established successfully (for example, the Kerberos configuration is not correct, or the server's hostname is not recorded in the KDC principal DB, or the user name and password provided by `Authenticator` is wrong), then the second strongest scheme will be automatically used.

> **Note:**
>
> If `http.auth.preference` is set to SPNEGO or Kerberos, then SPNEGO assumes you only want to try the Negotiate scheme even if it fails. SPNEGO will not fallback to any other scheme and your program will throw an `IOException` saying it received a 401 or 407 error from the HTTP response.

HTTP/SPNEGO Authentication Example

Assume that you have an IIS Server running on a Windows Server within an Active Directory. A web page on this server is configured to be protected by Integrated Windows Authentication. This means the server will prompt for both Negotiate and NTLM authentication.

You need to prepare these files to get the protected file:

RunHttpSpnego.java

```
import java.io.BufferedReader;
import java.io.InputStream;
import java.io.InputStreamReader;
import java.net.Authenticator;
import java.net.PasswordAuthentication;
import java.net.URL;

public class RunHttpSpnego {

    static final String kuser = "username"; // your account name
```

```
    static final String kpass = "password"; // your password for the
account

    static class MyAuthenticator extends Authenticator {
        public PasswordAuthentication getPasswordAuthentication() {
            // I haven't checked getRequestingScheme() here, since for NTLM
            // and Negotiate, the usrname and password are all the same.
            System.err.println("Feeding username and password for " +
getRequestingScheme());
            return (new PasswordAuthentication(kuser,
kpass.toCharArray()));
        }
    }

    public static void main(String[] args) throws Exception {
        Authenticator.setDefault(new MyAuthenticator());
        URL url = new URL(args[0]);
        InputStream ins = url.openConnection().getInputStream();
        BufferedReader reader = new BufferedReader(new
InputStreamReader(ins));
        String str;
        while((str = reader.readLine()) != null)
            System.out.println(str);
    }
}
```

krb.conf

```
[libdefaults]
    default_realm = AD.LOCAL
[realms]
    AD.LOCAL = {
        kdc = kdc.ad.local
    }
```

login.conf

```
com.sun.security.jgss.krb5.initiate {
  com.sun.security.auth.module.Krb5LoginModule required  doNotPrompt=false
useTicketCache=true;
};
```

Compiling and Running the Example

1. Compile `RunHttpSpnego.java`.
2. Run `RunHttpSpnego.java`:

```
java -Djava.security.krb5.conf=krb5.conf \
    -Djava.security.auth.login.config=login.conf \
    -Djavax.security.auth.useSubjectCredsOnly=false \
    RunHttpSpnego \
    http://www.ad.local/hello/hello.html
```

You will see the following:

```
Feeding username and password for Negotiate
<h1>Hello, You got me!</h1>
```

In fact, if you are running on a Windows machine as a domain user, or, you are running on a Linux or Solaris machine that has already issued the kinit command and got the credential cache. The class `MyAuthenticator` will be completely ignored, and the output will be simply:

```
<h1>Hello, You got me!</h1>
```

which shows the user name and password are not consulted. This is the so-called Single Sign-On.

Also, you can just run

```
java RunHttpSpnego http://www.ad.local/hello/hello.html
```

to see how the fallback is done, in which case you will see

```
Feeding username and password for ntlm
<h1>Hello, You got me!</h1>
```

Source Code for Advanced Security Programming in Java SE Authentication, Secure Communication and Single Sign-On

Jass.java

```
import javax.security.auth.Subject;
import javax.security.auth.login.*;
import javax.security.auth.callback.CallbackHandler;
import java.security.*;
import com.sun.security.auth.callback.TextCallbackHandler;
import java.io.File;

public class Jaas {
    private static String name;
    private static final boolean verbose = false;

    public static void main(String[] args) throws Exception {
        if (args.length > 0) {
            name = args[0];
        } else {
            name = "client";
        }

        // Create action to perform
        PrivilegedExceptionAction action = new MyAction();
```

```
        loginAndAction(name, action);
    }

    static void loginAndAction(String name, PrivilegedExceptionAction
action)
        throws LoginException, PrivilegedActionException {

        // Create a callback handler
        CallbackHandler callbackHandler = new TextCallbackHandler();

        LoginContext context = null;

        try {
            // Create a LoginContext with a callback handler
            context = new LoginContext(name, callbackHandler);

            // Perform authentication
            context.login();
        } catch (LoginException e) {
            System.err.println("Login failed");
            e.printStackTrace();
            System.exit(-1);
        }

        // Perform action as authenticated user
        Subject subject = context.getSubject();
        if (verbose) {
            System.out.println(subject.toString());
        } else {
            System.out.println("Authenticated principal: " +
                subject.getPrincipals());
        }

        Subject.doAs(subject, action);

        context.logout();
    }

    // Action to perform
    static class MyAction implements PrivilegedExceptionAction {
        MyAction() {
        }

        public Object run() throws Exception {
            // Replace the following with an action to be performed
            // by authenticated user
            System.out.println("Performing secure action ...");
            return null;
        }
    }
}
```

jass-krb5.conf

```
client {
    com.sun.security.auth.module.Krb5LoginModule required
    principal="test";
};

server {
    com.sun.security.auth.module.Krb5LoginModule required
    useKeyTab=true
    storeKey=true
    keyTab=sample.keytab
    principal="host/machineName";
};
```

AppConnection.java

```
import java.io.*;
import java.net.Socket;

class AppConnection {
    public static final int AUTH_CMD = 100;
    public static final int DATA_CMD = 200;

    public static final int SUCCESS = 0;
    public static final int AUTH_INPROGRESS = 1;
    public static final int FAILURE = 2;

    private DataInputStream inStream;
    private DataOutputStream outStream;
    private Socket socket;

    // Client application
    AppConnection(String hostName, int port) throws IOException {
        socket = new Socket(hostName, port);

        inStream = new DataInputStream(socket.getInputStream());
        outStream = new DataOutputStream(socket.getOutputStream());

        System.out.println("Connected to address " +
            socket.getInetAddress());
    }

    // Server side application
    AppConnection(Socket socket) throws IOException {
        this.socket = socket;
        inStream = new DataInputStream(socket.getInputStream());
        outStream = new DataOutputStream(socket.getOutputStream());

        System.out.println("Got connection from client " +
            socket.getInetAddress());
    }
```

```
byte[] receive(int expected) throws IOException {
    if (expected != -1) {
        int cmd = inStream.readInt();
        if (expected != cmd) {
            throw new IOException("Received unexpected code: " + cmd);
        }
        //System.out.println("Read cmd: " + cmd);
    }

    byte[] reply = null;
    int len;
    try {
        len = inStream.readInt();
        //System.out.println("Read length: " + len);

    } catch (IOException e) {
        len = 0;
    }
    if (len > 0) {
        reply = new byte[len];
        inStream.readFully(reply);
    } else {
        reply = new byte[0];
    }
    return reply;
}

AppReply send(int cmd, byte[] bytes) throws IOException {
    //System.out.println("Write cmd: " + cmd);
    outStream.writeInt(cmd);
    if (bytes != null) {
        //System.out.println("Write length: " + bytes.length);
        outStream.writeInt(bytes.length);
        if (bytes.length > 0) {
            outStream.write(bytes);
        }
    } else {
        //System.out.println("Write length: " + 0);
        outStream.writeInt(0);
    }

    outStream.flush();

    if (cmd == SUCCESS || cmd == FAILURE) {
        return null;   // Done
    }

    int returnCode = inStream.readInt();
    //System.out.println("Read cmd: " + returnCode);

    byte[] reply = null;
    if (returnCode != FAILURE) {
        reply = receive(-1);
    }
    return new AppReply(returnCode, reply);
```

```
    }

    static class AppReply {
        private int code;
        private byte[] bytes;

        AppReply(int code, byte[] bytes) {
            this.bytes = bytes;
            this.code = code;
        }

        int getStatus() {
            return code;
        }

        byte[] getBytes() {
            return bytes;
        }
    }

    void close() {
        try {
            socket.close();
        } catch (IOException e) {
        }
    }
}
```

GssServer.java

```
import org.ietf.jgss.*;
import java.io.*;
import java.net.Socket;
import java.net.ServerSocket;
import java.security.*;
import java.util.Date;

/**
 * A sample server application that uses JGSS to do mutual authentication
 * with a client using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the client.
 *
 * Every message exchanged with the client includes a 4-byte application-
 * level header that contains the big-endian integer value for the number
 * of bytes that will follow as part of the JGSS token.
 *
 * The protocol is:
 *    1.  Context establishment loop:
 *         a. client sends init sec context token to server
 *         b. server sends accept sec context token to client
 *         ....
 *    2. client sends a wrap token to the server.
 *    3. server sends a wrap token back to the client.
 *
```

```
 * Start GssServer first before starting GssClient.
 *
 * Usage:  java <options> GssServer
 *
 * Example: java -Djava.security.auth.login.config=jaas-krb5.conf \
 *               GssServer
 *
 * Add -Djava.security.krb5.conf=krb5.conf to specify application-specific
 * Kerberos configuration (different from operating system's Kerberos
 * configuration).
 */

public class GssServer  {
    private static final int PORT = 4567;
    private static final boolean verbose = false;
    private static final int LOOP_LIMIT = 1;
    private static int loopCount = 0;

    public static void main(String[] args) throws Exception {

        PrivilegedExceptionAction action = new GssServerAction(PORT);

        Jaas.loginAndAction("server", action);
    }

    static class GssServerAction implements PrivilegedExceptionAction {
        private int localPort;

        GssServerAction(int port) {
            this.localPort = port;
        }

        public Object run() throws Exception {

            ServerSocket ss = new ServerSocket(localPort);

            // Get own Kerberos credentials for accepting connection
            GSSManager manager = GSSManager.getInstance();
            Oid krb5Mechanism = new Oid("1.2.840.113554.1.2.2");
            GSSCredential serverCreds = manager.createCredential(null,

GSSCredential.DEFAULT_LIFETIME,
                                                krb5Mechanism,
                                                GSSCredential.ACCEPT_ONLY);
            while (loopCount++ < LOOP_LIMIT) {

                System.out.println("Waiting for incoming connection...");

                Socket socket = ss.accept();
                DataInputStream inStream =
                    new DataInputStream(socket.getInputStream());

                DataOutputStream outStream =
                    new DataOutputStream(socket.getOutputStream());
```

```
                System.out.println("Got connection from client " +
                    socket.getInetAddress());

                /*
                 * Create a GSSContext to receive the incoming request
                 * from the client. Use null for the server credentials
                 * passed in. This tells the underlying mechanism
                 * to use whatever credentials it has available that
                 * can be used to accept this connection.
                 */

                GSSContext context = manager.createContext(
                    (GSSCredential)serverCreds);

                // Do the context establishment loop

                byte[] token = null;

                while (!context.isEstablished()) {

                    if (verbose) {
                        System.out.println("Reading ...");
                    }
                    token = new byte[inStream.readInt()];

                    if (verbose) {
                        System.out.println("Will read input token of size
" +
                            token.length + " for processing by
acceptSecContext");
                    }
                    inStream.readFully(token);

                    if (token.length == 0) {
                        if (verbose) {
                            System.out.println("skipping zero length
token");
                        }
                        continue;
                    }
                    if (verbose) {
                        System.out.println("Token = " +
getHexBytes(token));
                        System.out.println("acceptSecContext..");
                    }
                    token = context.acceptSecContext(token, 0,
token.length);

                    // Send a token to the peer if one was generated by
                    // acceptSecContext
                    if (token != null) {
                        if (verbose) {
                            System.out.println("Will send token of size " +
                                token.length + " from acceptSecContext.");
                        }
```

```
                        outStream.writeInt(token.length);
                        outStream.write(token);
                        outStream.flush();
                    }
                }

                System.out.println("Context Established! ");
                System.out.println("Client principal is " +
context.getSrcName());
                System.out.println("Server principal is " +
context.getTargName());

                /*
                 * If mutual authentication did not take place, then
                 * only the client was authenticated to the
                 * server. Otherwise, both client and server were
                 * authenticated to each other.
                 */
                if (context.getMutualAuthState())
                    System.out.println("Mutual authentication took
place!");

                /*
                 * Create a MessageProp which unwrap will use to return
                 * information such as the Quality-of-Protection that was
                 * applied to the wrapped token, whether or not it was
                 * encrypted, etc. Since the initial MessageProp values
                 * are ignored, just set them to the defaults of 0 and
false.
                 */
                MessageProp prop = new MessageProp(0, false);

                /*
                 * Read the token. This uses the same token byte array
                 * as that used during context establishment.
                 */
                token = new byte[inStream.readInt()];
                if (verbose) {
                    System.out.println("Will read token of size " +
token.length);
                }
                inStream.readFully(token);

                byte[] input = context.unwrap(token, 0, token.length,
prop);
                String str = new String(input, "UTF-8");

                System.out.println("Received data \"" +
                    str + "\" of length " + str.length());

                System.out.println("Confidentiality applied: " +
                    prop.getPrivacy());

                /*
```

```
                 * Now generate reply that is the concatenation of the
                 * incoming string with the current time.
                 */

                /*
                 * First reset the QOP of the MessageProp to 0
                 * to ensure the default Quality-of-Protection
                 * is applied.
                 */
                prop.setQOP(0);

                String now = new Date().toString();
                byte[] nowBytes = now.getBytes("UTF-8");
                int len = input.length + 1 + nowBytes.length;
                byte[] reply = new byte[len];
                System.arraycopy(input, 0, reply, 0, input.length);
                reply[input.length] = ' ';
                System.arraycopy(nowBytes, 0, reply, input.length+1,
                    nowBytes.length);

                System.out.println("Sending: " + new String(reply,
"UTF-8"));
                token = context.wrap(reply, 0, reply.length, prop);

                outStream.writeInt(token.length);
                outStream.write(token);
                outStream.flush();

                System.out.println("Closing connection with client " +
                    socket.getInetAddress());
                context.dispose();
                socket.close();
            }
            return null;
        }
    }

    private static final String getHexBytes(byte[] bytes, int pos, int
len) {

        StringBuffer sb = new StringBuffer();
        for (int i = pos; i < (pos+len); i++) {

            int b1 = (bytes[i]>>4) & 0x0f;
            int b2 = bytes[i] & 0x0f;

            sb.append(Integer.toHexString(b1));
            sb.append(Integer.toHexString(b2));
            sb.append(' ');
        }
        return sb.toString();
    }

    private static final String getHexBytes(byte[] bytes) {
        return getHexBytes(bytes, 0, bytes.length);
```

```
    }
}
```

GssClient.java

```
import org.ietf.jgss.*;
import java.net.Socket;
import java.io.IOException;
import java.io.DataInputStream;
import java.io.DataOutputStream;
import java.security.*;
import javax.security.auth.login.LoginException;

/**
 * A sample client application that uses JGSS to do mutual authentication
 * with a server using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the server.
 *
 * Every message sent to the server includes a 4-byte application-level
 * header that contains the big-endian integer value for the number
 * of bytes that will follow as part of the JGSS token.
 *
 * The protocol is:
 *    1.  Context establishment loop:
 *         a. client sends init sec context token to server
 *         b. server sends accept sec context token to client
 *         ....
 *    2. client sends a wrapped token to the server.
 *    3. server sends a wrapped token back to the client for the
application
 *
 * Start GssServer first before starting GssClient.
 *
 * Usage:  java <options> GssClient <service> <serverName>
 *
 * Example: java -Djava.security.auth.login.config=jaas-krb5.conf \
 *               GssClient host machine.imc.org
 *
 * Add -Djava.security.krb5.conf=krb5.conf to specify application-specific
 * Kerberos configuration (different from operating system's Kerberos
 * configuration).
 */

public class GssClient {
    private static final int PORT = 4567;
    private static final boolean verbose = false;

    public static void main(String[] args) throws Exception {

        // Obtain the command-line arguments and parse the server's
principal

        if (args.length < 2) {
            System.err.println(
```

```
            "Usage: java <options> GssClient <service> <serverName>");
        System.exit(-1);
    }

    String serverPrinc = args[0] + "@" + args[1];

    PrivilegedExceptionAction action =
        new GssClientAction(serverPrinc, args[1], PORT);

    Jaas.loginAndAction("client", action);
}

static class GssClientAction implements PrivilegedExceptionAction {
    private String serverPrinc;
    private String hostName;
    private int port;

    GssClientAction(String serverPrinc, String hostName, int port) {
        this.serverPrinc = serverPrinc;
        this.hostName = hostName;
        this.port = port;
    }

    public Object run() throws Exception {
        Socket socket = new Socket(hostName, port);
        DataInputStream inStream =
            new DataInputStream(socket.getInputStream());
        DataOutputStream outStream =
            new DataOutputStream(socket.getOutputStream());

        System.out.println("Connected to address " +
            socket.getInetAddress());

        /*
         * This Oid is used to represent the Kerberos version 5 GSS-API
         * mechanism. It is defined in RFC 1964. We will use this Oid
         * whenever we need to indicate to the GSS-API that it must
         * use Kerberos for some purpose.
         */
        Oid krb5Oid = new Oid("1.2.840.113554.1.2.2");

        GSSManager manager = GSSManager.getInstance();

        /*
         * Create a GSSName out of the server's name.
         */
        GSSName serverName = manager.createName(serverPrinc,
            GSSName.NT_HOSTBASED_SERVICE);

        /*
         * Create a GSSContext for mutual authentication with the
         * server.
         *    - serverName is the GSSName that represents the server.
         *    - krb5Oid is the Oid that represents the mechanism to
         *      use. The client chooses the mechanism to use.
```

```
             *    - null is passed in for client credentials
             *    - DEFAULT_LIFETIME lets the mechanism decide how long the
             *      context can remain valid.
             * Note: Passing in null for the credentials asks GSS-API to
             * use the default credentials. This means that the mechanism
             * will look among the credentials stored in the current
Subject
             * to find the right kind of credentials that it needs.
             */
            GSSContext context = manager.createContext(serverName,
                krb5Oid,
                null,
                GSSContext.DEFAULT_LIFETIME);

            // Set the desired optional features on the context. The client
            // chooses these options.

            context.requestMutualAuth(true);  // Mutual authentication
            context.requestConf(true);  // Will use confidentiality later
            context.requestInteg(true); // Will use integrity later

            // Do the context eastablishment loop

            byte[] token = new byte[0];

            while (!context.isEstablished()) {

                // token is ignored on the first call
                token = context.initSecContext(token, 0, token.length);

                // Send a token to the server if one was generated by
                // initSecContext
                if (token != null) {
                    if (verbose) {
                        System.out.println("Will send token of size " +
                            token.length + " from initSecContext.");
                        System.out.println("writing token = " +
                            getHexBytes(token));
                    }

                    outStream.writeInt(token.length);
                    outStream.write(token);
                    outStream.flush();
                }

                // If the client is done with context establishment
                // then there will be no more tokens to read in this loop
                if (!context.isEstablished()) {
                    token = new byte[inStream.readInt()];
                    if (verbose) {
                        System.out.println("reading token = " +
                            getHexBytes(token));
                        System.out.println("Will read input token of size
" +
                            token.length + " for processing by
```

```
initSecContext");
                    }
                    inStream.readFully(token);
                }
            }

            System.out.println("Context Established! ");
            System.out.println("Client principal is " +
context.getSrcName());
            System.out.println("Server principal is " +
context.getTargName());

            /*
             * If mutual authentication did not take place, then only the
             * client was authenticated to the server. Otherwise, both
             * client and server were authenticated to each other.
             */
            if (context.getMutualAuthState())
                System.out.println("Mutual authentication took place!");

            byte[] messageBytes = "Hello There!".getBytes("UTF-8");

            /*
             * The first MessageProp argument is 0 to request
             * the default Quality-of-Protection.
             * The second argument is true to request
             * privacy (encryption of the message).
             */
            MessageProp prop =  new MessageProp(0, true);

            /*
             * Encrypt the data and send it across. Integrity protection
             * is always applied, irrespective of confidentiality
             * (i.e., encryption).
             * You can use the same token (byte array) as that used when
             * establishing the context.
             */

            System.out.println("Sending message: " +
                new String(messageBytes, "UTF-8"));
            token = context.wrap(messageBytes, 0, messageBytes.length,
prop);
            outStream.writeInt(token.length);
            outStream.write(token);
            outStream.flush();

            /*
             * Now we will allow the server to decrypt the message,
             * append a time/date on it, and send then it back.
             */

            token = new byte[inStream.readInt()];
            System.out.println("Will read token of size " + token.length);
            inStream.readFully(token);
            byte[] replyBytes = context.unwrap(token, 0, token.length,
```

```
prop);

            System.out.println("Received message: " +
                new String(replyBytes, "UTF-8"));

            System.out.println("Done.");
            context.dispose();
            socket.close();

            return null;
        }
    }

    private static final String getHexBytes(byte[] bytes, int pos, int
len) {

        StringBuffer sb = new StringBuffer();
        for (int i = pos; i < (pos+len); i++) {

            int b1 = (bytes[i]>>4) & 0x0f;
            int b2 = bytes[i] & 0x0f;

            sb.append(Integer.toHexString(b1));
            sb.append(Integer.toHexString(b2));
            sb.append(' ');
        }
        return sb.toString();
    }

    private static final String getHexBytes(byte[] bytes) {
        return getHexBytes(bytes, 0, bytes.length);
    }
}
```

SaslTestServer.java

```
import javax.security.sasl.*;
import javax.security.auth.callback.*;
import java.security.*;
import java.util.HashMap;
import java.net.*;
import java.util.Date;

/**
 * A sample server application that uses SASL to authenticate clients
 * using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the client.
 *
 * This sample program uses a ficticious application-level protocol.
 * Every message exchanged between the client and server an 8-byte
 * header that consists of two integers: the first integer represesents
 * the application-level command or status code while the second integer
 * indicates the length of the SASL buffer. This header is followed by
 * the SASL buffer.
```

```
 *
 * The protocol is:
 *    1.  Authentication
 *         a. client sends initial response to server containing
authentication
 *            information
 *         b. server accepts and evaluates response to generate challenge;
it
 *            sends the challenge to the server.
 *         c. client evaluates challenge to generate response; it sends the
 *            response;
 *         d. Steps b and c are repeated until authentication succeeds or
fails.
 *    2. client sends an encrypted message to the server.
 *    3. server decryptes the message and sends an encrypted one back
 *       that contains the original message plus the current time.
 *
 * Start SaslTestServer first before starting SaslTestClient.
 *
 * Usage:  java <options> SaslTestServer service serverName
 *
 * Example: java -Djava.security.auth.login.config=jaas-krb5.conf \
 *               SaslTestServer host machine.imc.org
 *
 * Add -Djava.security.krb5.conf=krb5.conf to specify application-specific
 * Kerberos configuration (different from operating system's Kerberos
 * configuration).
 */


public class SaslTestServer {
    private static final String MECH = "GSSAPI"; // SASL name for GSS-API/
Kerberos
    private static final int PORT = 4568;
    private static final int LOOP_LIMIT = 1;
    private static int loopCount = 0;

    public static void main(String[] args) throws Exception {
        // Obtain the command-line arguments and parse the server's
principal

        if (args.length < 2) {
            System.err.println(
                "Usage: java <options> SaslTestServer <service> <host>");
            System.exit(-1);
        }

        PrivilegedExceptionAction action =
            new SaslServerAction(args[0], args[1], PORT);

        Jaas.loginAndAction("server", action);
    }

    static class SaslServerAction implements PrivilegedExceptionAction {
        private String service;      // used for SASL authentication
```

```
        private String serverName;   // named used for SASL authentication
        private int localPort;
        private CallbackHandler cbh = new TestCallbackHandler();

        SaslServerAction(String service, String serverName, int port) {
            this.service = service;
            this.serverName = serverName;
            this.localPort = port;
        }

        public Object run() throws Exception {
            ServerSocket ss = new ServerSocket(localPort);

            HashMap<String,Object> props = new HashMap<String,Object>();
            props.put(Sasl.QOP, "auth-conf,auth-int,auth");

            // Loop, accepting requests from any client
            while (loopCount++ < LOOP_LIMIT) {
                System.out.println("Waiting for incoming connection...");
                Socket socket = ss.accept();

                // Create application-level connection to handle request
                AppConnection conn = new AppConnection(socket);

                // Normally, the application protocol will negotiate which
                // SASL mechanism to use. In this simplified example, we
                // will always use "GSSAPI", the name of the mechanism
that does
                // Kerberos via GSS-API

                // Create SaslServer to perform authentication
                SaslServer srv = Sasl.createSaslServer(MECH,
                    service, serverName, props, cbh);

                if (srv == null) {
                    throw new Exception(
                        "Unable to find server implementation for " +
MECH);
                }

                boolean auth = false;

                // Read initial response from client
                byte[] response = conn.receive(AppConnection.AUTH_CMD);
                AppConnection.AppReply clientMsg;

                while (!srv.isComplete()) {
                    try {
                        // Generate challenge based on response
                        byte[] challenge = srv.evaluateResponse(response);

                        if (srv.isComplete()) {
                            conn.send(AppConnection.SUCCESS, challenge);
                            auth = true;
                        } else {
```

```
                            clientMsg =
conn.send(AppConnection.AUTH_INPROGRESS,
                                challenge);
                            response = clientMsg.getBytes();
                        }
                    } catch (SaslException e) {
                        // e.printStackTrace();
                        // Send failure notification to client
                        conn.send(AppConnection.FAILURE, null);
                        break;
                    }
                }

                // Check status of authentication
                if (srv.isComplete() && auth) {
                    System.out.print("Client authenticated; ");
                    System.out.println("authorized client is: " +
                        srv.getAuthorizationID());
                } else {
                    // Go get another client
                    System.out.println("Authentication failed. ");
                    continue;
                }

                String qop = (String) srv.getNegotiatedProperty(Sasl.QOP);
                System.out.println("Negotiated QOP: " + qop);

                // Now try to use security layer
                boolean sl = (qop.equals("auth-conf") || qop.equals("auth-
int"));

                byte[] msg = conn.receive(AppConnection.DATA_CMD);
                byte[] realMsg = (sl ? srv.unwrap(msg, 0, msg.length) :
msg);

                System.out.println("Received: " + new String(realMsg,
"UTF-8"));

                // Construct reply to send to client
                String now = new Date().toString();
                byte[] nowBytes = now.getBytes("UTF-8");
                int len = realMsg.length + 1 + nowBytes.length;
                byte[] reply = new byte[len];
                System.arraycopy(realMsg, 0, reply, 0, realMsg.length);
                reply[realMsg.length] = ' ';
                System.arraycopy(nowBytes, 0, reply, realMsg.length+1,
                    nowBytes.length);

                System.out.println("Sending: " + new String(reply,
"UTF-8"));

                byte[] realReply = (sl ? srv.wrap(reply, 0,
reply.length) : reply);

                conn.send(AppConnection.SUCCESS, realReply);
```

```
            }
            return null;
        }
    }

    static class TestCallbackHandler implements CallbackHandler {

        public void handle(Callback[] callbacks)
            throws UnsupportedCallbackException {

            AuthorizeCallback acb = null;

            for (int i = 0; i < callbacks.length; i++) {
                if (callbacks[i] instanceof AuthorizeCallback) {
                    acb = (AuthorizeCallback) callbacks[i];
                } else {
                    throw new UnsupportedCallbackException(callbacks[i]);
                }
            }

            if (acb != null) {
                String authid = acb.getAuthenticationID();
                String authzid = acb.getAuthorizationID();
                if (authid.equals(authzid)) {
                    // Self is always authorized
                    acb.setAuthorized(true);

                } else {
                    // Should check some database for mapping and decide.
                    // Current simplified policy is to reject authzids that
                    // don't match authid

                    acb.setAuthorized(false);
                }

                if (acb.isAuthorized()) {
                    // Set canonicalized name.
                    // Should look up database for canonical names

                    acb.setAuthorizedID(authzid);
                }
            }
        }
    }
}
```

SaslTestClient.java

```
import javax.security.sasl.*;
import javax.security.auth.callback.*;
import java.security.*;
import javax.security.auth.Subject;
import javax.security.auth.login.*;
import com.sun.security.auth.callback.*;
```

```
import java.util.HashMap;

/**
 * A sample client application that uses SASL to authenticate to
 * a server using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the server.
 *
 * This sample program uses a ficticious application-level protocol.
 * Every message exchanged between the client and server an 8-byte
 * header that consists of two integers: the first integer represesents
 * the application-level command or status code while the second integer
 * indicates the length of the SASL buffer. This header is followed by
 * the SASL buffer.
 *
 * The protocol is:
 *    1.  Authentication
 *         a. client sends initial response to server containing
authentication
 *            information
 *         b. server accepts and evaluates response to generate challenge;
it
 *            sends the challenge to the server.
 *         c. client evaluates challenge to generate response; it sends the
 *            response;
 *         d. Steps b and c are repeated until authentication succeeds or
fails.
 *    2. client sends an encrypted message to the server.
 *    3. server decryptes the message and sends an encrypted one back
 *       that contains the original message plus the current time.
 *
 * Start SaslTestServer first before starting SaslTestClient.
 *
 * Usage:  java <options> SaslTestClient service serverName
 *
 * Example: java -Djava.security.auth.login.config=jaas-krb5.conf \
 *               SaslTestClient host machine.imc.org
 *
 * Add -Djava.security.krb5.conf=krb5.conf to specify application-specific
 * Kerberos configuration (different from operating system's Kerberos
 * configuration).
 */

public class SaslTestClient {
    private static final String MECH = "GSSAPI"; // SASL name for GSS-API/
Kerberos
    private static final int PORT = 4568;

    private static final byte[] EMPTY = new byte[0];

    public static void main(String[] args) throws Exception {
        // Obtain the command-line arguments and parse the server's
principal

        if (args.length < 2) {
            System.err.println(
```

```
                "Usage: java <options> SaslTestClient <service>
<serverName>");
            System.exit(-1);
        }

        PrivilegedExceptionAction action =
            new SaslClientAction(args[0], args[1], PORT);

        Jaas.loginAndAction("client", action);
    }

    static class SaslClientAction implements PrivilegedExceptionAction {
        private String service;      // used for SASL authentication
        private String serverName;   // name used for SASL authentication
        private int port;
        private CallbackHandler cbh = null; // Don't need handler for
GSSAPI

        SaslClientAction(String service, String serverName, int port) {
            this.service = service;
            this.serverName = serverName;
            this.port = port;
        }

        public Object run() throws Exception {
            // Create application-level connection
            AppConnection conn = new AppConnection(serverName, port);

            HashMap<String,Object> props = new HashMap<String,Object>();
            // Request confidentiality
            props.put(Sasl.QOP, "auth-conf");

            // Create SaslClient to perform authentication
            SaslClient clnt = Sasl.createSaslClient(
                new String[]{MECH}, null, service, serverName, props, cbh);

            if (clnt == null) {
                throw new Exception(
                    "Unable to find client implementation for " + MECH);
            }

            byte[] response;
            byte[] challenge;

            // Get initial response for authentication
            response = clnt.hasInitialResponse() ?
                clnt.evaluateChallenge(EMPTY) : EMPTY;

            // Send initial response to server
            AppConnection.AppReply reply =
                conn.send(AppConnection.AUTH_CMD, response);

            // Repeat until authentication terminates
            while (!clnt.isComplete() &&
                (reply.getStatus() == AppConnection.AUTH_INPROGRESS ||
```

```
                 reply.getStatus() == AppConnection.SUCCESS)) {

                // Evaluate challenge to generate response
                challenge = reply.getBytes();
                response = clnt.evaluateChallenge(challenge);

                if (reply.getStatus() == AppConnection.SUCCESS) {
                    if (response != null) {
                        throw new Exception("Protocol error interacting
with SASL");
                    }
                    break;
                }

                // Send response to server and read server's next challenge
                reply = conn.send(AppConnection.AUTH_CMD, response);
            }

            // Check status of authentication
            if (clnt.isComplete() && reply.getStatus() ==
AppConnection.SUCCESS) {
                System.out.println("Client authenticated.");
            } else {
                throw new Exception("Authentication failed: " +
                    " connection status? " + reply.getStatus());
            }

            String qop = (String) clnt.getNegotiatedProperty(Sasl.QOP);
            System.out.println("Negotiated QOP: " + qop);

            // Try out security layer
            boolean sl = (qop.equals("auth-conf") || qop.equals("auth-
int"));

            byte[] msg = "Hello There!".getBytes("UTF-8");
            System.out.println("Sending: " + new String(msg, "UTF-8"));

            byte[] encrypted = (sl ? clnt.wrap(msg, 0, msg.length) : msg);

            reply = conn.send(AppConnection.DATA_CMD, encrypted);

            if (reply.getStatus() == AppConnection.SUCCESS) {
                byte[] encryptedReply = reply.getBytes();

                byte[] clearReply = (sl ? clnt.unwrap(encryptedReply,
                    0, encryptedReply.length) : encryptedReply);

                System.out.println("Received: " + new String(clearReply,
"UTF-8"));
            } else {
                System.out.println("Failed exchange: " +
reply.getStatus());
            }

            conn.close();
```

```
            return null;
        }
    }
}
```

JsseServer.java

```
import java.io.*;
import java.net.*;
import javax.net.ssl.*;
import java.util.Date;
import java.security.PrivilegedExceptionAction;
import java.security.Principal;

/*
 * Tests support for RFC 2712. Specify use of only a KRB5 cipher for both
 * client and server, by first doing a JAAS login for the server
 * without first doing a JAAS login for the client.
 */

public class JsseServer {

    private static final String KRB5_CIPHER =
"TLS_KRB5_WITH_3DES_EDE_CBC_SHA";


    private static final int PORT = 4569;
    private static final boolean verbose = false;
    private static final int LOOP_LIMIT = 1;
    private static int loopCount = 0;

    public static void main(String[] args) throws Exception {

        PrivilegedExceptionAction action = new JsseServerAction(PORT);

        Jaas.loginAndAction("server", action);
    }

    static class JsseServerAction implements PrivilegedExceptionAction {
        private int localPort;

        JsseServerAction(int port) {
            this.localPort = port;
        }

        public Object run() throws Exception {

            SSLServerSocketFactory sslssf =
                (SSLServerSocketFactory)
SSLServerSocketFactory.getDefault();
            SSLServerSocket sslServerSocket =
                (SSLServerSocket) sslssf.createServerSocket(localPort);
```

```
            // Enable only a KRB5 cipher suite.
            String enabledSuites[] = { KRB5_CIPHER };
            sslServerSocket.setEnabledCipherSuites(enabledSuites);
            // Should check for exception if enabledSuites is not supported

            while (loopCount++ < LOOP_LIMIT) {
                System.out.println("Waiting for incoming connection...");

                SSLSocket sslSocket = (SSLSocket) sslServerSocket.accept();

                System.out.println("Got connection from client " +
                    sslSocket.getInetAddress());

                BufferedReader in = new BufferedReader(new
InputStreamReader(
                    sslSocket.getInputStream()));
                BufferedWriter out = new BufferedWriter(new
OutputStreamWriter(
                    sslSocket.getOutputStream()));

                String inStr = in.readLine();
                System.out.println("Received " + inStr);

                String outStr = inStr + " " + new Date().toString() + "\n";
                out.write(outStr);
                System.out.println("Sending " + outStr);
                out.flush();

                String cipherSuiteChosen =
sslSocket.getSession().getCipherSuite();
                System.out.println("Cipher suite in use: " +
cipherSuiteChosen);
                Principal self =
sslSocket.getSession().getLocalPrincipal();
                System.out.println("I am: " + self.toString());
                Principal peer = sslSocket.getSession().getPeerPrincipal();
                System.out.println("Client is: " + peer.toString());

                sslSocket.close();
            }
            return null;
        }
    }
}
```

JsseClient.java

```
import java.io.*;
import java.net.*;
import javax.net.ssl.*;
import java.security.PrivilegedExceptionAction;
import java.security.Principal;

public class JsseClient {
```

```
    private static final String KRB5_CIPHER =
"TLS_KRB5_WITH_3DES_EDE_CBC_SHA";

    private static final int PORT = 4569;
    private static final boolean verbose = false;

    public static void main(String[] args) throws Exception {
        // Obtain the command-line arguments and parse the server's name

        if (args.length < 1) {
            System.err.println(
                "Usage: java <options> JsseClient <serverName>");
            System.exit(-1);
        }

        PrivilegedExceptionAction action = new JsseClientAction(args[0],
PORT);

        Jaas.loginAndAction("client", action);
    }

    static class JsseClientAction implements PrivilegedExceptionAction {
        private String server;
        private int port;

        JsseClientAction(String server, int port) {
            this.port = port;
            this.server = server;
        }

        public Object run() throws Exception {
            SSLSocketFactory sslsf =
                (SSLSocketFactory) SSLSocketFactory.getDefault();
            SSLSocket sslSocket = (SSLSocket) sslsf.createSocket(server,
port);

            // Enable only a KRB5 cipher suite.
            String enabledSuites[] = { KRB5_CIPHER };
            sslSocket.setEnabledCipherSuites(enabledSuites);
            // Should check for exception if enabledSuites is not supported

            BufferedReader in = new BufferedReader(new InputStreamReader(
                sslSocket.getInputStream()));
            BufferedWriter out = new BufferedWriter(new OutputStreamWriter(
                sslSocket.getOutputStream()));

            String outStr = "Hello There!\n";
            out.write(outStr);
            out.flush();
            System.out.print("Sending " + outStr);

            String inStr = in.readLine();
            System.out.println("Received " + inStr);
```

```
            String cipherSuiteChosen =
sslSocket.getSession().getCipherSuite();
            System.out.println("Cipher suite in use: " +
cipherSuiteChosen);
            Principal self = sslSocket.getSession().getLocalPrincipal();
            System.out.println("I am: " + self.toString());
            Principal peer = sslSocket.getSession().getPeerPrincipal();
            System.out.println("Server is: " + peer.toString());

            sslSocket.close();
            return null;
        }
    }
}
```

krb5.conf

```
# krb5.conf template
# In order to complete this configuration file
# you will need to replace the __<name>__ placeholders
# with appropriate values for your network.
#
[libdefaults]
        default_realm = J1LABS.EXAMPLE.COM
    forwardable = true

default_tkt_enctypes = aes128-cts rc4-hmac des3-cbc-sha1
default_tgs_enctypes = aes128-cts rc4-hmac des3-cbc-sha1
permitted_enctypes   = aes128-cts rc4-hmac des3-cbc-sha1

[realms]
        J1LABS.EXAMPLE.COM = {
        kdc = j1hol-1280
                kdc = j1hol-004
                admin_server = j1hol-1280
        }

[domain_realm]
    .example.com = J1LABS.EXAMPLE.COM

[logging]
        default = FILE:/var/krb5/kdc.log
        kdc = FILE:/var/krb5/kdc.log
    kdc_rotate = {

# How often to rotate kdc.log. Logs will get rotated no more
# often than the period, and less often if the KDC is not used
# frequently.

        period = 1d

# how many versions of kdc.log to keep around (kdc.log.0, kdc.log.1, ...)

        versions = 10
```

```
    }

[appdefaults]
    gkadmin = {
                help_url = http://localhost:8888/ab2/coll.384.1/SEAM
        }

    kinit = {
        renewable = true
        forwardable= true
    }

        rlogin = {
                forwardable= true
        }
        rsh = {
                forwardable= true
        }
        telnet = {
                autologin = true
                forwardable= true
        }
```

GssSpNegoClient.java

```
import org.ietf.jgss.*;
import java.net.Socket;
import java.io.IOException;
import java.io.DataInputStream;
import java.io.DataOutputStream;
import java.security.*;
import javax.security.auth.login.LoginException;

/**
 * A sample client application that uses JGSS to do mutual authentication
 * with a server using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the server.
 *
 * Every message sent to the server includes a 4-byte application-level
 * header that contains the big-endian integer value for the number
 * of bytes that will follow as part of the JGSS token.
 *
 * The protocol is:
 *    1.  Context establishment loop:
 *         a. client sends init sec context token to server
 *         b. server sends accept sec context token to client
 *         ....
 *    2. client sends a wrapped token to the server.
 *    3. server sends a wrapped token back to the client for the
application
 *
 * Start GssServer first before starting GssClient.
 *
 * Usage:  java <options> GssSpNegoClient <service> <serverName>
```

```
 *
 * Example: java -Djava.security.auth.login.config=jaas-krb5.conf \
 *               GssSpNegoClient host machine.imc.org
 *
 * Add -Djava.security.krb5.conf=krb5.conf to specify application-specific
 * Kerberos configuration (different from operating system's Kerberos
 * configuration).
 */

public class GssSpNegoClient {
    private static final int PORT = 4567;
    private static final boolean verbose = false;

    public static void main(String[] args) throws Exception {

        // Obtain the command-line arguments and parse the server's
principal

        if (args.length < 2) {
            System.err.println(
                "Usage: java <options> GssSpNegoClient <service>
<serverName>");
            System.exit(-1);
        }

        String serverPrinc = args[0] + "@" + args[1];

        PrivilegedExceptionAction action =
            new GssClientAction(serverPrinc, args[1], PORT);

        Jaas.loginAndAction("client", action);
    }

    static class GssClientAction implements PrivilegedExceptionAction {
        private String serverPrinc;
        private String hostName;
        private int port;

        GssClientAction(String serverPrinc, String hostName, int port) {
            this.serverPrinc = serverPrinc;
            this.hostName = hostName;
            this.port = port;
        }

        public Object run() throws Exception {
            Socket socket = new Socket(hostName, port);
            DataInputStream inStream =
                new DataInputStream(socket.getInputStream());
            DataOutputStream outStream =
                new DataOutputStream(socket.getOutputStream());

            System.out.println("Connected to address " +
                socket.getInetAddress());

            /*
```

```
                 * This Oid is used to represent the SPNEGO GSS-API
                 * mechanism. It is defined in RFC 2478. We will use this Oid
                 * whenever we need to indicate to the GSS-API that it must
                 * use SPNEGO for some purpose.
                 */
                Oid spnegoOid = new Oid("1.3.6.1.5.5.2");

                GSSManager manager = GSSManager.getInstance();

                /*
                 * Create a GSSName out of the server's name.
                 */
                GSSName serverName = manager.createName(serverPrinc,
                    GSSName.NT_HOSTBASED_SERVICE, spnegoOid);

                /*
                 * Create a GSSContext for mutual authentication with the
                 * server.
                 *    - serverName is the GSSName that represents the server.
                 *    - krb5Oid is the Oid that represents the mechanism to
                 *      use. The client chooses the mechanism to use.
                 *    - null is passed in for client credentials
                 *    - DEFAULT_LIFETIME lets the mechanism decide how long the
                 *      context can remain valid.
                 * Note: Passing in null for the credentials asks GSS-API to
                 * use the default credentials. This means that the mechanism
                 * will look among the credentials stored in the current
Subject
                 * to find the right kind of credentials that it needs.
                 */
                GSSContext context = manager.createContext(serverName,
                    spnegoOid,
                    null,
                    GSSContext.DEFAULT_LIFETIME);

                // Set the desired optional features on the context. The client
                // chooses these options.

                context.requestMutualAuth(true);  // Mutual authentication
                context.requestConf(true);  // Will use confidentiality later
                context.requestInteg(true); // Will use integrity later

                // Do the context eastablishment loop

                byte[] token = new byte[0];

                while (!context.isEstablished()) {

                    // token is ignored on the first call
                    token = context.initSecContext(token, 0, token.length);

                    // Send a token to the server if one was generated by
                    // initSecContext
                    if (token != null) {
                        if (verbose) {
```

```
                        System.out.println("Will send token of size " +
                            token.length + " from initSecContext.");
                        System.out.println("writing token = " +
                            getHexBytes(token));
                    }

                    outStream.writeInt(token.length);
                    outStream.write(token);
                    outStream.flush();
                }

                // If the client is done with context establishment
                // then there will be no more tokens to read in this loop
                if (!context.isEstablished()) {
                    token = new byte[inStream.readInt()];
                    if (verbose) {
                        System.out.println("reading token = " +
                            getHexBytes(token));
                        System.out.println("Will read input token of size
" +
                            token.length + " for processing by
initSecContext");
                    }
                    inStream.readFully(token);
                }
            }

            System.out.println("Context Established! ");
            System.out.println("Client principal is " +
context.getSrcName());
            System.out.println("Server principal is " +
context.getTargName());

            /*
             * If mutual authentication did not take place, then only the
             * client was authenticated to the server. Otherwise, both
             * client and server were authenticated to each other.
             */
            if (context.getMutualAuthState())
                System.out.println("Mutual authentication took place!");

            byte[] messageBytes = "Hello There!".getBytes("UTF-8");

            /*
             * The first MessageProp argument is 0 to request
             * the default Quality-of-Protection.
             * The second argument is true to request
             * privacy (encryption of the message).
             */
            MessageProp prop =  new MessageProp(0, true);

            /*
             * Encrypt the data and send it across. Integrity protection
             * is always applied, irrespective of confidentiality
             * (i.e., encryption).
```

```
             * You can use the same token (byte array) as that used when
             * establishing the context.
             */

            System.out.println("Sending message: " +
                new String(messageBytes, "UTF-8"));
            token = context.wrap(messageBytes, 0, messageBytes.length,
prop);
            outStream.writeInt(token.length);
            outStream.write(token);
            outStream.flush();

            /*
             * Now we will allow the server to decrypt the message,
             * append a time/date on it, and send then it back.
             */

            token = new byte[inStream.readInt()];
            System.out.println("Will read token of size " + token.length);
            inStream.readFully(token);
            byte[] replyBytes = context.unwrap(token, 0, token.length,
prop);

            System.out.println("Received message: " +
                new String(replyBytes, "UTF-8"));

            System.out.println("Done.");
            context.dispose();
            socket.close();

            return null;
        }
    }

    private static final String getHexBytes(byte[] bytes, int pos, int
len) {

        StringBuffer sb = new StringBuffer();
        for (int i = pos; i < (pos+len); i++) {

            int b1 = (bytes[i]>>4) & 0x0f;
            int b2 = bytes[i] & 0x0f;

            sb.append(Integer.toHexString(b1));
            sb.append(Integer.toHexString(b2));
            sb.append(' ');
        }
        return sb.toString();
    }

    private static final String getHexBytes(byte[] bytes) {
        return getHexBytes(bytes, 0, bytes.length);
    }
}
```

GssSpNegoServer.java

```
import org.ietf.jgss.*;
import java.io.*;
import java.net.Socket;
import java.net.ServerSocket;
import java.security.*;
import java.util.Date;

/**
 * A sample server application that uses JGSS to do mutual authentication
 * with a client using Kerberos as the underlying mechanism. It then
 * exchanges data securely with the client.
 *
 * Every message exchanged with the client includes a 4-byte application-
 * level header that contains the big-endian integer value for the number
 * of bytes that will follow as part of the JGSS token.
 *
 * The protocol is:
 *    1.  Context establishment loop:
 *         a. client sends init sec context token to server
 *         b. server sends accept sec context token to client
 *         ....
 *    2. client sends a wrap token to the server.
 *    3. server sends a wrap token back to the client.
 *
 * Start GssSpNegoServer first before starting GssClient.
 *
 * Usage:  java <options> GssSpNegoServer
 *
 * Example: java -Djava.security.auth.login.config=jaas-krb5.conf \
 *               GssSpNegoServer
 *
 * Add -Djava.security.krb5.conf=krb5.conf to specify application-specific
 * Kerberos configuration (different from operating system's Kerberos
 * configuration).
 */

public class GssSpNegoServer  {
    private static final int PORT = 4567;
    private static final boolean verbose = false;
    private static final int LOOP_LIMIT = 1;
    private static int loopCount = 0;

    public static void main(String[] args) throws Exception {

        PrivilegedExceptionAction action = new GssServerAction(PORT);

        Jaas.loginAndAction("server", action);
    }

    static class GssServerAction implements PrivilegedExceptionAction {
        private int localPort;
```

```
        GssServerAction(int port) {
            this.localPort = port;
        }

        public Object run() throws Exception {

            ServerSocket ss = new ServerSocket(localPort);

            // Get own Kerberos credentials for accepting connection
            GSSManager manager = GSSManager.getInstance();
            Oid spnegoOid = new Oid("1.3.6.1.5.5.2");
            GSSCredential serverCreds = manager.createCredential(null,

GSSCredential.DEFAULT_LIFETIME,
                                                spnegoOid,
                                                GSSCredential.ACCEPT_ONLY);
            while (loopCount++ < LOOP_LIMIT) {

                System.out.println("Waiting for incoming connection...");

                Socket socket = ss.accept();
                DataInputStream inStream =
                    new DataInputStream(socket.getInputStream());

                DataOutputStream outStream =
                    new DataOutputStream(socket.getOutputStream());

                System.out.println("Got connection from client " +
                    socket.getInetAddress());

                /*
                 * Create a GSSContext to receive the incoming request
                 * from the client. Use null for the server credentials
                 * passed in. This tells the underlying mechanism
                 * to use whatever credentials it has available that
                 * can be used to accept this connection.
                 */

                GSSContext context = manager.createContext(
                    (GSSCredential)serverCreds);

                // Do the context establishment loop

                byte[] token = null;

                while (!context.isEstablished()) {

                    if (verbose) {
                        System.out.println("Reading ...");
                    }
                    token = new byte[inStream.readInt()];

                    if (verbose) {
                        System.out.println("Will read input token of size
" +
```

```
                            token.length + " for processing by
acceptSecContext");
                    }
                    inStream.readFully(token);

                    if (token.length == 0) {
                        if (verbose) {
                            System.out.println("skipping zero length
token");
                        }
                        continue;
                    }
                    if (verbose) {
                        System.out.println("Token = " +
getHexBytes(token));
                        System.out.println("acceptSecContext..");
                    }
                    token = context.acceptSecContext(token, 0,
token.length);

                    // Send a token to the peer if one was generated by
                    // acceptSecContext
                    if (token != null) {
                        if (verbose) {
                            System.out.println("Will send token of size " +
                                token.length + " from acceptSecContext.");
                        }

                        outStream.writeInt(token.length);
                        outStream.write(token);
                        outStream.flush();
                    }
                }

                System.out.println("Context Established! ");
                System.out.println("Client principal is " +
context.getSrcName());
                System.out.println("Server principal is " +
context.getTargName());

                /*
                 * If mutual authentication did not take place, then
                 * only the client was authenticated to the
                 * server. Otherwise, both client and server were
                 * authenticated to each other.
                 */
                if (context.getMutualAuthState())
                    System.out.println("Mutual authentication took
place!");

                /*
                 * Create a MessageProp which unwrap will use to return
                 * information such as the Quality-of-Protection that was
                 * applied to the wrapped token, whether or not it was
                 * encrypted, etc. Since the initial MessageProp values
```

```
                 * are ignored, just set them to the defaults of 0 and
false.
                 */
                MessageProp prop = new MessageProp(0, false);

                /*
                 * Read the token. This uses the same token byte array
                 * as that used during context establishment.
                 */
                token = new byte[inStream.readInt()];
                if (verbose) {
                    System.out.println("Will read token of size " +
token.length);
                }
                inStream.readFully(token);

                byte[] input = context.unwrap(token, 0, token.length,
prop);
                String str = new String(input, "UTF-8");

                System.out.println("Received data \"" +
                    str + "\" of length " + str.length());

                System.out.println("Confidentiality applied: " +
                    prop.getPrivacy());

                /*
                 * Now generate reply that is the concatenation of the
                 * incoming string with the current time.
                 */

                /*
                 * First reset the QOP of the MessageProp to 0
                 * to ensure the default Quality-of-Protection
                 * is applied.
                 */
                prop.setQOP(0);

                String now = new Date().toString();
                byte[] nowBytes = now.getBytes("UTF-8");
                int len = input.length + 1 + nowBytes.length;
                byte[] reply = new byte[len];
                System.arraycopy(input, 0, reply, 0, input.length);
                reply[input.length] = ' ';
                System.arraycopy(nowBytes, 0, reply, input.length+1,
                    nowBytes.length);

                System.out.println("Sending: " + new String(reply,
"UTF-8"));
                token = context.wrap(reply, 0, reply.length, prop);

                outStream.writeInt(token.length);
                outStream.write(token);
                outStream.flush();
```

```
                System.out.println("Closing connection with client " +
                    socket.getInetAddress());
                context.dispose();
                socket.close();
            }
            return null;
        }
    }

    private static final String getHexBytes(byte[] bytes, int pos, int
len) {

        StringBuffer sb = new StringBuffer();
        for (int i = pos; i < (pos+len); i++) {

            int b1 = (bytes[i]>>4) & 0x0f;
            int b2 = bytes[i] & 0x0f;

            sb.append(Integer.toHexString(b1));
            sb.append(Integer.toHexString(b2));
            sb.append(' ');
        }
        return sb.toString();
    }

    private static final String getHexBytes(byte[] bytes) {
        return getHexBytes(bytes, 0, bytes.length);
    }
}
```

Appendix A: Setting up Kerberos Accounts

Kerberos accounts are set up on the Key Distribution Center (KDC). Each entry in the Kerberos database contains a Kerberos principal. You should create a host-based principal for the machine that you will be running the servers (for example, "host/j1hol-001") and a client principal (for example, "test") for accessing the servers.

For Solaris, please refer to following documentation on how to setup Kerberos principals.

- Administering Kerberos Principals
- System Administration Guide: Security Services

For Windows, please refer to Microsoft documentation. Here are some pointers.

- How To Create an Active Directory Server in Windows Server 2003
- Microsoft Kerberos
- Interoperability with Microsoft Windows 2000 Active Directory and Kerberos Services

The exercises assume that the operating system has been configured to use the correct Kerberos server. This configuration typically requires administration privileges. If you cannot configure the operating system, then you can use a Kerberos configuration file with your `java` command by using the `-Djava.security.krb5.conf`

option. Here is an example of how to invoke one of the commands from the exercises to use the `krb5.conf` configuration file.

```
% java -Djava.security.auth.login.config=jaas-krb5.conf\
  -Djava.security.krb5.conf=krb5.conf Jaas client
```

The Kerberos 5 GSS-API Mechanism

This section describes and lists security features regarding Java Generic Security Services (Java GSS) for Kerberos 5. It also describes the Object Identifier (OID) for the Kerberos V5 mechanism, the encryption types, and the `krb5.conf` settings supported by Java GSS.

The Generic Security Services Application Program Interface (GSS-API) mechanism is defined by RFC 1964 and supplemented with RFC 4121 under the Internet Standards process.

The OID for the Kerberos V5 Mechanism

According to RFC 1964 section 1, the OID for Java Generic Security Services (Java GSS) for Kerberos 5 is defined as 1.2.840.113554.1.2.2; see also GSSAPI Mechanisms in *Java Security Standard Algorithm Names*.

Java GSS/Kerberos Supported Encryption Types

The following table lists the preferred order of Java GSS/Kerberos supported encryption types.

Table 7-1 Java GSS/Kerberos Supported Encryption Types

Name	Aliases	etype Number
aes256-cts-hmac-sha1-96	aes256-sha1, aes256-cts	18
aes128-cts-hmac-sha1-96	aes128-sha1, aes128-cts	17
aes256-cts-hmac-sha384-192	aes256-sha2	20
aes128-cts-hmac-sha256-128	aes128-sha2	19
des3-cbc-sha1	des3-hmac-sha1	16
arcfour-hmac-md5	arcfour-hmac, rc4-hmac	23
des-cbc-crc	*None*	1
des-cbc-md5	*None*	3

> **Note:**
>
> The AES-256 encryption type is enabled by default. The DES-based encryption types, including `des-cbc-crc` and `dec-cbc-md5`, are disabled by default.

A user can restrict the usage of encryption for various purposes in `krb5.conf`, in the `[libdefaults]` section.

Supported krb5.conf Settings

The following parameters are supported:

```
include FILENAME
includedir DIRNAME

[libdefaults]
default_realm
allow_weak_crypto

dns_lookup_kdc
dns_lookup_realm
dns_fallback

default_checksum
safe_checksum_type
ap_req_checksum_type
default_keytab_name

default_tkt_enctypes
permitted_enctypes
default_tgs_enctypes

no_addresses
noaddresses

renewable
proxiable
forwardable

kdc_default_options
clockskew

kdc_timeout
udp_preference_limit

max_retries
renew_lifetime
ticket_lifetime

[realms]
  REALM.NAME = {
      kdc
      kdc_timeout
      udp_preference_limit
      max_retries
  }

[capaths]
  A = {
    I = .
    B = I
  }
```

```
[domain_realm]
  domain=REALM
```

The following are the defaults for the `krb5.conf` file parameters:

```
no_addresses = true
noaddresses = true
dns_lookup_kdc = true
dns_lookup_realm = false
allow_weak_crypto = false
kdc_timeout = 30s
max_retries = 3
udp_preference_limit = 1465
clockskew = 300
renewable = false
proxiable = false
forwardable = false
```

8

Java Secure Socket Extension (JSSE) Reference Guide

The Java Secure Socket Extension (JSSE) enables secure Internet communications. It provides a framework and an implementation for a Java version of the TLS and DTLS protocols and includes functionality for data encryption, server authentication, message integrity, and optional client authentication.

Introduction to JSSE

Data that travels across a network can easily be accessed by someone who is not the intended recipient. When the data includes private information, such as passwords and credit card numbers, steps must be taken to make the data unintelligible to unauthorized parties. It is also important to ensure that the data has not been modified, either intentionally or unintentionally, during transport. The Transport Layer Security (TLS) protocol was designed to help protect the privacy and integrity of data while it is being transferred across a network.

The Java Secure Socket Extension (JSSE) enables secure Internet communications. It provides a framework and an implementation for a Java version of the TLS protocol and includes functionality for data encryption, server authentication, message integrity, and optional client authentication. Using JSSE, developers can provide for the secure passage of data between a client and a server running any application protocol (such as HTTP, Telnet, or FTP) over TCP/IP.

By abstracting the complex underlying security algorithms and handshaking mechanisms, JSSE minimizes the risk of creating subtle but dangerous security vulnerabilities. Furthermore, it simplifies application development by serving as a building block that developers can integrate directly into their applications.

JSSE provides both an application programming interface (API) framework and an implementation of that API. The JSSE API supplements the core network and cryptographic services defined by the `java.security` and `java.net` packages by providing extended networking socket classes, trust managers, key managers, SSL contexts, and a socket factory framework for encapsulating socket creation behavior. Because the `SSLSocket` class is based on a blocking I/O model, the Java Development Kit (JDK) includes a nonblocking `SSLEngine` class to enable implementations to choose their own I/O methods.

The JSSE API supports the following security protocols:

- DTLS: versions 1.0 and 1.2
- TLS: version 1.0, 1.1, 1.2, and 1.3
- SSL (Secure Socket Layer): version 3.0

These security protocols encapsulate a normal bidirectional stream socket, and the JSSE API adds transparent support for authentication, encryption, and integrity protection.

JSSE is a security component of the Java SE platform, and is based on the same design principles found elsewhere in the Java Cryptography Architecture (JCA) Reference Guide framework. This framework for cryptography-related security components allows them to have implementation independence and, whenever possible, algorithm independence. JSSE uses the Cryptographic Service Providers defined by the JCA framework.

Other security components in the Java SE platform include the Java Authentication and Authorization Service (JAAS) Reference Guide and the Java Security Tools. JSSE encompasses many of the same concepts and algorithms as those in JCA but automatically applies them underneath a simple stream socket API.

The JSSE API was designed to allow other SSL/TLS/DTLS protocols and Public Key Infrastructure (PKI) implementations to be plugged in seamlessly. Developers can also provide alternative logic to determine if remote hosts should be trusted or what authentication key material should be sent to a remote host.

JSSE Features and Benefits

JSSE includes the following important benefits and features:

- Included as a standard component of the JDK
- Extensible, provider-based architecture
- Implemented in 100% pure Java
- Provides API support for TLS/DTLS
- Provides implementations of SSL 3.0, TLS (versions 1.0, 1.1, 1.2, and 1.3), and DTLS (versions 1.0 and 1.2)
- Includes classes that can be instantiated to create secure channels (`SSLSocket`, `SSLServerSocket`, and `SSLEngine`)
- Provides support for cipher suite negotiation, which is part of the TLS/DTLS handshaking used to initiate or verify secure communications
- Provides support for client and server authentication, which is part of the normal TLS/DTLS handshaking
- Provides support for HTTP encapsulated in the TLS protocol, which allows access to data such as web pages using HTTPS
- Provides server session management APIs to manage memory-resident SSL sessions
- Provides support for the certificate status request extension (OCSP stapling), which saves client certificate validation round-trips and resources
- Provides support for the Server Name Indication (SNI) rxtension, which extends the TLS/DTLS protocols to indicate what server name the client is attempting to connect to during handshaking
- Provides support for endpoint identification during handshaking, which prevents man-in-the-middle attacks
- Provides support for cryptographic algorithm constraints, which provides fine-grained control over algorithms negotiated by JSSE

JSSE Standard API

The JSSE standard API, available in the `javax.net` and `javax.net.ssl` packages, provides:

- Secure sockets tailored to client and server-side applications.
- A non-blocking engine for producing and consuming streams of TLS/DTLS data (`SSLEngine`).
- Factories for creating sockets, server sockets, SSL sockets, and SSL server sockets. By using socket factories, you can encapsulate socket creation and configuration behavior.
- A class representing a secure socket context that acts as a factory for secure socket factories and engines.
- Key and trust manager interfaces (including X.509-specific key and trust managers), and factories that can be used for creating them.
- A class for secure HTTP URL connections (HTTPS).

SunJSSE Provider

Oracle's implementation of Java SE includes a JSSE provider named *SunJSSE*, which comes preinstalled and preregistered with the JCA. This provider supplies the following cryptographic services:

- An implementation of the SSL 3.0, TLS (versions 1.0, 1.1, 1.2, and 1.3), and DTLS (versions 1.0 and 1.2) security protocols.
- An implementation of the most common TLS and DTLS cipher suites. This implementation encompasses a combination of authentication, key agreement, encryption, and integrity protection.
- An implementation of an X.509-based key manager that chooses appropriate authentication keys from a standard JCA keystore.
- An implementation of an X.509-based trust manager that implements rules for certificate chain validation.

See The SunJSSE Provider.

JSSE Related Documentation

The following list contains links to online documentation and names of books about related subjects:

JSSE API Documentation

- `javax.net` package
- `javax.net.ssl` package

Java SE Security

- The Java SE Security home page
- The Security Features in Java SE trail of the Java Tutorial

- Java PKI Programmer's Guide
- Inside Java 2 Platform Security, Second Edition: Architecture, API Design and Implementation

Transport Layer Security (TLS)

- The TLS Protocol Version 1.0
- The TLS Protocol Version 1.1
- The TLS Protocol Version 1.2
- The (TLS) Protocol Version 1.3
- Transport Layer Security (TLS) Extensions
- HTTP Over TLS

Datagram Transport Layer Security (DTLS)

- The DTLS Protocol Version 1.0
- The DTLS Protocol Version 1.2

U.S. Encryption Policies

- U.S. Department of Commerce
- Technology CEO Council
- Current export policies: Encryption and Export Administration Regulations (EAR)
- NIST Computer Security Publications

JSSE Classes and Interfaces

To communicate securely, both sides of the connection must be SSL-enabled. In the JSSE API, the endpoint classes of the connection are `SSLSocket` and `SSLEngine`. In Figure 8-1, the major classes used to create `SSLSocket` and `SSLEngine` are laid out in a logical ordering.

Figure 8-1 JSSE Classes Used to Create SSLSocket and SSLEngine

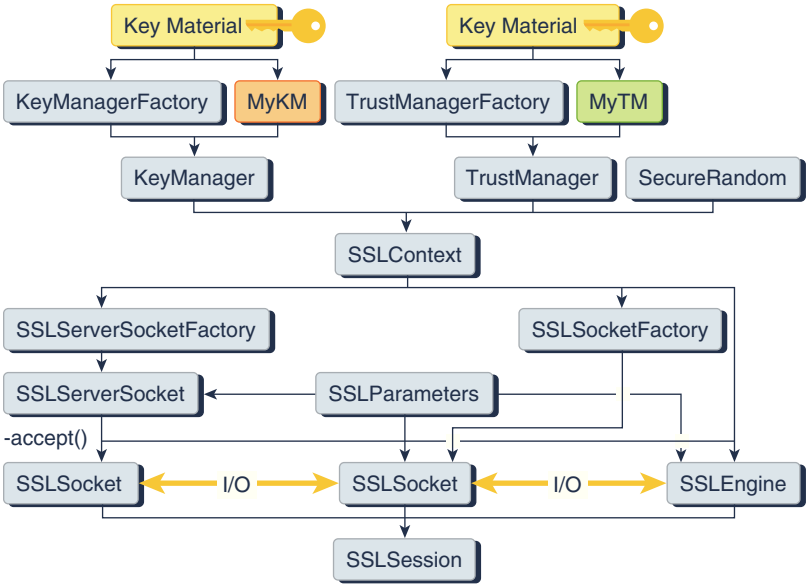

An `SSLSocket` is created either by an `SSLSocketFactory` or by an `SSLServerSocket` accepting an inbound connection. In turn, an `SSLServerSocket` is created by an `SSLServerSocketFactory`. Both `SSLSocketFactory` and `SSLServerSocketFactory` objects are created by an `SSLContext`. An `SSLEngine` is created directly by an `SSLContext`, and relies on the application to handle all I/O.

> **Note:**
>
> When using raw `SSLSocket` or `SSLEngine` classes, you should always check the peer's credentials before sending any data. Since JDK 7, endpoint identification/verification procedures can be handled during SSL/TLS handshaking. See the method `SSLParameters.setEndpointIdentificationAlgorithm`.
> For example, the host name in a URL should match the host name in the peer's credentials. An application could be exploited with URL spoofing if the host name is not verified.

JSSE Core Classes and Interfaces

The core JSSE classes are part of the `javax.net` and `javax.net.ssl` packages.

SocketFactory and ServerSocketFactory Classes

The abstract `javax.net.SocketFactory` class is used to create sockets. Subclasses of this class are factories that create particular subclasses of sockets and thus provide a general framework for the addition of public socket-level functionality. For example, see SSLSocketFactory and SSLServerSocketFactory Classes.

The abstract `javax.net.ServerSocketFactory` class is analogous to the `SocketFactory` class, but is used specifically for creating server sockets.

Socket factories are a simple way to capture a variety of policies related to the sockets being constructed, producing such sockets in a way that does not require special configuration of the code that asks for the sockets:

- Due to polymorphism of both factories and sockets, different kinds of sockets can be used by the same application code just by passing different kinds of factories.
- Factories can themselves be customized with parameters used in socket construction. For example, factories could be customized to return sockets with different networking timeouts or security parameters already configured.
- The sockets returned to the application can be subclasses of `java.net.Socket` (or `javax.net.ssl.SSLSocket`), so that they can directly expose new APIs for features such as compression, security, record marking, statistics collection, or firewall tunneling.

SSLSocketFactory and SSLServerSocketFactory Classes

The `javax.net.ssl.SSLSocketFactory` class acts as a factory for creating secure sockets. This class is an abstract subclass of `javax.net.SocketFactory`.

Secure socket factories encapsulate the details of creating and initially configuring secure sockets. This includes authentication keys, peer certificate validation, enabled cipher suites, and the like.

The `javax.net.ssl.SSLServerSocketFactory` class is analogous to the `SSLSocketFactory` class, but is used specifically for creating server sockets.

Obtaining an SSLSocketFactory

The following ways can be used to obtain an `SSLSocketFactory`:

- Get the default factory by calling the `SSLSocketFactory.getDefault()` static method.
- Receive a factory as an API parameter. That is, code that must create sockets but does not care about the details of how the sockets are configured can include a method with an `SSLSocketFactory` parameter that can be called by clients to specify which `SSLSocketFactory` to use when creating sockets (for example, `javax.net.ssl.HttpsURLConnection`).
- Construct a new factory with specifically configured behavior.

The default factory is typically configured to support server authentication only so that sockets created by the default factory do not leak any more information about the client than a normal TCP socket would.

Many classes that create and use sockets do not need to know the details of socket creation behavior. Creating sockets through a socket factory passed in as a parameter is a good way of isolating the details of socket configuration, and increases the reusability of classes that create and use sockets.

You can create new socket factory instances either by implementing your own socket factory subclass or by using another class which acts as a factory for socket factories.

One example of such a class is `SSLContext`, which is provided with the JSSE implementation as a provider-based configuration class.

SSLSocket and SSLServerSocket Classes

The `javax.net.ssl.SSLSocket` class is a subclass of the standard Java `java.net.Socket` class. It supports all of the standard socket methods and adds methods specific to secure sockets. Instances of this class encapsulate the `SSLContext` under which they were created. See SSLContext Class. There are APIs to control the creation of secure socket sessions for a socket instance, but trust and key management are not directly exposed.

The `javax.net.ssl.SSLServerSocket` class is analogous to the `SSLSocket` class, but is used specifically for creating server sockets.

To prevent peer spoofing, you should always verify the credentials presented to an `SSLSocket`. See Cipher Suite Choice and Remote Entity Verification.

> **Note:**
>
> Due to the complexity of the SSL and TLS protocols, it is difficult to predict whether incoming bytes on a connection are handshake or application data, and how that data might affect the current connection state (even causing the process to block). In the Oracle JSSE implementation, the `available()` method on the object obtained by `SSLSocket.getInputStream()` returns a count of the number of application data bytes successfully decrypted from the SSL connection but not yet read by the application.

Obtaining an SSLSocket

Instances of `SSLSocket` can be obtained in one of the following ways:

- An `SSLSocket` can be created by an instance of `SSLSocketFactory` via one of the several `createSocket` methods of that class.
- An `SSLSocket` can be created through the `accept` method of the `SSLServerSocket` class.

Cipher Suite Choice and Remote Entity Verification

The SSL/TLS protocols define a specific series of steps to ensure a *protected* connection. However, the choice of cipher suite directly affects the type of security that the connection enjoys. For example, if an anonymous cipher suite is selected, then the application has no way to verify the remote peer's identity. If a suite with no encryption is selected, then the privacy of the data cannot be protected. Additionally, the SSL/TLS protocols do not specify that the credentials received must match those that peer might be expected to send. If the connection were somehow redirected to a rogue peer, but the rogue's credentials were acceptable based on the current trust material, then the connection would be considered valid.

When using raw `SSLSocket` and `SSLEngine` classes, you should always check the peer's credentials before sending any data. The `SSLSocket` and `SSLEngine` classes do

not automatically verify that the host name in a URL matches the host name in the peer's credentials. An application could be exploited with URL spoofing if the host name is not verified. Since JDK 7, endpoint identification/verification procedures can be handled during SSL/TLS handshaking. See the `SSLParameters.getEndpointIdentificationAlgorithm` method.

Protocols such as HTTPS (HTTP Over TLS) do require host name verification. Since JDK 7, the HTTPS endpoint identification is enforced during handshaking for `HttpsURLConnection` by default. See the `SSLParameters.getEndpointIdentificationAlgorithm` method. Alternatively, applications can use the `HostnameVerifier` interface to override the default HTTPS host name rules. See HostnameVerifier Interface and HttpsURLConnection Class.

SSLEngine Class

TLS/DTLS is becoming increasingly popular. It is being used in a wide variety of applications across a wide range of computing platforms and devices. Along with this popularity come demands to use TLS/DTLS with different I/O and threading models to satisfy the applications' performance, scalability, footprint, and other requirements. There are demands to use TLS/DTLS with blocking and nonblocking I/O channels, asynchronous I/O, arbitrary input and output streams, and byte buffers. There are demands to use it in highly scalable, performance-critical environments, requiring management of thousands of network connections.

Abstraction of the I/O transport mechanism using the `SSLEngine` class in Java SE allows applications to use the TLS/DTLS protocols in a transport-independent way, and thus frees application developers to choose transport and computing models that best meet their needs. Not only does this abstraction allow applications to use nonblocking I/O channels and other I/O models, it also accommodates different threading models. This effectively leaves the I/O and threading decisions up to the application developer. Because of this flexibility, the application developer must manage I/O and threading (complex topics in and of themselves), as well as have some understanding of the TLS/DTLS protocols. The abstraction is therefore an advanced API: beginners should use `SSLSocket`.

Users of other Java programming language APIs such as the Java Generic Security Services (Java GSS-API) and the Java Simple Authentication Security Layer (Java SASL) will notice similarities in that the application is also responsible for transporting data.

The core class is `javax.net.ssl.SSLEngine`. It encapsulates a TLS/DTLS state machine and operates on inbound and outbound byte buffers supplied by the user of the `SSLEngine` class. Figure 8-2 illustrates the flow of data from the application, through `SSLEngine`, to the transport mechanism, and back.

Figure 8-2 Flow of Data Through SSLEngine

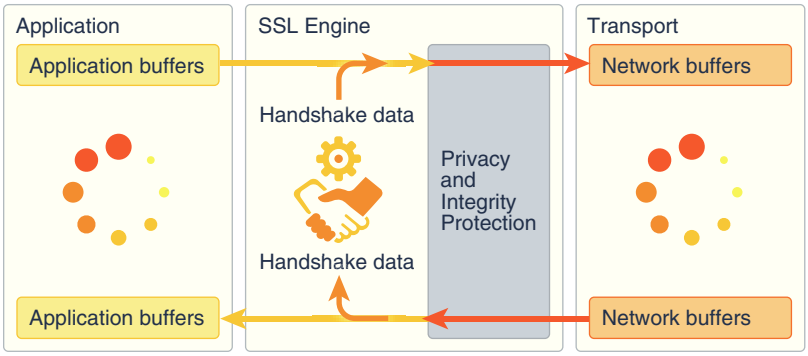

The application, shown on the left, supplies application (plaintext) data in an application buffer and passes it to `SSLEngine`. The `SSLEngine` object processes the data contained in the buffer, or any handshaking data, to produce TLS/DTLS encoded data and places it to the network buffer supplied by the application. The application is then responsible for using an appropriate transport (shown on the right) to send the contents of the network buffer to its peer. Upon receiving TLS/DTLS encoded data from its peer (via the transport), the application places the data into a network buffer and passes it to `SSLEngine`. The `SSLEngine` object processes the network buffer's contents to produce handshaking data or application data.

An instance of the `SSLEngine` class can be in one of the following states:

Creation
The `SSLEngine` has been created and initialized, but has not yet been used. During this phase, an application may set any `SSLEngine`-specific settings (enabled cipher suites, whether the `SSLEngine` should handshake in client or server mode, and so on). Once handshaking has begun, though, any new settings (except client/server mode) will be used for the next handshake.

Initial handshaking
The initial handshake is a procedure by which the two peers exchange communication parameters until an `SSLSession` is established. Application data can't be sent during this phase.

Application data
After the communication parameters have been established and the handshake is complete, application data can flow through the `SSLEngine`. Outbound application messages are encrypted and integrity protected, and inbound messages reverse the process.

Rehandshaking
Either side can request a renegotiation of the session at any time during the Application Data phase. New handshaking data can be intermixed among the application data. Before starting the rehandshake phase, the application may reset the TLS/DTLS communication parameters such as the list of enabled cipher suites and whether to use client authentication, but can not change between client/server modes. As before, after handshaking has begun, any new `SSLEngine` configuration settings won't be used until the next handshake.

Closure
When the connection is no longer needed, the application should close the `SSLEngine` and should send/receive any remaining messages to the peer before closing the underlying transport mechanism. Once an engine is closed, it is not reusable: a new `SSLEngine` must be created.

Understanding SSLEngine Operation Statuses

The status of the SSLEngine is represented by `SSLEngineResult.Status`.

To indicate the status of the engine and what actions the application should take, the `SSLEngine.wrap()` and `SSLEngine.unwrap()` methods return an `SSLEngineResult` instance, as shown in Example 8-5. This `SSLEngineResult` object contains two pieces of status information: the overall status of the engine and the handshaking status.

The possible overall statuses are represented by the `SSLEngineResult.Status` enum. The following statuses are available:

`OK`
There was no error.

`CLOSED`
The operation closed the `SSLEngine` or the operation could not be completed because it was already closed.

`BUFFER_UNDERFLOW`
The input buffer had insufficient data, indicating that the application must obtain more data from the peer (for example, by reading more data from the network).

`BUFFER_OVERFLOW`
The output buffer had insufficient space to hold the result, indicating that the application must clear or enlarge the destination buffer.

Example 8-1 illustrates how to handle the `BUFFER_UNDERFLOW` and `BUFFER_OVERFLOW` statuses of the `SSLEngine.unwrap()` method. It uses `SSLSession.getApplicationBufferSize()` and `SSLSession.getPacketBufferSize()` to determine how large to make the byte buffers.

The possible handshaking statuses are represented by the `SSLEngineResult.HandshakeStatus` enum. They represent whether handshaking has completed, whether the caller must obtain more handshaking data from the peer or send more handshaking data to the peer, and so on. The following handshake statuses are available:

`FINISHED`
The `SSLEngine` has just finished handshaking.

`NEED_TASK`
The `SSLEngine` needs the results of one (or more) delegated tasks before handshaking can continue.

`NEED_UNWRAP`
The `SSLEngine` needs to receive data from the remote side before handshaking can continue.

NEED_UNWRAP_AGAIN
The `SSLEngine` needs to unwrap before handshaking can continue. This value indicates that not-yet-interpreted data has been previously received from the remote side and does not need to be received again; the data has been brought into the JSSE framework but has not been processed yet.

NEED_WRAP
The `SSLEngine` must send data to the remote side before handshaking can continue, so `SSLEngine.wrap()` should be called.

NOT_HANDSHAKING
The `SSLEngine` is not currently handshaking.

Having two statuses per result allows the `SSLEngine` to indicate that the application must take two actions: one in response to the handshaking and one representing the overall status of the `wrap()` and `unwrap()` methods. For example, the engine might, as the result of a single `SSLEngine.unwrap()` call, return `SSLEngineResult.Status.OK` to indicate that the input data was processed successfully and `SSLEngineResult.HandshakeStatus.NEED_UNWRAP` to indicate that the application should obtain more TLS/DTLS encoded data from the peer and supply it to `SSLEngine.unwrap()` again so that handshaking can continue. As you can see, the previous examples were greatly simplified; they would need to be expanded significantly to properly handle all of these statuses.

Example 8-3 and Example 8-2 illustrate how to process handshaking data by checking handshaking status and the overall status of the `wrap()` and `unwrap()` methods.

Example 8-1 Sample Code for Handling BUFFER_UNDERFLOW and BUFFER_OVERFLOW

The following code sample illustrates how to handle BUFFER_UNDERFLOW and BUFFER_OVERFLOW status:

```
    SSLEngineResult res = engine.unwrap(peerNetData, peerAppData);
    switch (res.getStatus()) {

    case BUFFER_OVERFLOW:
            // Maybe need to enlarge the peer application data buffer.
        if (engine.getSession().getApplicationBufferSize() >
peerAppData.capacity()) {
            // enlarge the peer application data buffer
        } else {
            // compact or clear the buffer
        }
        // retry the operation
    break;

    case BUFFER_UNDERFLOW:
        // Maybe need to enlarge the peer network packet buffer
        if (engine.getSession().getPacketBufferSize() >
peerNetData.capacity()) {
        // enlarge the peer network packet buffer
        } else {
        // compact or clear the buffer
        }
        // obtain more inbound network data and then retry the operation
```

```
        break;

        // Handle other status: CLOSED, OK
        // ...
    }
```

Example 8-2 Sample Code for Checking and Processing Handshaking Statuses and Overall Statuses

The following code sample illustrates how to process handshaking data by checking handshaking status and the overall status of the `wrap()` and `unwrap()` methods:

```
void doHandshake(SocketChannel socketChannel, SSLEngine engine,
    ByteBuffer myNetData, ByteBuffer peerNetData) throws Exception {

    // Create byte buffers to use for holding application data
    int appBufferSize = engine.getSession().getApplicationBufferSize();
    ByteBuffer myAppData = ByteBuffer.allocate(appBufferSize);
    ByteBuffer peerAppData = ByteBuffer.allocate(appBufferSize);

    // Begin handshake
    engine.beginHandshake();
    SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();

    // Process handshaking message
    while (hs != SSLEngineResult.HandshakeStatus.FINISHED &&
        hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {

        switch (hs) {

        case NEED_UNWRAP:
            // Receive handshaking data from peer
            if (socketChannel.read(peerNetData) < 0) {
                // The channel has reached end-of-stream
            }

            // Process incoming handshaking data
            peerNetData.flip();
            SSLEngineResult res = engine.unwrap(peerNetData, peerAppData);
            peerNetData.compact();
            hs = res.getHandshakeStatus();

            // Check status
            switch (res.getStatus()) {
            case OK :
                // Handle OK status
                break;

            // Handle other status: BUFFER_UNDERFLOW, BUFFER_OVERFLOW,
CLOSED
            // ...
            }
            break;

        case NEED_WRAP :
```

```
            // Empty the local network packet buffer.
            myNetData.clear();

            // Generate handshaking data
            res = engine.wrap(myAppData, myNetData);
            hs = res.getHandshakeStatus();

            // Check status
            switch (res.getStatus()) {
            case OK :
                myNetData.flip();

                // Send the handshaking data to peer
                while (myNetData.hasRemaining()) {
                    socketChannel.write(myNetData);
                }
                break;

            // Handle other status:  BUFFER_OVERFLOW, BUFFER_UNDERFLOW,
CLOSED
            // ...
            }
            break;

        case NEED_TASK :
            // Handle blocking tasks
            break;

            // Handle other status:  // FINISHED or NOT_HANDSHAKING
            // ...
        }
    }

    // Processes after handshaking
    // ...
}
```

Example 8-3 Sample Code for Handling DTLS handshake Status and Overall Status

The following code sample illustrates how to handle DTLS handshake status:

```
void handshake(SSLEngine engine, DatagramSocket socket,
               SocketAddress peerAddr) throws Exception {
    boolean endLoops = false;
    // private static int MAX_HANDSHAKE_LOOPS = 60;
    int loops = MAX_HANDSHAKE_LOOPS;
    engine.beginHandshake();
    while (!endLoops && (serverException == null) && (clientException ==
null)) {
        if (--loops < 0) {
            throw new RuntimeException("Too many loops to produce
handshake packets");
        }
        SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
```

```
        if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP ||
                hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP_AGAIN) {
            ByteBuffer iNet;
            ByteBuffer iApp;
            if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP) {
                // receive ClientHello request and other SSL/TLS/DTLS
records
                byte[] buf = new byte[1024];
                DatagramPacket packet = new DatagramPacket(buf,
buf.length);
                try {
                    socket.receive(packet);
                } catch (SocketTimeoutException ste) {
                    // retransmit the packet if timeout
                    List <Datagrampacket> packets =
                        onReceiveTimeout(engine, peerAddr);
                    for (DatagramPacket p : packets) {
                        socket.send(p);
                    }
                    continue;
                }
                iNet = ByteBuffer.wrap(buf, 0, packet.getLength());
                iApp = ByteBuffer.allocate(1024);
            } else {
                iNet = ByteBuffer.allocate(0);
                iApp = ByteBuffer.allocate(1024);
            }
            SSLEngineResult r = engine.unwrap(iNet, iApp);
            SSLEngineResult.Status rs = r.getStatus();
            hs = r.getHandshakeStatus();
            if (rs == SSLEngineResult.Status.BUFFER_OVERFLOW) {
                // the client maximum fragment size config does not work?
                throw new Exception("Buffer overflow: " +
                                    "incorrect client maximum fragment
size");
            } else if (rs == SSLEngineResult.Status.BUFFER_UNDERFLOW) {
                // bad packet, or the client maximum fragment size
                // config does not work?
                if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING)
{
                    throw new Exception("Buffer underflow: " +
                                        "incorrect client maximum fragment
size");
                } // otherwise, ignore this packet
            } else if (rs == SSLEngineResult.Status.CLOSED) {
                endLoops = true;
            }   // otherwise, SSLEngineResult.Status.OK:
            if (rs != SSLEngineResult.Status.OK) {
                continue;
            }
        } else if (hs == SSLEngineResult.HandshakeStatus.NEED_WRAP) {
            List <DatagramPacket> packets =
                // Call a function to produce handshake packets
                produceHandshakePackets(engine, peerAddr);
            for (DatagramPacket p : packets) {
```

```
                socket.send(p);
            }
        } else if (hs == SSLEngineResult.HandshakeStatus.NEED_TASK) {
            runDelegatedTasks(engine);
        } else if (hs == SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
            // OK, time to do application data exchange.
            endLoops = true;
        } else if (hs == SSLEngineResult.HandshakeStatus.FINISHED) {
            endLoops = true;
        }
    }
    SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
    if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
        throw new Exception("Not ready for application data yet");
    }
}
```

SSLEngine for TLS Protocols

This section shows you how to create an `SSLEngine` object and use it to generate and process TLS data.

Creating an SSLEngine Object

Use the `SSLContext.createSSLEngine()` method to create an `SSLEngine` object.

Before you create an `SSLEngine` object, you must configure the engine to act as a client or a server, and set other configuration parameters, such as which cipher suites to use and whether client authentication is required. The `SSLContext.createSSLEngine` method creates an `javax.net.ssl.SSLEngine` object.

Example 8-4 Sample Code for Creating an SSLEngine Client for TLS with JKS as Keystore

The following sample code creates an `SSLEngine` client for TLS that uses JKS as keystore.

> **Note:**
>
> In this sample, the server name and port number are not used for communicating with the server (all transport is the responsibility of the application). They are hints to the JSSE provider to use for TLS session caching.

```
import javax.net.ssl.*;
import java.security.*;

// Create and initialize the SSLContext with key material
char[] passphrase = "passphrase".toCharArray();

// First initialize the key and trust material
KeyStore ksKeys = KeyStore.getInstance("JKS");
ksKeys.load(new FileInputStream("testKeys"), passphrase);
```

```
KeyStore ksTrust = KeyStore.getInstance("JKS");
ksTrust.load(new FileInputStream("testTrust"), passphrase);

// KeyManagers decide which key material to use
KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
kmf.init(ksKeys, passphrase);

// TrustManagers decide whether to allow connections
TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
tmf.init(ksTrust);

// Get an instance of SSLContext for TLS protocols
sslContext = SSLContext.getInstance("TLS");
sslContext.init(kmf.getKeyManagers(), tmf.getTrustManagers(), null);

// Create the engine
SSLEngine engine = sslContext.createSSLengine(hostname, port);

// Use as client
engine.setUseClientMode(true);
```

Generating and Processing TLS Data

The two main `SSLEngine` methods are `wrap()` and `unwrap()`. They are responsible for generating and consuming network data respectively. Depending on the state of the `SSLEngine` object, this data might be handshake or application data.

Each `SSLEngine` object has several phases during its lifetime. Before application data can be sent or received, the TLS protocol requires a handshake to establish cryptographic parameters. This handshake requires a series of back-and-forth steps by the `SSLEngine` object.

During the initial handshaking, the `wrap()` and `unwrap()` methods generate and consume handshake data, and the application is responsible for transporting the data. The `wrap()` and `unwrap()` method sequence is repeated until the handshake is finished. Each `SSLEngine` operation generates an instance of the `SSLEngineResult` class, in which the `SSLEngineResult.HandshakeStatus` field is used to determine what operation must occur next to move the handshake along.

Figure 8-3 shows the state machine during a typical TLS handshake, with corresponding messages and statuses:

Figure 8-3 State Machine during TLS Handshake

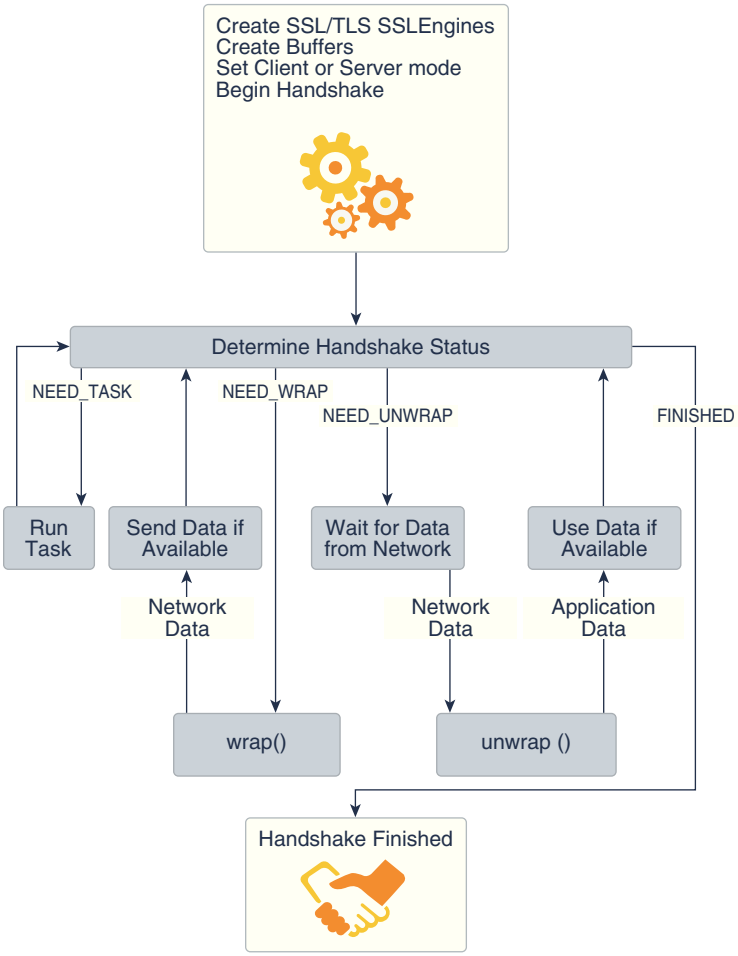

When handshaking is complete, further calls to `wrap()` will attempt to consume application data and package it for transport. The `unwrap()` method will attempt the opposite.

To send data to the peer, the application first supplies the data that it wants to send via `SSLEngine.wrap()` to obtain the corresponding TLS encoded data. The application then sends the encoded data to the peer using its chosen transport mechanism. When the application receives the TLS encoded data from the peer via the transport mechanism, it supplies this data to the `SSLEngine` via `SSLEngine.unwrap()` to obtain the plaintext data sent by the peer.

Example 8-5 Sample Code for Creating a Nonblocking SocketChannel

The following example is an SSL application that uses a non-blocking `SocketChannel` to communicate with its peer. It sends the string "hello" to the peer by encoding it using the `SSLEngine` created in Example 8-4 . It uses information from the `SSLSession` to determine how large to make the byte buffers.

> **Note:**
>
> The example can be made more robust and scalable by using a `Selector` with the nonblocking `SocketChannel`.

```
    // Create a nonblocking socket channel
    SocketChannel socketChannel = SocketChannel.open();
    socketChannel.configureBlocking(false);
    socketChannel.connect(new InetSocketAddress(hostname, port));

    // Complete connection
    while (!socketChannel.finishedConnect()) {
    // do something until connect completed
    }

    //Create byte buffers for holding application and encoded data

    SSLSession session = engine.getSession();
    ByteBuffer myAppData =
ByteBuffer.allocate(session.getApplicationBufferSize());
    ByteBuffer myNetData =
ByteBuffer.allocate(session.getPacketBufferSize());
    ByteBuffer peerAppData =
ByteBuffer.allocate(session.getApplicationBufferSize());
    ByteBuffer peerNetData =
ByteBuffer.allocate(session.getPacketBufferSize());

    // Do initial handshake
    doHandshake(socketChannel, engine, myNetData, peerNetData);

    myAppData.put("hello".getBytes());
    myAppData.flip();

    while (myAppData.hasRemaining()) {
    // Generate TLS/DTLS encoded data (handshake or application data)
    SSLEngineResult res = engine.wrap(myAppData, myNetData);

    // Process status of call
    if (res.getStatus() == SSLEngineResult.Status.OK) {
        myAppData.compact();

        // Send TLS/DTLS encoded data to peer
        while(myNetData.hasRemaining()) {
            int num = socketChannel.write(myNetData);
            if (num == 0) {
                // no bytes written; try again later
            }
        }
    }

    // Handle other status:  BUFFER_OVERFLOW, CLOSED
```

```
...
}
```

Example 8-6 Sample Code for Reading Data From Nonblocking SocketChannel

`SocketChannelSSLEngine`Example 8-4

```
// Read TLS/DTLS encoded data from peer
int num = socketChannel.read(peerNetData);
if (num == -1) {
    // The channel has reached end-of-stream
} else if (num == 0) {
    // No bytes read; try again ...
} else {
    // Process incoming data
    peerNetData.flip();
    res = engine.unwrap(peerNetData, peerAppData);

    if (res.getStatus() == SSLEngineResult.Status.OK) {
        peerNetData.compact();

    if (peerAppData.hasRemaining()) {
        // Use peerAppData
    }
}
// Handle other status:  BUFFER_OVERFLOW, BUFFER_UNDERFLOW, CLOSED
...
}
```

SSLEngine for DTLS Protocols

This section shows you how to create an `SSLEngine` object and use it to handle a DTLS handshake, generate and process DTLS data, and handle retransmissions in DTLS connections.

Creating an SSLEngine Object for DTLS

The following examples illustrate how to create an `SSLEngine` object for DTLS.

> **Note:**
>
> The server name and port number are not used for communicating with the server (all transport is the responsibility of the application). They are hints to the JSSE provider to use for DTLS session caching, and for Kerberos-based cipher suite implementations to determine which server credentials should be obtained.

Example 8-7 Sample Code for Creating an SSLEngine Client for DTLS with PKCS12 as Keystore

The following sample code creates an `SSLEngine` client for DTLS that uses PKCS12 as keystore:

```
import javax.net.ssl.*;
import java.security.*;

// Create and initialize the SSLContext with key material
char[] passphrase = "passphrase".toCharArray();

// First initialize the key and trust material
KeyStore ksKeys = KeyStore.getInstance("PKCS12");
ksKeys.load(new FileInputStream("testKeys"), passphrase);
KeyStore ksTrust = KeyStore.getInstance("PKCS12");
ksTrust.load(new FileInputStream("testTrust"), passphrase);

// KeyManagers decide which key material to use
KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
kmf.init(ksKeys, passphrase);

// TrustManagers decide whether to allow connections
TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
tmf.init(ksTrust);

// Get an instance of SSLContext for DTLS protocols
sslContext = SSLContext.getInstance("DTLS");
sslContext.init(kmf.getKeyManagers(), tmf.getTrustManagers(), null);

// Create the engine
SSLEngine engine = sslContext.createSSLengine(hostname, port);

// Use engine as client
engine.setUseClientMode(true);
```

Example 8-8 Sample Code for Creating an SSLEngine Server for DTLS with PKCS12 as Keystore

`SSLEngine`

```
import javax.net.ssl.*;
import java.security.*;

// Create and initialize the SSLContext with key material
char[] passphrase = "passphrase".toCharArray();

// First initialize the key and trust material
KeyStore ksKeys = KeyStore.getInstance("PKCS12");
ksKeys.load(new FileInputStream("testKeys"), passphrase);
KeyStore ksTrust = KeyStore.getInstance("PKCS12");
ksTrust.load(new FileInputStream("testTrust"), passphrase);

// KeyManagers decide which key material to use
KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
```

```
kmf.init(ksKeys, passphrase);

// TrustManagers decide whether to allow connections
TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
tmf.init(ksTrust);

// Get an SSLContext for DTLS Protocol without authentication
sslContext = SSLContext.getInstance("DTLS");
sslContext.init(null, null, null);

// Create the engine
SSLEngine engine = sslContext.createSSLeEngine(hostname, port);

// Use the engine as server
engine.setUseClientMode(false);

// Require client authentication
engine.setNeedClientAuth(true);
```

Generating and Processing DTLS Data

A DTLS handshake and a TLS handshake generate and process data similarly. (See Generating and Processing TLS Data.) They both use the `SSLEngine.wrap()` and `SSLEngine.wrap()` methods to generate and consume network data, respectively.

The following diagram shows the state machine during a typical DTLS handshake, with corresponding messages and statuses:

Figure 8-4 State Machine during DTLS Handshake

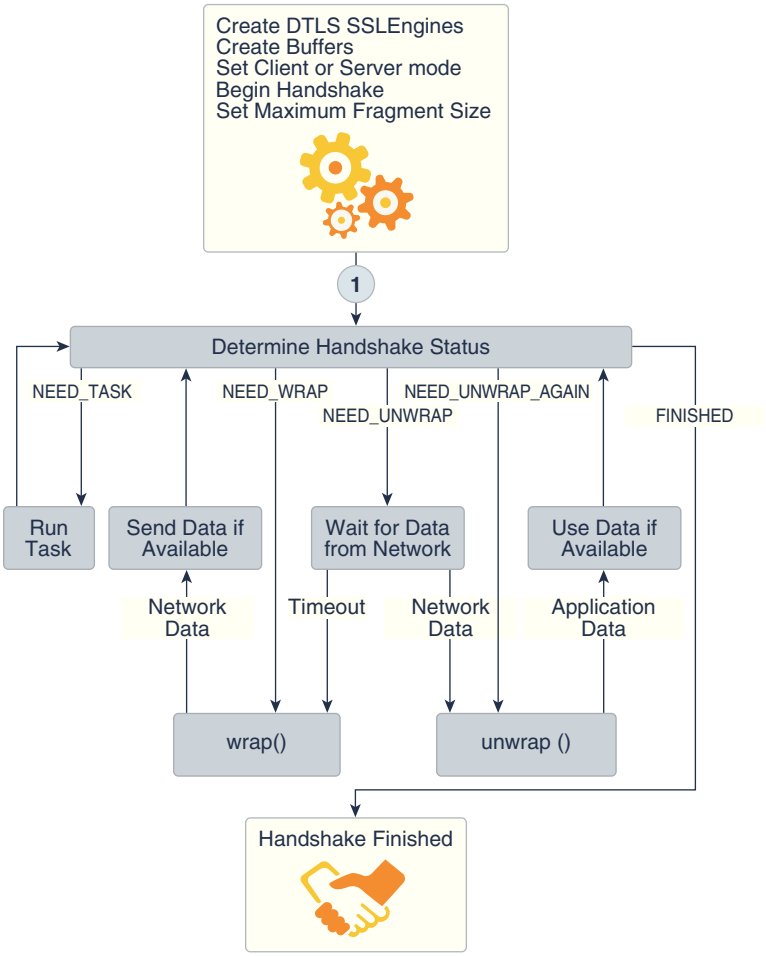

Example 8-9 Sample Code for Handling DTLS Handshake Status and Overall Status

This sample demonstrates how to handle DTLS handshake status (from the `SSLEngine.getHandshakeStatus` method) and overall status (from the `SSLEngineResult.getStatus` method).

```
void handshake(SSLEngine engine, DatagramSocket socket, SocketAddress
peerAddr) throws Exception {
    boolean endLoops = false;
    // private static int MAX_HANDSHAKE_LOOPS = 60;
    int loops = MAX_HANDSHAKE_LOOPS;
    engine.beginHandshake();
    while (!endLoops && (serverException == null) && (clientException ==
null)) {
        if (--loops < 0) {
            throw new RuntimeException("Too many loops to produce
handshake packets");
        }
        SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
        if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP ||
```

```
                hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP_AGAIN) {
            ByteBuffer iNet;
            ByteBuffer iApp;
            if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP) {
                // Receive ClientHello request and other SSL/TLS/DTLS
records
                byte[] buf = new byte[1024];
                DatagramPacket packet = new DatagramPacket(buf,
buf.length);
                try {
                    socket.receive(packet);
                } catch (SocketTimeoutException ste) {
                    // Retransmit the packet if timeout
                    List <Datagrampacket> packets =
onReceiveTimeout(engine, peerAddr);
                    for (DatagramPacket p : packets) {
                        socket.send(p);
                    }
                    continue;
                }
                iNet = ByteBuffer.wrap(buf, 0, packet.getLength());
                iApp = ByteBuffer.allocate(1024);
            } else {
                iNet = ByteBuffer.allocate(0);
                iApp = ByteBuffer.allocate(1024);
            }
            SSLEngineResult r = engine.unwrap(iNet, iApp);
            SSLEngineResult.Status rs = r.getStatus();
            hs = r.getHandshakeStatus();
            if (rs == SSLEngineResult.Status.BUFFER_OVERFLOW) {
                // The client maximum fragment size config does not work?
                throw new Exception("Buffer overflow: " +
                                    "incorrect client maximum fragment
size");
            } else if (rs == SSLEngineResult.Status.BUFFER_UNDERFLOW) {
                // Bad packet, or the client maximum fragment size
                // config does not work?
                if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING)
{
                    throw new Exception("Buffer underflow: " +
                                        "incorrect client maximum fragment
size");
                } // Otherwise, ignore this packet
            } else if (rs == SSLEngineResult.Status.CLOSED) {
                endLoops = true;
            } // Otherwise, SSLEngineResult.Status.OK
            if (rs != SSLEngineResult.Status.OK) {
                continue;
            }
        } else if (hs == SSLEngineResult.HandshakeStatus.NEED_WRAP) {
            // Call a function to produce handshake packets
            List <DatagramPacket> packets =
produceHandshakePackets(engine, peerAddr);
            for (DatagramPacket p : packets) {
                socket.send(p);
```

```
            }
        } else if (hs == SSLEngineResult.HandshakeStatus.NEED_TASK) {
            runDelegatedTasks(engine);
        } else if (hs == SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
            // OK, time to do application data exchange
            endLoops = true;
        } else if (hs == SSLEngineResult.HandshakeStatus.FINISHED) {
            endLoops = true;
        }
    }
    SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
    if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
        throw new Exception("Not ready for application data yet");
    }
}
```

Difference Between the TLS and DTLS SSLEngine.wrap() Methods

The `SSLEngine.wrap()` method for DTLS is different from TLS as follows:

- In the TLS implementation of `SSLEngine`, the output buffer of `SSLEngine.wrap()` contains one or more TLS records (due to the TLSv1 BEAST Cipher Block Chaining vulnerability).
- In the DTLS implementation of `SSLEngine`, the output buffer of `SSLEngine.wrap()` contains at most one record, so that every DTLS record can be marshaled and delivered to the datagram layer individually.

> **Note:**
>
> Each record produced by `SSLEngine.wrap()` should comply to the maximum packet size limitation as specified by `SSLParameters.getMaximumPacketSize()`.

Handling Retransmissions in DTLS Connections

In SSL/TLS over a reliable connection, data is guaranteed to arrive in the proper order, and retransmission is unnecessary. However, for DTLS, which often works over unreliable media, missing or delayed handshake messages must be retransmitted.

The `SSLEngine` class operates in a completely transport-neutral manner, and the application layer performs all I/O. Because the `SSLEngine` class isn't responsible for I/O, the application instead is responsible for providing timers and signalling the `SSLEngine` class when a retransmission is needed. The application layer must determine the right timeout value and when to trigger the timeout event. During handshaking, if an `SSLEngine` object is in `HandshakeStatus.NEED_UNWRAP` state, a call to `SSLEngine.wrap()` means that the previous packets were lost, and must be retransmitted. For such cases, the DTLS implementation of the `SSLEngine` class takes the responsibility to wrap the previous necessary handshaking messages again if necessary.

> **Note:**
>
> In a DTLS engine, only handshake messages must be properly exchanged. Application data can handle packet loss without the need for timers.

Handling Retransmission in an Application

`SSLEngine.unwrap()` and `SSLEngine.wrap()` can be used together to handle retransmission in an application.

Figure 8-5 shows a typical scenario for handling DTLS handshaking retransmission:

Figure 8-5 DTLS Handshake Retransmission State Flow

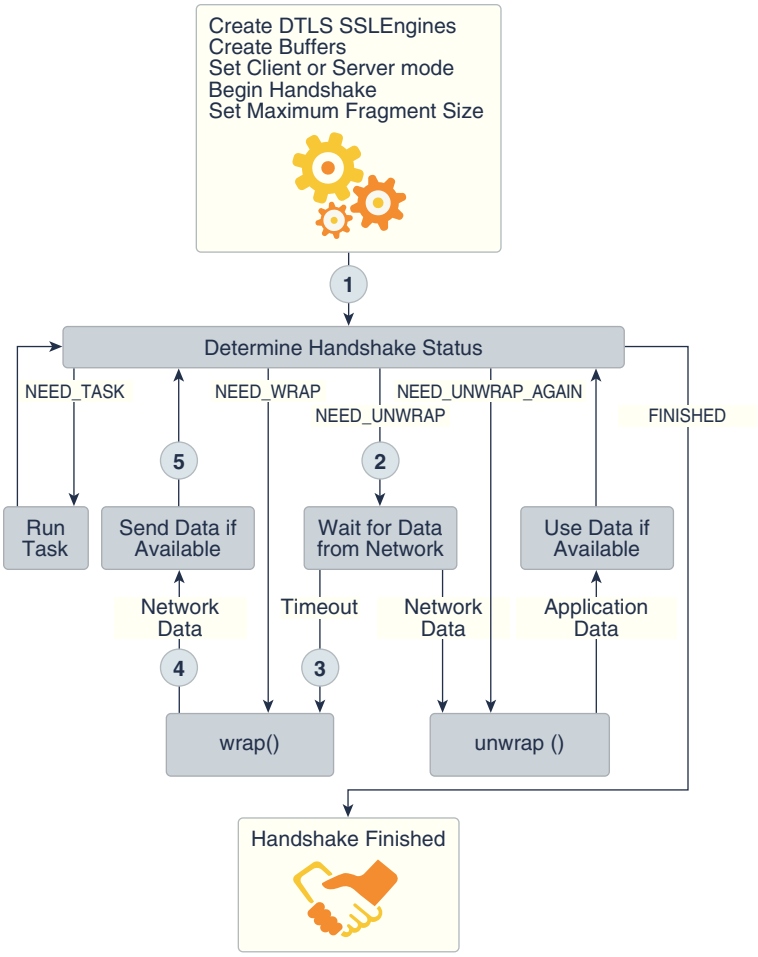

1. Create and initialize an instance of DTLS `SSLEngine`.

   See Creating an SSLEngine Object. The DTLS handshake process begins.

2. If the handshake status is `HandshakeStatus.NEED_UNWRAP`, wait for data from network.

3. If the timer times out, it indicates that the previous delivered handshake messages may have been lost.

> **Note:**
>
> In DTLS handshaking retransmission, the determined handshake status isn't necessarily `HandshakeStatus.NEED_WRAP` for the call to `SSLEngine.wrap()`.

4. Call `SSLEngine.wrap()`.
5. The wrapped packets are delivered.

Handling a Buffered Handshake Message in an Application

Datagram transport doesn't require or provide reliable or in-order delivery of data. Handshake messages may be lost or need to be reordered. In the DTLS implementation, a handshake message may need to be buffered for future handling before all previous messages have been received.

The DTLS implementation of `SSLEngine` takes the responsibility to reorder handshake messages. Handshake message buffering and reordering are transparent to applications.

However, applications must manage `HandshakeStatus.NEED_UNWRAP_AGAIN` status. This status indicates that for the next `SSLEngine.unwrap()` operation no additional data from the remote side is required.

Figure 8-6 shows a typical scenario for using the `HandshakeStatus.NEED_UNWRAP_AGAIN`.

Figure 8-6 State Machine of DTLS Buffered Handshake with NEED_UNWRAP_AGAIN

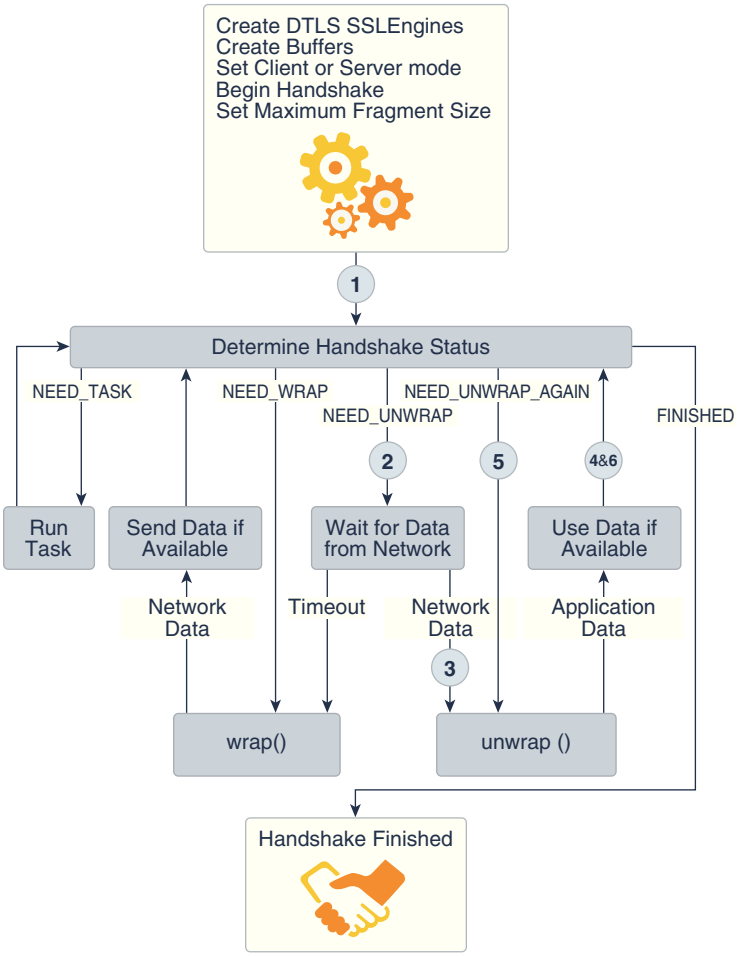

1. Create and initialize an instance of DTLS `SSLEngine`.

   See Creating an SSLEngine Object.

2. Optional: If the handshake status is `HandshakeStatus.NEED_UNWRAP`, wait for data from network.

3. Optional: If you received the network data, call `SSLEngine.unwrap()`.

4. Determine the handshake status for next processing. The handshake status can be `HandshakeStatus.NEED_UNWRAP_AGAIN`, `HandshakeStatus.NEED_UNWRAP`, or `HandshakeStatus.NEED_WRAP`.

   - If the handshake status is `HandshakeStatus.NEED_UNWRAP_AGAIN`, call `SSLEngine.unwrap()`.

> **Note:**
>
> For `HandshakeStatus.NEED_UNWRAP_AGAIN` status, no additional data from the network is required for an `SSLEngine.unwrap()` operation.

5. Determine the handshake status for further processing. The handshake status can be `HandshakeStatus.NEED_UNWRAP_AGAIN`, `HandshakeStatus.NEED_UNWRAP`, or `HandshakeStatus.NEED_WRAP`.

Dealing With Blocking Tasks

During handshaking, an `SSLEngine` might encounter tasks that can block or take a long time. For example, a `TrustManager` may need to connect to a remote certificate validation service, or a `KeyManager` might need to prompt a user to determine which certificate to use as part of client authentication. To preserve the nonblocking nature of `SSLEngine`, when the engine encounters such a task, it will return `SSLEngineResult.HandshakeStatus.NEED_TASK`. Upon receiving this status, the application should invoke `SSLEngine.getDelegatedTask()` to get the task, and then, using the threading model appropriate for its requirements, process the task. The application might, for example, obtain threads from a thread pool to process the tasks, while the main thread handles other I/O.

The following code executes each task in a newly created thread:

```
if (res.getHandshakeStatus() == SSLEngineResult.HandshakeStatus.NEED_TASK)
{
    Runnable task;
    while ((task = engine.getDelegatedTask()) != null) {
        new Thread(task).start();
    }
}
```

The `SSLEngine` will block future `wrap()` and `unwrap()` calls until all of the outstanding tasks are completed.

Shutting Down a TLS/DTLS Connection

For an orderly shutdown of a TLS/DTLS connection, the TLS/DTLS protocols require transmission of close messages. Therefore, when an application is done with the TLS/DTLS connection, it should first obtain the close messages from the `SSLEngine`, then transmit them to the peer using its transport mechanism, and finally shut down the transport mechanism. Example 8-10 illustrates this.

In addition to an application explicitly closing the `SSLEngine`, the `SSLEngine` might be closed by the peer (via receipt of a close message while it is processing handshake data), or by the `SSLEngine` encountering an error while processing application or handshake data, indicated by throwing an `SSLException`. In such cases, the application should invoke `SSLEngine.wrap()` to get the close message and send it to the peer until `SSLEngine.isOutboundDone()` returns `true` (as shown in Example 8-10), or until the `SSLEngineResult.getStatus()` returns `CLOSED`.

In addition to orderly shutdowns, there can also be unexpected shutdowns when the transport link is severed before close messages are exchanged. In the previous examples, the application might get `-1` or `IOException` when trying to read from the nonblocking `SocketChannel`, or get `IOException` when trying to write to the non-blocking `SocketChannel`. When you get to the end of your input data, you should call `engine.closeInbound()`, which will verify with the `SSLEngine` that the remote peer has closed cleanly from the TLS/DTLS perspective. Then the application should still try to shut down cleanly by using the procedure in Example 8-10. Obviously, unlike

SSLSocket, the application using SSLEngine must deal with more state transitions, statuses, and programming. See Sample Code Illustrating the Use of an SSLEngine.

Example 8-10 Sample Code for Shutting Down a SSL/TLS/DTLS Connection

The following code sample illustrates how to shut down a TLS/DTLS connection:

```
// Indicate that application is done with engine
engine.closeOutbound();

while (!engine.isOutboundDone()) {
    // Get close message
    SSLEngineResult res = engine.wrap(empty, myNetData);

    // Check res statuses

    // Send close message to peer
    while(myNetData.hasRemaining()) {
        int num = socketChannel.write(myNetData);
        if (num == 0) {
            // no bytes written; try again later
        }
        myNetData().compact();
    }
}

// Close transport
socketChannel.close();
```

SSLSession and ExtendedSSLSession

The javax.net.ssl.SSLSession interface represents a security context negotiated between the two peers of an SSLSocket or SSLEngine connection. After a session has been arranged, it can be shared by future SSLSocket or SSLEngine objects connected between the same two peers.

In some cases, parameters negotiated during the handshake are needed later in the handshake to make decisions about trust. For example, the list of valid signature algorithms might restrict the certificate types that can be used for authentication. The SSLSession can be retrieved *during* the handshake by calling getHandshakeSession() on an SSLSocket or SSLEngine. Implementations of TrustManager or KeyManager can use the getHandshakeSession() method to get information about session parameters to help them make decisions.

A fully initialized SSLSession contains the cipher suite that will be used for communications over a secure socket as well as a nonauthoritative hint as to the network address of the remote peer, and management information such as the time of creation and last use. A session also contains a shared master secret negotiated between the peers that is used to create cryptographic keys for encrypting and guaranteeing the integrity of the communications over an SSLSocket or SSLEngine connection. The value of this master secret is known only to the underlying secure socket implementation and is not exposed through the SSLSession API.

ExtendedSSLSession extends the SSLSession interface to support additional session attributes. The ExtendedSSLSession class adds methods that describe the signature

algorithms that are supported by the local implementation and the peer. The `getRequestedServerNames()` method called on an `ExtendedSSLSession` instance is used to obtain a list of `SNIServerName` objects in the requested Server Name Indication (SNI) Extension. The server should use the requested server names to guide its selection of an appropriate authentication certificate, and/or other aspects of the security policy. The client should use the requested server names to guide its endpoint identification of the peer's identity, and/or other aspects of the security policy.

Calls to the `getPacketBufferSize()` and `getApplicationBufferSize()` methods on `SSLSession` are used to determine the appropriate buffer sizes used by `SSLEngine`.

> **Note:**
>
> The TLS protocols specify that implementations are to produce packets containing at most 16 kilobytes (KB) of plain text. However, some implementations violate the specification and generate large records up to 32 KB. If the `SSLEngine.unwrap()` code detects large inbound packets, then the buffer sizes returned by `SSLSession` will be updated dynamically. Applications should always check the BUFFER_OVERFLOW and BUFFER_UNDERFLOW statuses and enlarge the corresponding buffers if necessary. See Understanding SSLEngine Operation Statuses. SunJSSE will always send standard compliant 16 KB records and allow incoming 32 KB records. For a workaround, see the System property `jsse.SSLEngine.acceptLargeFragments` in Customizing JSSE.

HttpsURLConnection Class

The `javax.net.ssl.HttpsURLConnection` class extends the `java.net.HttpURLConnection` class and adds support for HTTPS-specific features.

The HTTPS protocol is similar to HTTP, but HTTPS first establishes a secure channel through TLS sockets and then verifies the identity of the peer (see Cipher Suite Choice and Remote Entity Verification) before requesting or receiving data. The `javax.net.ssl.HttpsURLConnection` class extends the `java.net.HttpURLConnection` class and adds support for HTTPS-specific features. To know more about how HTTPS URLs are constructed and used, see the`java.net.URL`, `java.net.URLConnection`, `java.net.HttpURLConnection`, and `javax.net.ssl.HttpsURLConnection` classes.

Upon obtaining an `HttpsURLConnection` instance, you can configure a number of HTTP and HTTPS parameters before actually initiating the network connection via the `URLConnection.connect()` method. Of particular interest are:

- Setting the Assigned SSLSocketFactory
- Setting the Assigned HostnameVerifier

Setting the Assigned SSLSocketFactory

In some situations, it is desirable to specify the `SSLSocketFactory` that an `HttpsURLConnection` instance uses. For example, you might want to tunnel through a proxy type that is not supported by the default implementation. The new

`SSLSocketFactory` could return sockets that have already performed all necessary tunneling, thus allowing `HttpsURLConnection` to use additional proxies.

The `HttpsURLConnection` class has a default `SSLSocketFactory` that is assigned when the class is loaded (this is the factory returned by the `SSLSocketFactory.getDefault()` method). Future instances of `HttpsURLConnection` will inherit the current default `SSLSocketFactory` until a new default `SSLSocketFactory` is assigned to the class via the static `HttpsURLConnection.setDefaultSSLSocketFactory()` method. Once an instance of `HttpsURLConnection` has been created, the inherited `SSLSocketFactory` on this instance can be overridden with a call to the `setSSLSocketFactory()` method.

> **Note:**
>
> Changing the default static `SSLSocketFactory` has no effect on existing instances of `HttpsURLConnection`. A call to the `setSSLSocketFactory()` method is necessary to change the existing instances.

You can obtain the per-instance or per-class `SSLSocketFactory` by making a call to the `getSSLSocketFactory()` or `getDefaultSSLSocketFactory()` method, respectively.

Setting the Assigned HostnameVerifier

If the host name of the URL does not match the host name in the credentials received as part of the TLS handshake, then it is possible that URL spoofing has occurred. If the implementation cannot determine a host name match with reasonable certainty, then the TLS implementation performs a callback to the instance's assigned `HostnameVerifier` for further checking. The host name verifier can take whatever steps are necessary to make the determination, such as performing host name pattern matching or perhaps opening an interactive dialog box. An unsuccessful verification by the host name verifier closes the connection. For more information regarding host name verification, see RFC 2818: HTTP over TLS.

The `setHostnameVerifier()` and `setDefaultHostnameVerifier()` methods operate in a similar manner to the `setSSLSocketFactory()` and `setDefaultSSLSocketFactory()` methods, in that `HostnameVerifier` objects are assigned on a per-instance and per-class basis, and the current values can be obtained by a call to the `getHostnameVerifier()` or `getDefaultHostnameVerifier()` method.

Support Classes and Interfaces

The classes and interfaces in this section are provided to support the creation and initialization of `SSLContext` objects, which are used to create `SSLSocketFactory`, `SSLServerSocketFactory`, and `SSLEngine` objects. The support classes and interfaces are part of the `javax.net.ssl` package.

Three of the classes described in this section (SSLContext Class, KeyManagerFactory Class, and TrustManagerFactory Class) are *engine classes*. An engine class is an API class for specific algorithms (or protocols, in the case of `SSLContext`), for which

implementations may be provided in one or more Cryptographic Service Provider (provider) packages. See JCA Design Principles and Engine Classes and Algorithms.

The SunJSSE provider that comes standard with JSSE provides `SSLContext`, `KeyManagerFactory`, and `TrustManagerFactory` implementations, as well as implementations for engine classes in the standard `java.security` API. Table 8-1 lists implementations supplied by SunJSSE.

Table 8-1 Implementations Supplied by SunJSSE

Engine Class Implemented	Algorithm or Protocol
`KeyStore`	PKCS12
`KeyManagerFactory`	PKIX, SunX509
`TrustManagerFactory`	PKIX (X509 or SunPKIX), SunX509
`SSLContext`	SSLv3[1], TLSv1, TLSv1.1, TLSv1.2, TLSv1.3, DTLSv1.0, DTLSv1.2

[1] Starting with JDK 8u31, the SSLv3 protocol (Secure Socket Layer) has been deactivated and is not available by default. See the `java.security.Security` property `jdk.tls.disabledAlgorithms` in the `<`*`java_home`*`>/conf/security/java.security` file. If SSLv3 is absolutely required, the protocol can be reactivated by removing SSLv3 from the `jdk.tls.disabledAlgorithms` property in the `java.security` file or by dynamically setting this Security Property before JSSE is initialized.

SSLContext Class

The `javax.net.ssl.SSLContext` class is an engine class for an implementation of a secure socket protocol. An instance of this class acts as a factory for `SSLSocket`, `SSLServerSocket`, and `SSLEngine`. An `SSLContext` object holds all of the state information shared across all objects created under that context. For example, session state is associated with the `SSLContext` when it is negotiated through the handshake protocol by sockets created by socket factories provided by the context. These cached sessions can be reused and shared by other sockets created under the same context.

Each instance is configured through its `init` method with the keys, certificate chains, and trusted root CA certificates that it needs to perform authentication. This configuration is provided in the form of key and trust managers. These managers provide support for the authentication and key agreement aspects of the cipher suites supported by the context.

Currently, only X.509-based managers are supported.

Obtaining and Initializing the SSLContext Class

The `SSLContext` class is used to create the `SSLSocketFactory` or `SSLServerSocketFactory` class.

There are two ways to obtain and initialize an `SSLContext`:

- The simplest way is to call the static `SSLContext.getDefault` method on either the `SSLSocketFactory` or `SSLServerSocketFactory` class. This method creates a default `SSLContext` with a default `KeyManager`, `TrustManager`, and `SecureRandom` (a secure random number generator). A default `KeyManagerFactory` and `TrustManagerFactory` are used to create the `KeyManager` and `TrustManager`, respectively. The key material used is found in the default keystore and truststore,

as determined by system properties described in Customizing the Default Keystores and Truststores, Store Types, and Store Passwords.

- The approach that gives the caller the most control over the behavior of the created context is to call the static method `SSLContext.getDefault` on the `SSLContext` class, and then initialize the context by calling the instance's proper `init()` method. One variant of the `init()` method takes three arguments: an array of `KeyManager` objects, an array of `TrustManager` objects, and a `SecureRandom` object. The `KeyManager` and `TrustManager` objects are created by either implementing the appropriate interfaces or using the `KeyManagerFactory` and `TrustManagerFactory` classes to generate implementations. The `KeyManagerFactory` and `TrustManagerFactory` can then each be initialized with key material contained in the `KeyStore` passed as an argument to the `init()` method of the `TrustManagerFactory` or `KeyManagerFactory` classes. Finally, the `getTrustManagers()` method (in `TrustManagerFactory`) and `getKeyManagers()` method (in `KeyManagerFactory`) can be called to obtain the array of trust managers or key managers, one for each type of trust or key material.

Once a TLS connection is established, an `SSLSession` is created which contains various information, such as identities established and cipher suite used. The `SSLSession` is then used to describe an ongoing relationship and state information between two entities. Each TLS connection involves one session at a time, but that session may be used on many connections between those entities, simultaneously or sequentially.

Creating an SSLContext Object

Like other JCA provider-based engine classes, `SSLContext` objects are created using the `getInstance()` factory methods of the `SSLContext` class. These static methods each return an instance that implements *at least* the requested secure socket protocol. The returned instance may implement other protocols, too. For example, `getInstance("TLSv1")` may return an instance that implements TLSv1, TLSv1.1, and TLSv1.2. The `getSupportedProtocols()` method returns a list of supported protocols when an `SSLSocket`, `SSLServerSocket`, or `SSLEngine` is created from this context. You can control which protocols are actually enabled for an SSL connection by using the `setEnabledProtocols(String[] protocols)` method.

> **Note:**
>
> An `SSLContext` object is automatically created, initialized, and statically assigned to the `SSLSocketFactory` class when you call the `SSLSocketFactory.getDefault()` method. Therefore, you do not have to directly create and initialize an `SSLContext` object (unless you want to override the default behavior).

To create an `SSLContext` object by calling the `getInstance()` factory method, you must specify the protocol name. You may also specify which provider you want to supply the implementation of the requested protocol:

- `public static SSLContext getInstance(String protocol);`
- `public static SSLContext getInstance(String protocol, String provider);`

- `public static SSLContext getInstance(String protocol, Provider provider);`

If just a protocol name is specified, then the system will determine whether an implementation of the requested protocol is available in the environment. If there is more than one implementation, then it will determine whether there is a preferred one.

If both a protocol name and a provider are specified, then the system will determine whether an implementation of the requested protocol is in the provider requested. If there is no implementation, an exception will be thrown.

A protocol is a string (such as `"TLS"`) that describes the secure socket protocol desired. Common protocol names for `SSLContext` objects are defined in Java Security Standard Algorithm Names.

An `SSLContext` can be obtained as follows:

```
SSLContext sc = SSLContext.getInstance("TLS");
```

A newly created `SSLContext` should be initialized by calling the `init` method:

```
public void init(KeyManager[] km, TrustManager[] tm, SecureRandom random);
```

If the `KeyManager[]` parameter is null, then an empty `KeyManager` will be defined for this context. If the `TrustManager[]` parameter is null, then the installed security providers will be searched for the highest-priority implementation of the `TrustManagerFactory` class (see TrustManagerFactory Class), from which an appropriate `TrustManager` will be obtained. Likewise, the `SecureRandom` parameter may be null, in which case a default implementation will be used.

If the internal default context is used, (for example, an `SSLContext` is created by `SSLSocketFactory.getDefault()` or `SSLServerSocketFactory.getDefault()`), then a default KeyManager and TrustManager are created. The default `SecureRandom` implementation is also chosen.

TrustManager Interface

The primary responsibility of the `TrustManager` is to determine whether the presented authentication credentials should be trusted. If the credentials are not trusted, then the connection will be terminated. To authenticate the remote identity of a secure socket peer, you must initialize an `SSLContext` object with one or more `TrustManager` objects. You must pass one `TrustManager` for each authentication mechanism that is supported. If null is passed into the `SSLContext` initialization, then a trust manager will be created for you. Typically, a single trust manager supports authentication based on X.509 public key certificates (for example, `X509TrustManager`). Some secure socket implementations may also support authentication based on shared secret keys, Kerberos, or other mechanisms.

`TrustManager` objects are created either by a `TrustManagerFactory`, or by providing a concrete implementation of the interface.

TrustManagerFactory Class

The `javax.net.ssl.TrustManagerFactory` is an engine class for a provider-based service that acts as a factory for one or more types of `TrustManager` objects. Because

it is provider-based, additional factories can be implemented and configured to provide additional or alternative trust managers that provide more sophisticated services or that implement installation-specific authentication policies.

Creating a TrustManagerFactory

You create an instance of this class in a similar manner to `SSLContext`, except for passing an algorithm name string instead of a protocol name to the `getInstance()` method:

```
TrustManagerFactory tmf = TrustManagerFactory.getInstance(String
algorithm);
TrustManagerFactory tmf = TrustManagerFactory.getInstance(String
algorithm, String provider);
TrustManagerFactory tmf = TrustManagerFactory.getInstance(String
algorithm, Provider provider);
```

A sample call is as follows:

```
TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX",
"SunJSSE");
```

The preceding call creates an instance of the SunJSSE provider's PKIX trust manager factory. This factory can be used to create trust managers that provide X.509 PKIX-based certification path validity checking.

When initializing an `SSLContext`, you can use trust managers created from a trust manager factory, or you can write your own trust manager, for example, using the `CertPath` API. See Java PKI Programmer's Guide. You do not need to use a trust manager factory if you implement a trust manager using the `X509TrustManager` interface.

A newly created factory should be initialized by calling one of the `init()` methods:

```
public void init(KeyStore ks);
public void init(ManagerFactoryParameters spec);
```

Call whichever `init()` method is appropriate for the `TrustManagerFactory` you are using. If you are not sure, then ask the provider vendor.

For many factories, such as the SunX509 `TrustManagerFactory` from the SunJSSE provider, the `KeyStore` is the only information required to initialize the `TrustManagerFactory` and thus the first `init` method is the appropriate one to call. The `TrustManagerFactory` will query the `KeyStore` for information about which remote certificates should be trusted during authorization checks.

Sometimes, initialization parameters other than a `KeyStore` are needed by a provider. Users of that provider are expected to pass an implementation of the appropriate `ManagerFactoryParameters` as defined by the provider. The provider can then call the specified methods in the `ManagerFactoryParameters` implementation to obtain the needed information.

For example, suppose the `TrustManagerFactory` provider requires initialization parameters B, R, and S from any application that wants to use that provider. Like all

providers that require initialization parameters other than a `KeyStore`, the provider requires the application to provide an instance of a class that implements a particular `ManagerFactoryParameters` subinterface. In the example, suppose that the provider requires the calling application to implement and create an instance of `MyTrustManagerFactoryParams` and pass it to the second `init()` method. The following example illustrates what `MyTrustManagerFactoryParams` can look like:

```
public interface MyTrustManagerFactoryParams extends
ManagerFactoryParameters {
  public boolean getBValue();
  public float getRValue();
  public String getSValue();
}
```

Some trust managers can make trust decisions without being explicitly initialized with a `KeyStore` object or any other parameters. For example, they may access trust material from a local directory service via LDAP, use a remote online certificate status checking server, or access default trust material from a standard local location.

Check TrustManagerFactory Object's Certificates' Expiration Date

The following method, `filterTrustAnchors`, filters a set of trust anchors, removing those whose certificates will expire by the specified date, then initializes a `TrustManagerFactory` object with this set.

```
    public static void filterTrustAnchors (
        String truststore, String password, String validityDate)
        throws FileNotFoundException, KeyStoreException, IOException,
            ParseException, NoSuchAlgorithmException,
            InvalidAlgorithmParameterException, CertificateException,
            KeyManagementException {

        FileInputStream is = new FileInputStream(truststore);
        KeyStore keystore =
KeyStore.getInstance(KeyStore.getDefaultType());
        keystore.load(is, password.toCharArray());
        PKIXParameters params = new PKIXParameters(keystore);

        // Obtain CA root certificates
        Set<TrustAnchor> myTrustAnchors = params.getTrustAnchors();

        // Create new set of CA certificates that are still valid for
        // specified date
        Set<TrustAnchor> validTrustAnchors =
            myTrustAnchors.stream().filter(
                ta -> {
                    try {
                        ta.getTrustedCert().checkValidity(

DateFormat.getDateInstance().parse(validityDate));
                    } catch (CertificateException | ParseException e) {
                        return false;
                    }
                    return true; }).collect(Collectors.toSet());
```

```
        // Create PKIXBuilderParameters parameters
        PKIXBuilderParameters pkixParams =
            new PKIXBuilderParameters(validTrustAnchors, new
X509CertSelector());

        // Wrap PKIX parameters as trust manager parameters
        ManagerFactoryParameters trustParams =
            new CertPathTrustManagerParameters(pkixParams);

        // Create TrustManagerFactory for PKIX-compliant trust managers
        TrustManagerFactory factory =
TrustManagerFactory.getInstance("PKIX");

        // Pass parameters to factory to be passed to CertPath
implementation
        factory.init(trustParams);

        // Use factory
        SSLContext ctx = SSLContext.getInstance("TLS");
        ctx.init(null, factory.getTrustManagers(), null);
    }
```

PKIX TrustManager Support

The default trust manager algorithm is PKIX. It can be changed by editing the `ssl.TrustManagerFactory.algorithm` property in the `java.security` file.

The PKIX trust manager factory uses the CertPath PKIX implementation (see PKI Programmer's Guide Overview) from an installed security provider. The trust manager factory can be initialized using the normal `init(KeyStores)` method, or by passing `CertPath` parameters to the PKIX trust manager using the `CertPathTrustManagerParameters` class.

Example 8-11 illustrates how to get the trust manager to use a particular LDAP certificate store and enable revocation checking.

If the `TrustManagerFactory.init(KeyStore)` method is used, then default PKIX parameters are used with the exception that revocation checking is disabled. It can be enabled by setting the `com.sun.net.ssl.checkRevocation` system property to `true`. This setting requires that the CertPath implementation can locate revocation information by itself. The PKIX implementation in the provider can do this in many cases but requires that the system property `com.sun.security.enableCRLDP` be set to `true`. Note that the `TrustManagerFactory.init(ManagerFactoryParameters)` method has revocation checking enabled by default.

See PKIX Classes and The CertPath Class.

Example 8-11 Sample Code for Using a LDAP Certificate to Enable Revocation Checking

The following example illustrates how to get the trust manager to use a particular LDAP certificate store and enable revocation checking:

```
    import javax.net.ssl.*;
    import java.security.cert.*;
    import java.security.KeyStore;
    import java.io.FileInputStream;
    ...

    // Obtain Keystore password
    char[] pass = System.console().readPassword("Password: ");

    // Create PKIX parameters
    KeyStore anchors = KeyStore.getInstance("JKS");
    anchors.load(new FileInputStream(anchorsFile, pass));
    PKIXBuilderParameters pkixParams = new PKIXBuilderParameters(anchors,
new X509CertSelector());

    // Specify LDAP certificate store to use
    LDAPCertStoreParameters lcsp = new
LDAPCertStoreParameters("ldap.imc.org", 389);
    pkixParams.addCertStore(CertStore.getInstance("LDAP", lcsp));

    // Specify that revocation checking is to be enabled
    pkixParams.setRevocationEnabled(true);

    // Wrap PKIX parameters as trust manager parameters
    ManagerFactoryParameters trustParams = new
CertPathTrustManagerParameters(pkixParams);

    // Create TrustManagerFactory for PKIX-compliant trust managers
    TrustManagerFactory factory = TrustManagerFactory.getInstance("PKIX");

    // Pass parameters to factory to be passed to CertPath implementation
    factory.init(trustParams);

    // Use factory
    SSLContext ctx = SSLContext.getInstance("TLS");
    ctx.init(null, factory.getTrustManagers(), null);
```

X509TrustManager Interface

The `javax.net.ssl.X509TrustManager` interface extends the general `TrustManager` interface. It must be implemented by a trust manager when using X.509-based authentication.

To support X.509 authentication of remote socket peers through JSSE, an instance of this interface must be passed to the `init` method of an `SSLContext` object.

Creating an X509TrustManager

You can either implement this interface directly yourself or obtain one from a provider-based `TrustManagerFactory` (such as that supplied by the SunJSSE provider). You could also implement your own interface that delegates to a factory-generated trust manager. For example, you might do this to filter the resulting trust decisions and query an end-user through a graphical user interface.

If a null KeyStore parameter is passed to the SunJSSE PKIX or SunX509 `TrustManagerFactory`, then the factory uses the following process to try to find trust material:

1. If the `javax.net.ssl.trustStore` property is defined, then the `TrustManagerFactory` attempts to find a file using the file name specified by that system property, and uses that file for the KeyStore parameter. If the `javax.net.ssl.trustStorePassword` system property is also defined, then its value is used to check the integrity of the data in the truststore before opening it.

   If the `javax.net.ssl.trustStore` property is defined but the specified file does not exist, then a default `TrustManager` using an empty keystore is created.

2. If the `javax.net.ssl.trustStore` system property was not specified, then:
   - if the file *java-home*`/lib/security/jssecacerts` exists, that file is used;
   - if the file *java-home*`/lib/security/cacerts` exists, that file is used;
   - if neither of these files exists, then the TLS cipher suite is anonymous, does not perform any authentication, and thus does not need a truststore.

To know more about what *java-home* refers to, see Terms and Definitions.

The factory looks for a file specified via the `javax.net.ssl.trustStore` Security Property or for the `jssecacerts` file before checking for a `cacerts` file. Therefore, you can provide a JSSE-specific set of trusted root certificates separate from ones that might be present in `cacerts` for code-signing purposes.

Creating Your Own X509TrustManager

If the supplied `X509TrustManager` behavior is not suitable for your situation, then you can create your own `X509TrustManager` by either creating and registering your own `TrustManagerFactory` or by implementing the `X509TrustManager` interface directly.

Example 8-12 illustrates a `MyX509TrustManager` class that enhances the default SunJSSE `X509TrustManager` behavior by providing alternative authentication logic when the default `X509TrustManager` fails.

Once you have created such a trust manager, assign it to an `SSLContext` via the `init()` method, as in the following example. Future `SocketFactories` created from this `SSLContext` will use your new `TrustManager` when making trust decisions.

```
TrustManager[] myTMs = new TrustManager[] { new MyX509TrustManager() };
SSLContext ctx = SSLContext.getInstance("TLS");
ctx.init(null, myTMs, null);
```

Example 8-12 Sample Code for Creating a X509TrustManager

The following code sample illustrates `MyX509TrustManager` class that enhances the default SunJSSE `X509TrustManager` behavior by providing alternative authentication logic when the default `X509TrustManager` fails:

```
class MyX509TrustManager implements X509TrustManager {

     /*
      * The default PKIX X509TrustManager9.  Decisions are delegated
      * to it, and a fall back to the logic in this class is performed
      * if the default X509TrustManager does not trust it.
      */
     X509TrustManager pkixTrustManager;

     MyX509TrustManager() throws Exception {
         // create a "default" JSSE X509TrustManager.

         KeyStore ks = KeyStore.getInstance("JKS");
         ks.load(new FileInputStream("trustedCerts"),
"passphrase".toCharArray());

         TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
         tmf.init(ks);

         TrustManager tms [] = tmf.getTrustManagers();

         /*
          * Iterate over the returned trust managers, looking
          * for an instance of X509TrustManager.  If found,
          * use that as the default trust manager.
          */
         for (int i = 0; i < tms.length; i++) {
             if (tms[i] instanceof X509TrustManager) {
                 pkixTrustManager = (X509TrustManager) tms[i];
                 return;
             }
         }

         /*
          * Find some other way to initialize, or else the
          * constructor fails.
          */
         throw new Exception("Couldn't initialize");
     }

     /*
      * Delegate to the default trust manager.
      */
     public void checkClientTrusted(X509Certificate[] chain, String
authType)
                 throws CertificateException {
         try {
             pkixTrustManager.checkClientTrusted(chain, authType);
```

```
        } catch (CertificateException excep) {
            // do any special handling here, or rethrow exception.
        }
    }

    /*
     * Delegate to the default trust manager.
     */
    public void checkServerTrusted(X509Certificate[] chain, String
authType)
                throws CertificateException {
        try {
            pkixTrustManager.checkServerTrusted(chain, authType);
        } catch (CertificateException excep) {
            /*
             * Possibly pop up a dialog box asking whether to trust the
             * cert chain.
             */
        }
    }

    /*
     * Merely pass this through.
     */
    public X509Certificate[] getAcceptedIssuers() {
        return pkixTrustManager.getAcceptedIssuers();
    }
}
```

Updating the Keystore Dynamically

You can enhance `MyX509TrustManager` to handle dynamic keystore updates. When a `checkClientTrusted` or `checkServerTrusted` test fails and does not establish a trusted certificate chain, you can add the required trusted certificate to the keystore. You must create a new `pkixTrustManager` from the `TrustManagerFactory` initialized with the updated keystore. When you establish a new connection (using the previously initialized `SSLContext`), the newly added certificate will be used when making trust decisions.

X509ExtendedTrustManager Class

The `X509ExtendedTrustManager` class is an abstract implementation of the `X509TrustManager` interface. It adds methods for connection-sensitive trust management. In addition, it enables endpoint verification at the TLS layer.

In TLS 1.2 and later, both client and server can specify which hash and signature algorithms they will accept. To authenticate the remote side, authentication decisions must be based on both X509 certificates and the local accepted hash and signature algorithms. The local accepted hash and signature algorithms can be obtained using the `ExtendedSSLSession.getLocalSupportedSignatureAlgorithms()` method.

The `ExtendedSSLSession` object can be retrieved by calling the `SSLSocket.getHandshakeSession()` method or the `SSLEngine.getHandshakeSession()` method.

The `X509TrustManager` interface is not connection-sensitive. It provides no way to access `SSLSocket` or `SSLEngine` session properties.

Besides TLS 1.2 and later support, the `X509ExtendedTrustManager` class also supports algorithm constraints and SSL layer host name verification. For JSSE providers and trust manager implementations, the `X509ExtendedTrustManager` class is highly recommended over the legacy `X509TrustManager` interface.

Creating an X509ExtendedTrustManager

You can either create an `X509ExtendedTrustManager` subclass yourself (which is outlined in the following section) or obtain one from a provider-based `TrustManagerFactory` (such as that supplied by the SunJSSE provider). In Java SE 7, the PKIX or SunX509 `TrustManagerFactory` returns an `X509ExtendedTrustManager` instance.

Creating Your Own X509ExtendedTrustManager

This section outlines how to create a subclass of `X509ExtendedTrustManager` in nearly the same way as described for `X509TrustManager`.

Example 8-13 illustrates how to create a class that uses the PKIX `TrustManagerFactory` to locate a default `X509ExtendedTrustManager` that will be used to make decisions about trust.

Example 8-13 Sample Code for Creating a PKIX TrustManagerFactory

The following code sample illustrates how to create a class that uses the PKIX `TrustManagerFactory` to locate a default `X509ExtendedTrustManager` that will be used to make decisions about trust. If the default trust manager fails for any reason, then the subclass can add other behavior. In the sample, these locations are indicated by comments in the `catch` clauses.

```
import java.io.*;
import java.net.*;
import java.security.*;
import java.security.cert.*;
import javax.net.ssl.*;

public class MyX509ExtendedTrustManager extends X509ExtendedTrustManager {

  /*
   * The default PKIX X509ExtendedTrustManager.  Decisions are
   * delegated to it, and a fall back to the logic in this class is
   * performed if the default X509ExtendedTrustManager does not
   * trust it.
   */

  X509ExtendedTrustManager pkixTrustManager;

  MyX509ExtendedTrustManager() throws Exception {
    // create a "default" JSSE X509ExtendedTrustManager.

    KeyStore ks = KeyStore.getInstance("JKS");
    ks.load(new FileInputStream("trustedCerts"),
```

```
"passphrase".toCharArray());

    TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
    tmf.init(ks);

    TrustManager tms [] = tmf.getTrustManagers();

    /*
     * Iterate over the returned trust managers, looking
     * for an instance of X509ExtendedTrustManager. If found,
     * use that as the default trust manager.
     */
    for (int i = 0; i < tms.length; i++) {
      if (tms[i] instanceof X509ExtendedTrustManager) {
        pkixTrustManager = (X509ExtendedTrustManager) tms[i];
        return;
      }
    }

    /*
     * Find some other way to initialize, or else we have to fail the
     * constructor.
     */
    throw new Exception("Couldn't initialize");
  }

  /*
   * Delegate to the default trust manager.
   */
  public void checkClientTrusted(X509Certificate[] chain, String authType)
    throws CertificateException {
    try {
      pkixTrustManager.checkClientTrusted(chain, authType);
    } catch (CertificateException excep) {
      // do any special handling here, or rethrow exception.
    }
  }

  /*
   * Delegate to the default trust manager.
   */
  public void checkServerTrusted(X509Certificate[] chain, String authType)
    throws CertificateException {
    try {
      pkixTrustManager.checkServerTrusted(chain, authType);
    } catch (CertificateException excep) {
      /*
       * Possibly pop up a dialog box asking whether to trust the
       * cert chain.
       */
    }
  }

  /*
   * Connection-sensitive verification.
```

```
   */
  public void checkClientTrusted(X509Certificate[] chain, String authType,
Socket socket)
    throws CertificateException {
    try {
      pkixTrustManager.checkClientTrusted(chain, authType, socket);
    } catch (CertificateException excep) {
      // do any special handling here, or rethrow exception.
    }
  }

  public void checkClientTrusted(X509Certificate[] chain, String authType,
SSLEngine engine)
    throws CertificateException {
    try {
      pkixTrustManager.checkClientTrusted(chain, authType, engine);
    } catch (CertificateException excep) {
      // do any special handling here, or rethrow exception.
    }
  }

  public void checkServerTrusted(X509Certificate[] chain, String authType,
Socket socket)
    throws CertificateException {
    try {
      pkixTrustManager.checkServerTrusted(chain, authType, socket);
    } catch (CertificateException excep) {
      // do any special handling here, or rethrow exception.
    }
  }

  public void checkServerTrusted(X509Certificate[] chain, String authType,
SSLEngine engine)
    throws CertificateException {
    try {
      pkixTrustManager.checkServerTrusted(chain, authType, engine);
    } catch (CertificateException excep) {
      // do any special handling here, or rethrow exception.
    }
  }

  /*
   * Merely pass this through.
   */
  public X509Certificate[] getAcceptedIssuers() {
    return pkixTrustManager.getAcceptedIssuers();
  }
}
```

KeyManager Interface

The primary responsibility of the `KeyManager` is to select the authentication credentials that will eventually be sent to the remote host. To authenticate yourself (a local secure socket peer) to a remote peer, you must initialize an `SSLContext` object with one or

more KeyManager objects. You must pass one KeyManager for each different authentication mechanism that will be supported. If null is passed into the SSLContext initialization, then an empty KeyManager will be created. If the internal default context is used (for example, an SSLContext created by SSLSocketFactory.getDefault() or SSLServerSocketFactory.getDefault()), then a default KeyManager is created. See Customizing the Default Keystores and Truststores, Store Types, and Store Passwords. Typically, a single key manager supports authentication based on X.509 public key certificates. Some secure socket implementations may also support authentication based on shared secret keys, Kerberos, or other mechanisms.

KeyManager objects are created either by a KeyManagerFactory, or by providing a concrete implementation of the interface.

KeyManagerFactory Class

The javax.net.ssl.KeyManagerFactory class is an engine class for a provider-based service that acts as a factory for one or more types of KeyManager objects. The SunJSSE provider implements a factory that can return a basic X.509 key manager. Because it is provider-based, additional factories can be implemented and configured to provide additional or alternative key managers.

Creating a KeyManagerFactory

You create an instance of this class in a similar manner to SSLContext, except for passing an algorithm name string instead of a protocol name to the getInstance() method:

```
KeyManagerFactory kmf = getInstance(String algorithm);
KeyManagerFactory kmf = getInstance(String algorithm, String provider);
KeyManagerFactory kmf = getInstance(String algorithm, Provider provider);
```

A sample call as follows:

```
KeyManagerFactory kmf = KeyManagerFactory.getInstance("SunX509",
"SunJSSE");
```

The preceding call creates an instance of the SunJSSE provider's default key manager factory, which provides basic X.509-based authentication keys.

A newly created factory should be initialized by calling one of the init methods:

```
public void init(KeyStore ks, char[] password);
public void init(ManagerFactoryParameters spec);
```

Call whichever init method is appropriate for the KeyManagerFactory you are using. If you are not sure, then ask the provider vendor.

For many factories, such as the default SunX509 KeyManagerFactory from the SunJSSE provider, the KeyStore and password are the only information required to initialize the KeyManagerFactory and thus the first init method is the appropriate one to call. The KeyManagerFactory will query the KeyStore for information about which private key and matching public key certificates should be used for authenticating to a remote socket peer. The password parameter specifies the password that will be used

with the methods for accessing keys from the `KeyStore`. All keys in the `KeyStore` must be protected by the same password.

Sometimes initialization parameters other than a `KeyStore` and password are needed by a provider. Users of that provider are expected to pass an implementation of the appropriate `ManagerFactoryParameters` as defined by the provider. The provider can then call the specified methods in the `ManagerFactoryParameters` implementation to obtain the needed information.

Some factories can provide access to authentication material without being initialized with a `KeyStore` object or any other parameters. For example, they may access key material as part of a login mechanism such as one based on JAAS, the Java Authentication and Authorization Service.

As previously indicated, the SunJSSE provider supports a SunX509 factory that must be initialized with a `KeyStore` parameter.

X509KeyManager Interface

The `javax.net.ssl.X509KeyManager` interface extends the general `KeyManager` interface. It must be implemented by a key manager for X.509-based authentication. To support X.509 authentication to remote socket peers through JSSE, an instance of this interface must be passed to the `init()` method of an `SSLContext` object.

Creating an X509KeyManager

You can either implement this interface directly yourself or obtain one from a provider-based `KeyManagerFactory` (such as that supplied by the SunJSSE provider). You could also implement your own interface that delegates to a factory-generated key manager. For example, you might do this to filter the resulting keys and query an end-user through a graphical user interface.

Creating Your Own X509KeyManager

If the default `X509KeyManager` behavior is not suitable for your situation, then you can create your own `X509KeyManager` in a way similar to that shown in Creating Your Own X509TrustManager.

X509ExtendedKeyManager Class

The `X509ExtendedKeyManager` abstract class is an implementation of the `X509KeyManager` interface that allows for connection-specific key selection. It adds two methods that select a key alias for client or server based on the key type, allowed issuers, and current `SSLEngine`:

- `public String chooseEngineClientAlias(String[] keyType, Principal[] issuers, SSLEngine engine)`
- `public String chooseEngineServerAlias(String keyType, Principal[] issuers, SSLEngine engine)`

If a key manager is not an instance of the `X509ExtendedKeyManager` class, then it will not work with the `SSLEngine` class.

For JSSE providers and key manager implementations, the `X509ExtendedKeyManager` class is highly recommended over the legacy `X509KeyManager` interface.

In TLS 1.2 and later, both client and server can specify which hash and signature algorithms they will accept. To pass the authentication required by the remote side, local key selection decisions must be based on both X509 certificates and the remote accepted hash and signature algorithms. The remote accepted hash and signature algorithms can be retrieved using the `ExtendedSSLSession.getPeerSupportedSignatureAlgorithms()` method.

You can create your own `X509ExtendedKeyManager` subclass in a way similar to that shown in Creating Your Own X509TrustManager.

Support for the Server Name Indication (SNI) Extension on the server side enables the key manager to check the server name and select the appropriate key accordingly. For example, suppose there are three key entries with certificates in the keystore:

- `cn=www.example.com`
- `cn=www.example.org`
- `cn=www.example.net`

If the ClientHello message requests to connect to `www.example.net` in the SNI extension, then the server should be able to select the certificate with subject `cn=www.example.net`.

Relationship Between a TrustManager and a KeyManager

Historically, there has been confusion regarding the functionality of a `TrustManager` and a `KeyManager`.

A `TrustManager` determines whether the remote authentication credentials (and thus the connection) should be trusted.

A `KeyManager` determines which authentication credentials to send to the remote host.

Secondary Support Classes and Interfaces

These classes are provided as part of the JSSE API to support the creation, use, and management of secure sockets. They are less likely to be used by secure socket applications than are the core and support classes. The secondary support classes and interfaces are part of the `javax.net.ssl` and `javax.security.cert` packages.

SSLParameters Class

The `SSLParameters` class encapsulates the following parameters that affect a SSL/TLS/DTLS connection:

- The list of cipher suites to be accepted in a TLS/DTLS handshake
- The list of protocols to be allowed
- The endpoint identification algorithm during TLS/DTLS handshaking
- The server names and server name matchers (see Server Name Indication (SNI) Extension)
- The cipher suite preference to be used in a TLS/DTLS handshake

- Algorithm during TLS/DTLS handshaking
- The Server Name Indication (SNI)
- The maximum network packet size
- The algorithm constraints and whether TLS/DTLS servers should request or require client authentication

You can retrieve the current `SSLParameters` for an `SSLSocket` or `SSLEngine` by using the following methods:

- `getSSLParameters()` in an `SSLSocket`, `SSLServerSocket`, and `SSLEngine`
- `getDefaultSSLParameters()` and `getSupportedSSLParamters()` in an `SSLContext`

You can assign `SSLParameters` with the `setSSLParameters()` method in an `SSLSocket`, `SSLServerSocket` and `SSLEngine`.

You can explicitly set the server name indication with the `SSLParameters.setServerNames()` method. The server name indication in client mode also affects endpoint identification. In the implementation of `X509ExtendedTrustManager`, it uses the server name indication retrieved by the `ExtendedSSLSession.getRequestedServerNames()` method. See Example 8-14.

Example 8-14 Sample Code to Set Server Name Indication

This example uses the host name in the server name indication (`www.example.com`) to make endpoint identification against the peer's identity presented in the end-entity's X.509 certificate.

```
    SSLSocketFactory factory = ...
    SSLSocket sslSocket = factory.createSocket("172.16.10.6", 443);
    // SSLEngine sslEngine = sslContext.createSSLEngine("172.16.10.6",
443);

    SNIHostName serverName = new SNIHostName("www.example.com");
    List<SNIServerName> serverNames = new ArrayList<>(1);
    serverNames.add(serverName);

    SSLParameters params = sslSocket.getSSLParameters();
    params.setServerNames(serverNames);
    sslSocket.setSSLParameters(params);
    // sslEngine.setSSLParameters(params);
```

Cipher Suite Preference

During TLS handshaking, the client requests to negotiate a cipher suite from a list of cryptographic options that it supports, starting with its first preference. Then, the server selects a single cipher suite from the list of cipher suites requested by the client. Normally, the selection honors the client's preference. However, to mitigate the risks of using weak cipher suites, the server may select cipher suites based on its own preference rather than the client's preference, by invoking the method `SSLParameters.setUseCipherSuitesOrder(true)`.

SSLSessionContext Interface

The `javax.net.ssl.SSLSessionContext` interface is a grouping of SSLSession objects associated with a single entity. For example, it could be associated with a server or client that participates in many sessions concurrently. The methods in this interface enable the enumeration of all sessions in a context and allow lookup of specific sessions via their session IDs.

An `SSLSessionContext` may optionally be obtained from an `SSLSession` by calling the SSLSession `getSessionContext()` method. The context may be unavailable in some environments, in which case the `getSessionContext()` method returns null.

SSLSessionBindingListener Interface

The `javax.net.ssl.SSLSessionBindingListener` interface is implemented by objects that are notified when they are being bound or unbound from an SSLSession.

SSLSessionBindingEvent Class

The `javax.net.ssl.SSLSessionBindingEvent` class defines the event communicated to an `SSLSessionBindingListener` (see SSLSessionBindingListener Interface) when it is bound or unbound from an `SSLSession` (see SSLSession and ExtendedSSLSession).

HandShakeCompletedListener Interface

The `javax.net.ssl.HandShakeCompletedListener` interface is an interface implemented by any class that is notified of the completion of an SSL protocol handshake on a given `SSLSocket` connection.

HandShakeCompletedEvent Class

The `javax.net.ssl.HandShakeCompletedEvent` class defines the event communicated to a `HandShakeCompletedListener` (see HandShakeCompletedListener Interface) upon completion of an SSL protocol handshake on a given `SSLSocket` connection.

HostnameVerifier Interface

If the SSL/TLS implementation's standard host name verification logic fails, then the implementation calls the `verify()` method of the class that implements this interface and is assigned to this `HttpsURLConnection` instance. If the callback class can determine that the host name is acceptable given the parameters, it reports that the connection should be allowed. An unacceptable response causes the connection to be terminated. See Example 8-15.

See `HttpsURLConnection` for more information about how to assign the `HostnameVerifier` to the `HttpsURLConnection`.

Example 8-15 Sample Code for Implementing the HostnameVerifier Interface

The following example illustrates a class that implements `HostnameVerifier` interface:

```
public class MyHostnameVerifier implements HostnameVerifier {
```

```
    public boolean verify(String hostname, SSLSession session) {
        // pop up an interactive dialog box
        // or insert additional matching logic
        if (good_address) {
            return true;
        } else {
            return false;
        }
    }
}

//...deleted...

HttpsURLConnection urlc = (HttpsURLConnection)
  (new URL("https://www.example.com/")).openConnection();
urlc.setHostnameVerifier(new MyHostnameVerifier());
```

X509Certificate Class

Many secure socket protocols perform authentication using public key certificates, also called X.509 certificates. This is the default authentication mechanism for the TLS protocol.

The `java.security.cert.X509Certificate` abstract class provides a standard way to access the attributes of X.509 certificates.

> **Note:**
>
> The `javax.security.cert.X509Certificate` class is supported only for backward compatibility with previous (1.0.x and 1.1.x) versions of JSSE. New applications should use the `java.security.cert.X509Certificate` class instead.

AlgorithmConstraints Interface

The `java.security.AlgorithmConstraints` interface is used for controlling allowed cryptographic algorithms. `AlgorithmConstraints` defines three `permits()` methods. These methods tell whether an algorithm name or a key is permitted for certain cryptographic functions. Cryptographic functions are represented by a set of `CryptoPrimitive`, which is an enumeration containing fields like `STREAM_CIPHER`, `MESSAGE_DIGEST`, and `SIGNATURE`.

Thus, an `AlgorithmConstraints` implementation can answer questions like: Can I use this key with this algorithm for the purpose of a cryptographic operation?

An `AlgorithmConstraints` object can be associated with an `SSLParameters` object by using the new `setAlgorithmConstraints()` method. The current `AlgorithmConstraints` object for an `SSLParameters` object is retrieved using the `getAlgorithmConstraints()` method.

StandardConstants Class

The `StandardConstants` class is used to represent standard constants definitions in JSSE.

`StandardConstants.SNI_HOST_NAME` represents a domain name server (DNS) host name in a Server Name Indication (SNI) extension, which can be used when instantiating an `SNIServerName` or `SNIMatcher` object.

SNIServerName Class

An instance of the abstract `SNIServerName` class represents a server name in the Server Name Indication (SNI) extension. It is instantiated using the type and encoded value of the specified server name.

You can use the `getType()` and `getEncoded()` methods to return the server name type and a copy of the encoded server name value, respectively. The `equals()` method can be used to check if some other object is "equal" to this server name. The `hashCode()` method returns a hash code value for this server name. To get a string representation of the server name (including the server name type and encoded server name value), use the `toString()` method.

SNIMatcher Class

An instance of the abstract `SNIMatcher` class performs match operations on an `SNIServerName` object. Servers can use information from the Server Name Indication (SNI) extension to decide if a specific `SSLSocket` or `SSLEngine` should accept a connection. For example, when multiple "virtual" or "name-based" servers are hosted on a single underlying network address, the server application can use SNI information to determine whether this server is the exact server that the client wants to access. Instances of this class can be used by a server to verify the acceptable server names of a particular type, such as host names.

The `SNIMatcher` class is instantiated using the specified server name type on which match operations will be performed. To match a given `SNIServerName`, use the `matches()` method. To return the server name type of the given `SNIMatcher` object, use the `getType()` method.

SNIHostName Class

An instance of the `SNIHostName` class (which extends the `SNIServerName` class) represents a server name of type "host_name" (see StandardConstants Class) in the Server Name Indication (SNI) Extension. To instantiate an `SNIHostName`, specify the fully qualified DNS host name of the server (as understood by the client) as a `String` argument. The argument is illegal in the following cases:

- The argument is empty.
- The argument ends with a trailing period.
- The argument is not a valid Internationalized Domain Name (IDN) compliant with the RFC 3490 specification.

You can also instantiate an `SNIHostName` by specifying the encoded host name value as a byte array. This method is typically used to parse the encoded name value in a requested SNI extension. Otherwise, use the `SNIHostName(String hostname)` constructor. The `encoded` argument is illegal in the following cases:

- The argument is empty.
- The argument ends with a trailing period.
- The argument is not a valid Internationalized Domain Name (IDN) compliant with the RFC 3490 specification.
- The argument is not encoded in UTF-8 or US-ASCII.

> **Note:**
>
> The `encoded` byte array passed in as an argument is cloned to protect against subsequent modification.

To return the host name of an `SNIHostName` object in US-ASCII encoding, use the `getAsciiName()` method. To compare a server name to another object, use the `equals()` method (comparison is *not* case-sensitive). To return a hash code value of an `SNIHostName`, use the `hashCode()` method. To return a string representation of an `SNIHostName`, including the DNS host name, use the `toString()` method.

You can create an `SNIMatcher` object for an `SNIHostName` object by passing a regular expression representing one or more host names to match to the `createSNIMatcher()` method.

Customizing JSSE

JSSE includes a standard implementation that can be customized by plugging in different implementations or specifying the default keystore, and so on.

Table 8-2 and Table 8-3 summarize which aspects can be customized, what the defaults are, and which mechanisms are used to provide customization.

Some of the customizations are done by setting system property or Security Property values. Sections following the table explain how to set such property values.

> **Note:**
>
> Many of the properties shown in this table are currently used by the JSSE implementation, but there is no guarantee that they will continue to have the same names and types (system or security) or even that they will exist at all in future releases. All such properties are flagged with an asterisk (*). They are documented here for your convenience for use with the JSSE implementation.

Table 8-2 shows items that are customized by setting the `java.security.Security` property. See How to Specify a java.security.Security Property

Table 8-2 Security Properties and Customized Items

Security Property	Customized Item	Default Value	Notes
`cert.provider.x509v1`	Customizing the X509Certificate Implementation	X509Certificate implementation from Oracle	None
`security.provider.`*n*	Cryptographic service provider; see Customizing the Provider Implementation and Customizing the Encryption Algorithm Providers	The first five providers in order of priority are: 1. SUN 2. SunRsaSign 3. SunEC 4. SunJSSE 5. SunJCE	Specify the provider in the `security.provider.`*n*= line in security properties file, where *n* is an integer whose value is equal or greater than 1.
`*ssl.SocketFactory.provider`	Default `SSLSocketFactory` implementation	`SSLSocketFactory` implementation from Oracle	None
`*ssl.ServerSocketFactory.provider`	Default `SSLServerSocketFactory` implementation	`SSLServerSocketFactory` implementation from Oracle	None
`ssl.KeyManagerFactory.algorithm`	Default key manager factory algorithm name (see Customizing the Default Key Managers and Trust Managers)	SunX509	None
`*jdk.certpath.disabledAlgorithms`	Disabled certificate verification cryptographic algorithm (see Disabled and Restricted Cryptographic Algorithms)	MD2, MD5, SHA1 jdkCA & usage TLSServer, RSA keySize < 1024, DSA keySize < 1024, EC keySize < 224[1]	None
`ssl.TrustManagerFactory.algorithm`	Default trust manager factory algorithm name (see Customizing the Default Key Managers and Trust Managers)	PKIX	None
JCE encryption algorithms used by the SunJSSE provider	Give alternative JCE algorithm providers a higher preference order than the SunJCE provider	SunJCE implementations	None
`*jdk.tls.disabledAlgorithms`	Disabled and Restricted Cryptographic Algorithms	SSLv3, RC4, MD5withRSA, DH keySize < 1024, EC keySize < 224[1]	Disables specific algorithms (protocols versions, cipher suites, key exchange mechanisms, etc.) that will not be negotiated for TLS/DTLS connections, even if they are enabled explicitly in an application

Table 8-2 (Cont.) Security Properties and Customized Items

Security Property	Customized Item	Default Value	Notes
`*jdk.tls.legacyAlgorithms`	Legacy Cryptographic Algorithms	K_NULL, C_NULL, M_NULL, DH_anon, ECDH_anon, RC4_128, RC4_40, DES_CBC, DES40_CBC, 3DES_EDE_CBC[1]	Specifies which algorithms are considered legacy algorithms, which are not negotiated during TLS/DTLS security parameters negotiation unless there are no other candidates.
`jdk.tls.server.defaultDHEParameters`	Diffie-Hellman groups	Safe prime Diffie-Hellman groups in OpenJDK TLS/DTLS implementation	Defines default finite field Diffie-Hellman ephemeral (DHE) parameters for (Datagram) Transport Layer Security ((D)TLS) processing
`*jdk.tls.keyLimits`	Limiting Amount of Data Algorithms May Encrypt with a Set of Keys	AES/GCM/NoPadding KeyUpdate 2^37	Limits the amount of data an algorithm may encrypt with a specific set of keys; once this limit is reached, a KeyUpdate post-handshake message is sent, which requests that the current set of keys be updated.
`*ocsp.enable`	Client-Driven OCSP and OCSP Stapling	false	Enables client-driven Online Certificate Status Protocol (OCSP). You must also enable revocation checking; see Setting up a Java Client to use Client-Driven OCSP.

[1] The list of restricted, disabled, and legacy algorithms specified in these Security Properties may change; see the `java.security` file in your JDK installation for the latest values.

* This Security Property is currently used by the JSSE implementation, but it is not guaranteed to be examined and used by other implementations. If it *is* examined by another implementation, then that implementation should handle it in the same manner as the JSSE implementation does. There is no guarantee the property will continue to exist or be of the same type (system or security) in future releases.

Table 8-3 shows items that are customized by setting `java.lang.System` property. See How to Specify a java.lang.System Property.

Table 8-3 System Properties and Customized Items

System Property	Customized Item	Default	Notes
`java.protocol.handler.pkgs`	Specifying an Alternative HTTPS Protocol Implementation	Implementation from Oracle	None

Table 8-3 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
`*javax.net.ssl.keyStore`	Default keystore; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	None	The value `NONE` may be specified. This setting is appropriate if the keystore is not file-based (for example, it resides in a hardware token)
`*javax.net.ssl.keyStorePassword`	Default keystore password; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	None	It is inadvisable to specify the password in a way that exposes it to discovery by other users. For example, specifying the password on the command line. To keep the password secure, have the application prompt for the password, or specify the password in a properly protected option file
`*javax.net.ssl.keyStoreProvider`	Default keystore provider; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	None	None
`*javax.net.ssl.keyStoreType`	Default keystore type; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	`KeyStore.getDefaultType()`	None
`*javax.net.ssl.trustStore`	Default truststore; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	`jssecacerts`, if it exists. Otherwise, `cacerts`	None
`*javax.net.ssl.trustStorePassword`	Default truststore password; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	None	It is inadvisable to specify the password in a way that exposes it to discovery by other users. For example, specifying the password on the command line. To keep the password secure, have the application prompt for the password, or specify the password in a properly protected option file
`*javax.net.ssl.trustStoreProvider`	Default truststore provider; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	None	None

Table 8-3 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
`*javax.net.ssl.trustStoreType`	Default truststore type; see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	`KeyStore.getDefaultType()`	The value `NONE` may be specified. This setting is appropriate if the truststore is not file-based (for example, it resides in a hardware token)
`*https.proxyHost`	Default proxy host	None	None
`*https.proxyPort`	Default proxy port	80	None
`*jsse.enableSNIExtension`	Server Name Indication option	`true`	Server Name Indication (SNI) is a TLS extension, defined in RFC 6066. It enables TLS connections to virtual servers, in which multiple servers for different network names are hosted at a single underlying network address. Some very old TLS vendors may not be able handle TLS extensions. In this case, set this property to `false` to disable the SNI extension
`*https.cipherSuites`	Default cipher suites for HTTPS connections	Determined by the socket factory.	This contains a comma-separated list of cipher suite names specifying which cipher suites to enable for use on this `HttpsURLConnection`. See the `SSLSocket.setEnabledCipherSuites(String[])` method. Note that this method sets the preference order of the ClientHello cipher suites directly from the `String` array passed to it.
`*https.protocols`	Default handshaking protocols for HTTPS connections	Determined by the socket factory.	This contains a comma-separated list of protocol suite names specifying which protocol suites to enable on this `HttpsURLConnection`. See `SSLSocket.setEnabledProtocols(String[])`
* Customize via `port` field in the HTTPS URL.	Default HTTPS port	443	None

Table 8-3 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
`*jsse.SSLEngine.acceptLargeFragments`	Default sizing buffers for large TLS packets	None	Setting this system property to `true`, `SSLSession` will size buffers to handle large data packets by default. This may cause applications to allocate unnecessarily large `SSLEngine` buffers. Instead, applications should dynamically check for buffer overflow conditions and resize buffers as appropriate
`jdk.tls.rejectClientInitiatedRenegotiation`	Rejects client-initiated renegotiation on the server side. If this system property is `true`, then the server will not accept client initiated renegotiations and will fail with a fatal `handshake_failure` alert. Rejects server-side client-initialized renegotiation.	`false`	None
`jdk.tls.client.cipherSuites`	Client-side default enabled cipher suites; see Specifying Default Enabled Cipher Suites	See SunJSSE Cipher Suites for a list of currently implemented SunJSSE cipher suites for this JDK release, sorted by order of preference	**Caution**: These system properties can be used to configure weak cipher suites, or the configured cipher suites may be weak in the future. It is not recommended that you use these system properties without understanding the risks.

Table 8-3 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
`jdk.tls.client.protocols`	Default handshaking protocols for TLS/DTLS clients. See The SunJSSE Provider	None	To enable specific SunJSSE protocols on the client, specify them in a comma-separated list within quotation marks; all other supported protocols are not enabled on the client For example, • If `jdk.tls.client.protocols="TLSv1,TLSv1.1"`, then the default protocol settings on the client for TLSv1 and TLSv1.1 are enabled, while SSLv3, TLSv1.2, TLSv1.3, and SSLv2Hello are not enabled • If `jdk.tls.client.protocols="DTLSv1.2"` , then the protocol setting on the client for DTLS1.2 is enabled, while DTLS1.0 is not enabled
`jdk.tls.server.cipherSuites`	Server-side default enabled cipher suites. See Specifying Default Enabled Cipher Suites	See SunJSSE Cipher Suites to determine which cipher suites are enabled by default	**Caution**: These system properties can be used to configure weak cipher suites, or the configured cipher suites may be weak in the future. It is not recommended that you use these system properties without understanding the risks.

Table 8-3 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
`jdk.tls.server.protocols`	Default handshaking protocols for TLS/DTLS servers. See The SunJSSE Provider	None	To configure the default enabled protocol suite in the server side of a SunJSSE provider, specify the protocols in a comma-separated list within quotation marks. The protocols in this list are standard SSL protocol names as described in Java Security Standard Algorithm Names. Note that this System Property impacts only the default protocol suite (SSLContext of the algorithms SSL, TLS and DTLS). If an application uses a version-specific SSLContext (SSLv3, TLSv1, TLSv1.1, TLSv1.2, TLSv1.3, DTLSv1.0, or DTLSv1.2), or sets the enabled protocol version explicitly, this System Property has no impact.
`jdk.tls.ephemeralDHKeySize`	Customizing Size of Ephemeral Diffie-Hellman Keys	1024 bits	None
`jdk.tls.namedGroups`	Customizing the supported named groups for TLS/DTLS key exchange	If this System Property is not defined or the value is empty, then the implementation default groups and preferences will be used.	This contains a comma-separated list within quotation marks of enabled named groups in preference order. For example: `jdk.tls.namedGroups="secp521r1,secp256r1,ffdhe2048"`
`jsse.enableMFLNExtension`	Customizing Maximum Fragment Length Negotiation (MFLN) Extension	false	None

Table 8-3 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
`jsse.enableFFDHEExtension`	Enables or disables Finite Field Diffie-Hellman Ephemeral (FFDHE) parameters for TLS/DTLS key exchange	true	FFDHE is a TLS/DTLS extension defined in RFC 7919. It enables TLS/DTLS connections to use known finite field Diffie-Hellman groups. Some very old TLS vendors may not be able handle TLS extensions. In this case, set this property to false to disable the FFDHE extension.
`*com.sun.net.ssl.checkRevocation`	Revocation checking	false	You must enable revocation checking to enable client-driven OCSP; see Client-Driven OCSP and OCSP Stapling.

* This system property is currently used by the JSSE implementation, but it is not guaranteed to be examined and used by other implementations. If it *is* examined by another implementation, then that implementation should handle it in the same manner as the JSSE implementation does. There is no guarantee the property will continue to exist or be of the same type (system or security) in future releases.

How to Specify a java.lang.System Property

You can customize some aspects of JSSE by setting system properties. There are several ways to set these properties:

- To set a system property statically, use the `-D` option of the `java` command. For example, to run an application named MyApp and set the `javax.net.ssl.trustStore` system property to specify a truststore named MyCacertsFile. See truststore. Enter the following:

```
java -Djavax.net.ssl.trustStore=MyCacertsFile MyApp
```

- To set a system property dynamically, call the `java.lang.System.setProperty()` method in your code:

```
System.setProperty("propertyName", "propertyValue");
```

For example, a `setProperty()` call corresponding to the previous example for setting the `javax.net.ssl.trustStore` system property to specify a truststore named "`MyCacertsFile`" would be:

```
System.setProperty("javax.net.ssl.trustStore", "MyCacertsFile");
```

How to Specify a java.security.Security Property

You can customize some aspects of JSSE by setting Security Properties. You can set a Security Property either statically or dynamically:

- To set a Security Property statically, add a line to the security properties file. The security properties file is located at *`java-home`*`/conf/security/java.security`

  java-home
  See Terms and Definitions

  To specify a Security Property value in the security properties file, you add a line of the following form:

  ```
  propertyName=propertyValue
  ```

  For example, suppose that you want to specify a different key manager factory algorithm name than the default SunX509. You do this by specifying the algorithm name as the value of a Security Property named `ssl.KeyManagerFactory.algorithm`. For example, to set the value to MyX509, add the following line to the security properties file:

  ```
  ssl.KeyManagerFactory.algorithm=MyX509
  ```

- To set a Security Property dynamically, call the `java.security.Security.setProperty` method in your code:

  ```
  Security.setProperty("propertyName," "propertyValue");
  ```

  For example, a call to the `setProperty()` method corresponding to the previous example for specifying the key manager factory algorithm name would be:

  ```
  Security.setProperty("ssl.KeyManagerFactory.algorithm", "MyX509");
  ```

Customizing the X509Certificate Implementation

The X509Certificate implementation returned by the `X509Certificate.getInstance()` method is by default the implementation from the JSSE implementation.

To cause a different implementation to be returned:

Specify the name (and package) of the other implementation's class as the value of a How to Specify a java.security.Security Property named `cert.provider.x509v1`.

```
MyX509CertificateImplcom.cryptox


    cert.provider.x509v1=com.cryptox.MyX509CertificateImpl
```

Specifying Default Enabled Cipher Suites

You can specify the default enabled cipher suites in your application or with the system properties `jdk.tls.client.cipherSuites` and `jdk.tls.server.cipherSuites`.

> **Note:**
>
> The actual use of enabled cipher suites is restricted by algorithm constraints.

The set of cipher suites to enable by default is determined by one of the following ways in this order of preference:

1. Explicitly set by application
2. Specified by system property
3. Specified by JSSE provider defaults

For example, explicitly setting the default enabled cipher suites in your application overrides settings specified in `jdk.tls.client.cipherSuites` or `jdk.tls.server.cipherSuites` as well as JSSE provider defaults.

Explicitly Set by Application

You can set which cipher suites are enabled with one of the following methods:

- `SSLSocket.setEnabledCipherSuites(String[])`
- `SSLEngine.setEnabledCipherSuites(String[])`
- `SSLServerSocket.setEnabledCipherSuites(String[])`
- `SSLParameters(String[] cipherSuites)`
- `SSLParameters(String[] cipherSuites, String[] protocols)`
- `SSLParameters.setCipherSuites(String[])`
- `https.cipherSuites` system property for `HttpsURLConnection`

Specified by System Property

The system property `jdk.tls.client.cipherSuites` specifies the default enabled cipher suites on the client side; `jdk.tls.server.cipherSuites` specifies those on the server side.

The syntax of the value of these two system properties is a comma-separated list of supported cipher suite names. Unrecognized or unsupported cipher suite names that are specified in these properties are ignored. See Java Security Standard Algorithms for standard JSSE cipher suite names.

> **Note:**
>
> These system properties are currently supported by Oracle JDK and OpenJDK. They are not guaranteed to be supported by other JDK implementations.

> **Caution:**
>
> These system properties can be used to configure weak cipher suites, or the configured cipher suites may be weak in the future. It is not recommended that you use these system properties without understanding the risks.

Specified by JSSE Provider Defaults

Each JSSE provider has its own default enabled cipher suites. See The SunJSSE Provider in JDK Providers Documentation for the cipher suite names supported by the SunJSSE provider and which ones that are enabled by default.

Specifying an Alternative HTTPS Protocol Implementation

You can communicate securely with an SSL-enabled web server by using the HTTPS URL scheme for the `java.net.URL` class. The JDK provides a default HTTPS URL implementation.

If you want an alternative HTTPS protocol implementation to be used, set the `java.protocol.handler.pkgs` How to Specify a java.lang.System Property to include the new class name. This action causes the specified classes to be found and loaded before the JDK default classes. See the `URL` class for details.

> **Note:**
>
> In past JSSE releases, you had to set the `java.protocol.handler.pkgs` system property during JSSE installation. This step is no longer required unless you want to obtain an instance of `com.sun.net.ssl.HttpsURLConnection`.

Customizing the Provider Implementation

The JDK comes with a JSSE Cryptographic Service Provider, or *provider* for short, named SunJSSE. Providers are essentially packages that implement one or more engine classes for specific cryptographic algorithms.

The JSSE engine classes are `SSLContext`, `KeyManagerFactory`, and `TrustManagerFactory`. See Java Cryptography Architecture (JCA) Reference Guide to know more about providers and engine classes.

Before it can be used, a provider must be registered, either statically or dynamically. You do not need to register the SunJSSE provider because it is preregistered. If you want to use other providers, read the following sections to see how to register them.

Registering the Cryptographic Provider Statically

Register a provider statically by adding a line of the following form to the security properties file, `<java-home>/conf/security/java.security`:

```
security.provider.n=provName|className
```

This declares a provider, and specifies its preference order *n*. The preference order is the order in which providers are searched for requested algorithms when no specific provider is requested. The order is 1-based; 1 is the most preferred, followed by 2, and so on.

provName is the provider's name and *`className`* is the fully qualified class name of the provider.

Standard security providers are automatically registered for you in the `java.security` security properties file.

To use another JSSE provider, add a line registering the other provider, giving it whatever preference order you prefer.

You can have more than one JSSE provider registered at the same time. The registered providers may include different implementations for different algorithms for different engine classes, or they may have support for some or all of the same types of algorithms and engine classes. When a particular engine class implementation for a particular algorithm is searched for, if no specific provider is specified for the search, then the providers are searched in preference order and the implementation from the first provider that supplies an implementation for the specified algorithm is used.

See Step 8.1: Configure the Provider in Steps to Implement and Integrate a Provider.

Registering the Cryptographic Service Provider Dynamically

Instead of registering a provider statically, you can add the provider dynamically at runtime by calling either the `addProvider` or `insertProviderAt` method in the `Security` class. Note that this type of registration is not persistent and can only be done by code which is granted the `insertProvider.<provider name>` permission.

See Step 8.1: Configure the Provider in Steps to Implement and Integrate a Provider.

Provider Configuration

Some providers may require configuration. This is done using the `configure` method of the `Provider` class, prior to calling the `addProvider` method of the `Security` class. See SunPKCS11 Configuration for an example. The `Provider.configure()` method is new to Java SE 9.

Configuring the Preferred Provider for Specific Algorithms

Specify the preferred provider for a specific algorithm in the `jdk.security.provider.preferred` Security Property. By specifying a preferred provider you can configure providers that offer performance gains for specific

algorithms but are not the best performing provider for other algorithms. The ordered provider list specified using the `security.provider.n` property is not sufficient to order providers that offer performance gains for specific algorithms but are not the best performing provider for other algorithms. More flexibility is required for configuring the ordering of provider list to achieve performance gains.

The `jdk.security.provider.preferred` Security Property allows specific algorithms, or service types to be selected from a preferred set of providers before accessing the list of registered providers. See How to Specify a java.security.Security Property.

The `jdk.security.provider.preferred` Security Property does not register the providers. The ordered provider list must be Registering the Cryptographic Provider Statically using the `security.provider.n` property. Any provider that is not registered is ignored.

Specifying the Preferred Provider for an Algorithm

The syntax for specifying the preferred providers string in the `jdk.security.provider.preferred` Security Property is a comma-separated list of `ServiceType.Algorithm:Provider`

In this syntax:

ServiceType
The name of the service type (for example: `"MessageDigest"`). ServiceType is optional. If it isn't specified, then the algorithm applies to all service types.

Algorithm
The standard algorithm name. See Java Security Standard Algorithm Names. Algorithms can be specified as full standard name, (AES/CBC/PKCS5Padding) or as partial (AES, AES/CBC, AES//PKCS5Padding).

Provider
The name of the provider. Any provider that isn't listed in the registered list is ignored. See JDK Providers.

Entries containing errors such as parsing errors are ignored. Use the command `java -Djava.security.debug=jca` to debug errors.

Preferred Providers and FIPS

If you add a FIPS provider to the `security.provider.n` property, and specify the preferred provider ordering in the `jdk.security.provider.preferred` property then the preferred providers specified in `jdk.security.provider.preferred` are selected first.

Hence, it is recommended that you don't configure `jdk.security.provider.preferred` property for FIPS provider configurations.

jdk.security.provider.preferred Default Values

The `jdk.security.provider.preferred` property is not set by default and is used only for application performance tuning.

Example 8-16 Sample jdk.security.provider.preferred Property

The syntax for specifying the `jdk.security.provider.preferred` property is as follows:

```
jdk.security.provider.preferred=AES/GCM/NoPadding:SunJCE,
MessageDigest.SHA-256:SUN
```

In this syntax:

ServiceType
MessageDigest

Algorithm
AES/GCM/NoPadding, SHA-256

Provider
SunJCE, SUN

Customizing the Default Keystores and Truststores, Store Types, and Store Passwords

Whenever a default `SSLSocketFactory` or `SSLServerSocketFactory` is created (via a call to `SSLSocketFactory.getDefault` or `SSLServerSocketFactory.getDefault`), and this default `SSLSocketFactory` (or `SSLServerSocketFactory`) comes from the JSSE reference implementation, a default `SSLContext` is associated with the socket factory. (The default socket factory will come from the JSSE implementation.)

This default `SSLContext` is initialized with a default `KeyManager` and a default `TrustManager`. If a keystore is specified by the `javax.net.ssl.keyStore` system property and an appropriate `javax.net.ssl.keyStorePassword` system property (see How to Specify a java.lang.System Property), then the `KeyManager` created by the default `SSLContext` will be a `KeyManager` implementation for managing the specified keystore. (The actual implementation will be as specified in Customizing the Default Key Managers and Trust Managers.) If no such system property is specified, then the keystore managed by the `KeyManager` will be a new empty keystore.

Generally, the peer acting as the server in the handshake will need a keystore for its KeyManager in order to obtain credentials for authentication to the client. However, if one of the anonymous cipher suites is selected, then the server's `KeyManager` keystore is not necessary. And, unless the server requires client authentication, the peer acting as the client does not need a `KeyManager` keystore. Thus, in these situations it may be OK if no `javax.net.ssl.keyStore` system property value is defined.

Similarly, if a truststore is specified by the `javax.net.ssl.trustStore` system property, then the `TrustManager` created by the default `SSLContext` will be a `TrustManager` implementation for managing the specified truststore. In this case, if such a property exists but the file it specifies does not, then no truststore is used. If no `javax.net.ssl.trustStore` property exists, then a default truststore is searched for. If a truststore named *`java-home`*`/lib/security/jssecacerts` is found, it is used. If not, then a truststore named *`java-home`*`/lib/security/cacerts` is searched for and used (if it exists). Finally, if a truststore is still not found, then the truststore managed by the `TrustManager` will be a new empty truststore.

> **Note:**
>
> The JDK ships with a limited number of trusted root certificates in the `java-home/lib/security/cacerts` file. As documented in keytool in *Java Platform, Standard Edition Tools Reference*, it is your responsibility to maintain (that is, add and remove) the certificates contained in this file if you use this file as a truststore.
>
> Depending on the certificate configuration of the servers that you contact, you may need to add additional root certificates. Obtain the needed specific root certificates from the appropriate vendor.

If the `javax.net.ssl.keyStoreType` and/or `javax.net.ssl.keyStorePassword` system properties are also specified, then they are treated as the default `KeyManager` keystore type and password, respectively. If no type is specified, then the default type is that returned by the `KeyStore.getDefaultType()` method, which is the value of the `keystore.type` Security Property, or "jks" if no such Security Property is specified. If no keystore password is specified, then it is assumed to be a blank string "".

Similarly, if the `javax.net.ssl.trustStoreType` and/or `javax.net.ssl.trustStorePassword` system properties are also specified, then they are treated as the default truststore type and password, respectively. If no type is specified, then the default type is that returned by the `KeyStore.getDefaultType()` method. If no truststore password is specified, then it is assumed to be a blank string "".

> **Note:**
>
> This section describes the current JSSE reference implementation behavior. The system properties described in this section are not guaranteed to continue to have the same names and types (system or security) or even to exist at all in future releases. They are also not guaranteed to be examined and used by any other JSSE implementations. If they *are* examined by an implementation, then that implementation should handle them in the same manner as the JSSE reference implementation does, as described herein.

Customizing the Default Key Managers and Trust Managers

As noted in Customizing the Default Keystores and Truststores, Store Types, and Store Passwords, whenever a default `SSLSocketFactory` or `SSLServerSocketFactory` is created, and this default `SSLSocketFactory` (or `SSLServerSocketFactory`) comes from the JSSE reference implementation, a default `SSLContext` is associated with the socket factory.

This default `SSLContext` is initialized with a `KeyManager` and a `TrustManager`. The `KeyManager` and/or `TrustManager` supplied to the default `SSLContext` will be an implementation for managing the specified keystore or truststore, as described in the aforementioned section.

The `KeyManager` implementation chosen is determined by first examining the `ssl.KeyManagerFactory.algorithm` Security Property. If such a property value is specified, then a `KeyManagerFactory` implementation for the specified algorithm is searched for. The implementation from the first provider that supplies an implementation is used. Its `getKeyManagers()` method is called to determine the `KeyManager` to supply to the default `SSLContext`. Technically, `getKeyManagers()` returns an array of `KeyManager` objects, one `KeyManager` for each type of key material. If no such Security Property value is specified, then the default value of SunX509 is used to perform the search.

> **Note:**
>
> A `KeyManagerFactory` implementation for the SunX509 algorithm is supplied by the SunJSSE provider. The `KeyManager` that it specifies is a `javax.net.ssl.X509KeyManager` implementation.

Similarly, the `TrustManager` implementation chosen is determined by first examining the `ssl.TrustManagerFactory.algorithm` Security Property. If such a property value is specified, then a `TrustManagerFactory` implementation for the specified algorithm is searched for. The implementation from the first provider that supplies an implementation is used. Its `getTrustManagers()` method is called to determine the `TrustManager` to supply to the default `SSLContext`. Technically, `getTrustManagers()` returns an array of `TrustManager` objects, one `TrustManager` for each type of trust material. If no such Security Property value is specified, then the default value of PKIX is used to perform the search.

> **Note:**
>
> A `TrustManagerFactory` implementation for the PKIX algorithm is supplied by the SunJSSE provider. The `TrustManager` that it specifies is a `javax.net.ssl.X509TrustManager` implementation.

> **Note:**
>
> This section describes the current JSSE reference implementation behavior. The system properties described in this section are not guaranteed to continue to have the same names and types (system or security) or even to exist at all in future releases. They are also not guaranteed to be examined and used by any other JSSE implementations. If they *are* examined by an implementation, then that implementation should handle them in the same manner as the JSSE reference implementation does, as described herein.

Disabled and Restricted Cryptographic Algorithms

In some environments, certain algorithms or key lengths may be undesirable when using TLS/DTLS. The Oracle JDK uses the `jdk.certpath.disabledAlgorithms` and `jdk.tls.disabledAlgorithm` Security Properties to disable algorithms during TLS/DTLS protocol negotiation, including version negotiation, cipher suites selection, peer

authentication, and key exchange mechanisms. Note that these Security Properties are not guaranteed to be used by other JDK implementations. See the `<java-home>/conf/security/java.security` file for information about the syntax of these Security Properties and their current active values.

- **`jdk.certpath.disabledAlgorithms` Property**: `CertPath` code uses the `jdk.certpath.disabledAlgorithms` Security Property to determine which algorithms should not be allowed during `CertPath` checking. For example, when a TLS server sends an identifying certificate chain, a client `TrustManager` that uses a `CertPath` implementation to verify the received chain will not allow the stated conditions. For example, the following line blocks any MD2-based certificate, as well as SHA1 TLSServer certificates that chain to trust anchors that are pre-installed in the `cacaerts` keystore. Likewise, this line blocks any RSA key less than 1024 bits.

  ```
  jdk.certpath.disabledAlgorithms=MD2, SHA1 jdkCA & usage TLSServer, RSA
  keySize < 1024
  ```

- **`jdk.tls.disabledAlgorithms` Property**: `SunJSSE` code uses the `jdk.tls.disabledAlgorithms` Security Property to disable TLS/DTLS protocols, cipher suites, keys, and so on. The syntax is similar to the `jdk.certpath.disabledAlgorithms` Security Property. For example, the following line disables the SSLv3 algorithm and all of the TLS_*_RC4_* cipher suites:

  ```
  jdk.tls.disabledAlgorithms=SSLv3, RC4
  ```

If you require a particular condition, you can reactivate it by either removing the associated value in the Security Property in the `java.security` file or dynamically setting the proper Security Property before JSSE is initialized.

> **Note:**
>
> Contact your security architect before modifying these Security Properties or enabling a cipher suite that hasn't been enabled; this allows the use of cipher suites with weaker protections.

Note that these Security Properties effectively create a third set of cipher suites, Disabled. The following list describes these three sets:

- **Disabled**: If a cipher suite contains any components (for example, RC4) on the disabled list (for example, `RC4` is specified in the `jdk.tls.disabledAlgorithms` Security Property), then that cipher suite is disabled and will **not** be considered for a connection handshake.
- **Enabled**: A list of specific cipher suites that will be considered for a connection.
- **Not Enabled**: A list of non-disabled cipher suites that will **not** be considered for a connection. To re-enable these cipher suites, call the appropriate `setEnabledCipherSuites()` or `setSSLParameters()` methods.

If any application attempts to reenable a cipher suite, which has been disabled by the `jdk.tls.disabledAlgorithms` Security Property, through the `setEnabledCipherSuites()` or `setSSLParameters()` methods, then JSSE

allows the method call but does not allow the use of the disabled cipher suite during handshaking.

See SunJSSE Cipher Suites for a list of currently implemented SunJSSE cipher suites for this JDK release.

Legacy Cryptographic Algorithms

In some environments, a certain algorithm may be undesirable but it cannot be disabled because of its use in legacy applications. Legacy algorithms may still be supported, but applications should not use them as the security strength of legacy algorithms is usually not strong enough. During TLS/DTLS security parameters negotiation, legacy algorithms are not negotiated unless there are no other candidates. The Security Property `jdk.tls.legacyAlgorithms` specifies which algorithms the Oracle JDK considers as legacy algorithms. `<java-home>/conf/security/java.security` file for the syntax of this Security Property.

> **Note:**
>
> - If a legacy algorithm is also restricted through the `jdk.tls.disabledAlgorithms` property or the `java.security.AlgorithmConstraints` API (see the method `javax.net.ssl.SSLParameters.setAlgorithmConstraints`), then the algorithm is completely disabled and will not be negotiated.
> - If your application uses an algorithm specified in the Security Property `jdk.tls.legacyAlgorithms`, use an alternative algorithm as soon as possible; a future JDK release may specify a legacy algorithm as a restricted algorithm.

Customizing the Encryption Algorithm Providers

The SunJSSE provider uses the SunJCE implementation for all its cryptographic needs. Although it is recommended that you leave the provider at its regular position, you can use implementations from other JCA or JCE providers by registering them *before* the SunJCE provider.

The standard JCA mechanism (see How Provider Implementations Are Requested and Supplied) can be used to configure providers, either statically via the security properties file `<java-home>/conf/security/java.security`, or dynamically via the `addProvider()` or `insertProviderAt()` method in the `java.security.Security` class.

Customizing Size of Ephemeral Diffie-Hellman Keys

In TLS/DTLS connections, ephemeral Diffie-Hellman (DH) keys may be used internally during the handshaking. The SunJSSE provider provides a flexible approach to customize the strength of the ephemeral DH key size during TLS/DTLS handshaking.

Diffie-Hellman (DH) keys of sizes less than 1024 bits have been deprecated because of their insufficient strength. You can customize the ephemeral DH key size with the system property `jdk.tls.ephemeralDHKeySize`. This system property does not impact DH key sizes in `ServerKeyExchange` messages for exportable cipher suites. It impacts

only the DHE_RSA, DHE_DSS, and DH_anon-based cipher suites in the JSSE Oracle provider.

You can specify one of the following values for this property:

- Undefined: A DH key of size 1024 bits will be used always for non-exportable cipher suites. This is the default value for this property.
- `legacy`: The JSSE Oracle provider preserves the legacy behavior (for example, using ephemeral DH keys of sizes 512 bits and 768 bits) of JDK 7 and earlier releases.
- `matched`: For non-exportable anonymous cipher suites, the DH key size in ServerKeyExchange messages is 1024 bits. For X.509 certificate based authentication (of non-exportable cipher suites), the DH key size matching the corresponding authentication key is used, except that the size must be between 1024 bits and 2048 bits. For example, if the public key size of an authentication certificate is 2048 bits, then the ephemeral DH key size should be 2048 bits unless the cipher suite is exportable. This key sizing scheme keeps the cryptographic strength consistent between authentication keys and key-exchange keys.
- A valid integer between 1024 and 2048, inclusively: A fixed ephemeral DH key size of the specified value, in bits, will be used for non-exportable cipher suites.

The following table summaries the minimum and maximum acceptable DH key sizes for each of the possible values for the system property `jdk.tls.ephemeralDHKeySize`:

Table 8-4 DH Key Sizes for the System Property `jdk.tls.ephemeralDHKeySize`

Value of jdk.tls.ephemeralDHKeySize	Undefined	legacy	matched	Integer value (fixed)
Exportable DH key size	512	512	512	512
Non-exportable anonymous cipher suites	1024	768	1024	The fixed key size is specified by a valid integer property value, which must be between 1024 and 2048, inclusively.

Table 8-4 (Cont.) DH Key Sizes for the System Property `jdk.tls.ephemeralDHKeySize`

Value of jdk.tls.ephemeralDHKeySize	Undefined	legacy	matched	Integer value (fixed)
Authentication certificate	1024	768	The key size is the same as the authentication certificate, but must be between 1024 bits and 2048 bits, inclusively. However, the only DH key size that the SunJCE provider supports that is larger than 1024 bits is 2048 bits. Consequently, you may use the values 1024 or 2048 only.	The fixed key size is specified by a valid integer property value, which must be between 1024 and 2048, inclusively.

Customizing Maximum Fragment Length Negotiation (MFLN) Extension

In order to negotiate smaller maximum fragment lengths, clients have an option to include an extension of type `max_fragment_length` in the ClientHello message. A system property `jsse.enableMFLNExtension`, can be used to enable or disable the MFLN extension for TLS/DTLS.

Maximum Fragment Length Negotiation

It may be desirable for constrained TLS/DTLS clients to negotiate a smaller maximum fragment length due to memory limitations or bandwidth limitations. In order to negotiate smaller maximum fragment lengths, clients have an option to include an extension of type `max_fragment_length` in the (extended) ClientHello message. See RFC 6066.

Once a maximum fragment length has been successfully negotiated, the TLS/DTLS client and server can immediately begin fragmenting messages (including handshake messages) to ensure that no fragment larger than the negotiated length is sent.

System Property jsse.enableMFLNExtension

A system property `jsse.enableMFLNExtension` is defined to enable or disable the MFLN extension. The `jsse.enableMFLNExtension` is disabled by default.

The value of the system property can be set as follows:

Table 8-5 jsse.enableMFLNExtension system property

System Property	Description
`jsse.enableMFLNExtension=true`	Enable the MFLN extension. If the returned value of `SSLParameters.getMaximumPacketSize()` is less than (2^{12} + header-size) the maximum fragment length negotiation extension would be enabled.
`jsse.enableMFLNExtension=false`	Disable the MFLN extension.

Configuring the Maximum and Minimum Packet Size

Set the maximum expected network packet size in bytes for a TLS/DTLS record with the `SSLParameters.setMaximumPacketSize` method.

It is recommended that the packet size should not be less than 256 bytes so that small handshake messages, such as HelloVerifyRequests, are not fragmented.

Limiting Amount of Data Algorithms May Encrypt with a Set of Keys

You can specify a limit on the amount of data an algorithm may encrypt with a specific set of keys with the `jdk.tls.keyLimits` Security Property. Once this limit is reached, a KeyUpdate post-handshake message is sent, which requests that the current set of keys be updated. This Security Property is only for symmetrical ciphers with TLS 1.3.

The syntax for this property is as follows:

```
jdk.tls.keyLimits=KeyLimit { , KeyLimit }
```

KeyLimit

```
AlgorithmName KeyUpdate Length
```

AlgorithmName
A full algorithm transformation

Length
The amount of encrypted data in a session before a KeyUpdate message is sent. This value may be an integer value in bytes or as a power of two, for example, `2^37`.

For example, the following specifies that a KeyUpdate message is sent once the algorithm AES/GCM/NoPadding has encrypted 2^{37} bytes:

```
jdk.tls.keyLimits=AES/GCM/NoPadding KeyUpdate 2^37
```

Client-Driven OCSP and OCSP Stapling

Use the Online Certificate Status Protocol (OCSP) to determine the X.509 certificate revocation status during the Transport Layer Security (TLS) handshake.

X.509 certificates used in TLS can be revoked by the issuing Certificate Authority (CA) if there is reason to believe that a certificate is compromised. You can check the revocation status of certificates during the TLS handshake by using one of the following approaches.

- **Certificate Revocation List (CRL)**: A CRL is a simple list of revoked certificates. The application receiving a certificate gets the CRL from a CRL server and checks if the certificate received is on the list. There are two disadvantages to using CRLs that mean a certificate could be revoked:
  - CRLs can become very large so there can be a substantial increase in network traffic.
  - Many CRLs are created with longer validity periods, which increases the possibility of a certificate being revoked within that validity period and not showing up until the next CRL refresh.

  See Certificate/CRL Storage Classes in Java PKI Programmer's Guide.
- **Client-driven OCSP** : In client-driven OCSP, the client uses OCSP to contact an OCSP responder to check the certificate's revocation status. The amount of data required is usually less than that of a CRL, and the OCSP responder is likely to be more up-to-date with the revocation status than a CRL. Each client connecting to a server requires an OCSP response for each certificate being checked. If the server is a popular one, and many of the clients are using clientdriven OCSP, these OCSP requests can have a negative effect on the performance of the OCSP responder.
- **OCSP stapling**: OCSP stapling enables the server, rather than the client, to make the request to the OCSP responder. The server staples the OCSP response to the certificate and returns it to the client during the TLS handshake. This approach enables the presenter of the certificate, rather than the issuing CA, to bear the resource cost of providing OCSP responses. It also enables the server to cache the OCSP responses and supply them to all clients. This significantly reduces the load on the OCSP responder because the response can be cached and periodically refreshed by the server rather than by each client.

Client-Driven OCSP and Certificate Revocation

Client-driven Online Certificate Status Protocol (OCSP) enables the client to check the certificate revocation status by connecting to an OCSP responder during the Transport Layer Security (TLS) handshake.

The client-driven OCSP request occurs during the TLS handshake just after the client receives the certificate from the server and validates it.

TLS Handshake with Client-Driven OCSP

Client-driven OCSP is used during the TLS handshake between the client and the server to check the server certificate revocation status. After the client receives the certificate, it performs certificate validation. If the validation is successful, then the client verifies that the certificate was not revoked by the issuer. This is done by sending an OCSP request to an OCSP responder. After receiving the OCSP response, the client checks this response before completing the TLS handshake.

Usually the client finds the OCSP responder's URL by looking in the Authority Information Access (AIA) extension of the certificate, but it can be set to a static URL through the use of a system property.

Setting up a Java Client to use Client-Driven OCSP

Client-driven OCSP is enabled by enabling revocation checking and enabling OCSP.

To configure a Java client to use client-driven OCSP, the Java client must already be set up to connect to a server using TLS.

1. Enable revocation checking. You can do this in two different ways.
   - Set the system property `com.sun.net.ssl.checkRevocation` to `true`.
   - Use the `setRevocationEnabled` method on `PKIXParameters`. See The PKIXParameters Class.
2. Enable client-driven OCSP:

   Set the Security Property `ocsp.enable` to `true`.

Both steps are necessary. The `ocsp.enable` setting has no effect unless revocation checking is enabled.

OCSP Stapling and Certificate Revocation

Online Certificate Status Protocol (OCSP) stapling enables the presenter of a certificate, rather than the issuing Certificate Authority (CA), to bear the resource cost of providing the OCSP responses that contain the certificate's revocation status.

TLS Handshake with OCSP Stapling

OCSP stapling is used during the Transport Layer Security (TLS) handshake between the client and the server to check the server certificate revocation status. The server makes the OCSP request to the OCSP responder and staples the OCSP responses to the certificates returned to the client. By having the server make the request to the OCSP responder, the responses can be cached, and then used multiple times for many clients.

The client receiving the certificates with stapled OCSP responses validates each certificate, and then checks the OCSP responses before continuing with the handshake. If, from the client's perspective, the stapled OCSP response from the server for a certificate is missing, the client will attempt to use client-driven OCSP or Certificate Revocation Lists (CRLs) to get revocation information if the following are true:

- The RevocationEnabled flag is set to true through the `PKIXParameters.setRecovcationEnabled` method.
- OCSP checking is enabled by setting the `ocsp.enable` Security Property to true.

OCSP checking works in conjunction with CRLs during revocation checking. See Appendix C: OCSP Support in Java PKI Programmer's Guide.

Status Request Versus Multiple Status Request

The OCSP stapling feature implements the TLS Certificate Status Request extension (section 8 of RFC 6066) and the Multiple Certificate Status Request Extension (RFC 6961).

The TLS Certificate Status Request extension requests revocation information for only the server certificate in the certificate chain while the Multiple Certificate Status Request Extension requests revocation information for all certificates in the certificate

chain. In the case where only the server certificate's revocation information is sent to the client, other certificates in the chain may be verified using the Certificate Revocation Lists (CRLs) or client-driven OCSP (but the client will need to be set up to do this).

Although TLS allows the server to also request the client's certificate, there is no provision in OCSP stapling that enables the client to contact the appropriate OCSP responder and staple the response to the certificate sent to the server.

The OCSP Request and Response

OCSP request and response messages are usually sent over unencrypted HTTP. The response is signed by the CA.

If necessary, the stapled responses can be obtained in the client code by calling the `getStatusResponses` method on the `ExtendedSSLSession` object. The method signature is:

```
public List<byte[]> getStatusResponses();
```

The OCSP response is encoded using the Distinguished Encoding Rules (DER) in a format described by the ASN.1 found in RFC 6960.

Setting Up a Java Client to Use OCSP Stapling

Online Certificate Status Protocol (OCSP) stapling is enabled on the client side by setting the system property `jdk.tls.client.enableStatusRequestExtension` to `true` (its default value).

To configure a Java client to make use of the OCSP response stapled to the certificate returned by a server, the Java client must already be set up to connect to a server using TLS, and the server must be set up to staple an OCSP response to the certificate it returns part of the TLS handshake.

1. Enable OCSP stapling on the client:

   If necessary, set the system property `jdk.tls.client.enableStatusRequestExtension` to `true`.

2. Enable revocation checking. You can do this in two different ways.

   - Set the system property `com.sun.net.ssl.checkRevocation` to `true`. You can do this from the command line or in the code.
   - Use the `setRevocationEnabled` method on the `PKIXParameters` class. See The PKIXParameters Class.

   For the client to include the stapled responses received from the server in the certificate validation, revocation checking must be set to `true`. If revocation checking is not set to `true`, then the connection will be allowed to proceed regardless of the presence or status of the revocation information

Setting Up a Java Server to Use OCSP Stapling

Online Certificate Status Protocol (OCSP) stapling is enabled on the server by setting the system property `jdk.tls.server.enableStatusRequestExtension` to `true`. (It is set to `false` by default.)

The following steps can be used to configure a Java server to connect to an OCSP responder and staple the OCSP response to the certificate to be returned to the client. The Java server must already be set up to respond to clients using TLS.

1. Enable OCSP stapling on the server:

   Set the system property `jdk.tls.server.enableStatusRequestExtension` to `true`.

2. Optional: Set other properties as required. See OCSP Stapling Configuration Properties for a list of the valid properties.

OCSP Stapling Configuration Properties

This topic lists the effects of setting various properties when using the Online Certificate Status Protocol (OCSP). It shows the properties used in both client-driven OCSP and OCSP stapling.

Server-side Properties

Most of the properties are read at `SSLContext` instantiation time. This means that if you set a property, you must obtain a new `SSLContext` object so that an `SSLSocket` or `SSLEngine` object you obtain from that `SSLContext` object will reflect the property setting. The one exception is the `jdk.tls.stapling.responseTimeout` property. That property is evaluated when the `ServerHandshaker` object is created (essentially at the same time that an `SSLSocket` or `SSLEngine` object gets created).

Table 8-6 Server-Side OCSP stapling Properties

Property	Description	Default Value
`jdk.tls.server.enableStatusRequestExtension`	Enables the server-side support for OCSP stapling.	False
`jdk.tls.stapling.responseTimeout`	Controls the maximum amount of time the server will use to obtain OCSP responses, whether from the cache or by contacting an OCSP responder. The responses that are already received will be sent in a `CertificateStatus` message, if applicable based on the type of stapling being done.	5000 (integer value in milliseconds)
`jdk.tls.stapling.cacheSize`	Controls the maximum cache size in entries. If the cache is full and a new response needs to be cached, then the least recently used cache entry will be replaced with the new one. A value of zero or less for this property means that the cache will have no upper bound on the number of responses it can contain.	256 objects

Table 8-6 (Cont.) Server-Side OCSP stapling Properties

Property	Description	Default Value
`jdk.tls.stapling.cacheLifetime`	Controls the maximum life of a cached response. It is possible for responses to have shorter lifetimes than the value set with this property if the response has a **nextUpdate** field that expires sooner than the cache lifetime. A value of zero or less for this property disables the cache lifetime. If an object has no **nextUpdate** value and cache lifetimes are disabled, then the response will not be cached.	3600 seconds (1 hour)
`jdk.tls.stapling.responderURI`	Enables the administrator to set a default URI in the event that certificates used for TLS do not have the Authority Info Access (AIA) extension. It will not override the Authority Info Access extension value unless the `jdk.tls.stapling.responderOverride` property is set.	Not set
`jdk.tls.stapling.responderOverride`	Enables a URI provided through the `jdk.tls.stapling.responderURI` property to override any AIA extension value.	False
`jdk.tls.stapling.ignoreExtensions`	Disables the forwarding of OCSP extensions specified in the `status_request` or `status_request_v2` TLS extensions.	False

Client-Side Settings

Table 8-7 Client-Side Settings Used in OCSP Stapling

PKIXBuilderParameters	checkRevocation Property	PKIXRevocationChecker	Result
Default	Default	Default	Revocation checking is disabled.
Default	True	Default	Revocation checking is enabled.[1]
Instantiated	Default	Default	Revocation checking is enabled.[1]
Instantiated	Default	Instantiated, added to `PKIXBuilderParameters` object.	Revocation checking is enabled and[1]will behave according to the `PKIXRevocationChecker` settings.

Footnote 1 Note that client-side OCSP fallback will occur only if the `ocsp.enable` Security Property is set to `true`.

Developers have some flexibility in how to handle the responses provided through OCSP stapling. OCSP stapling makes no changes to the current methodologies involved in certificate path checking and revocation checking. This means that it is possible to have both client and server assert the `status_request` extensions, obtain OCSP responses through the `CertificateStatus` message, and provide user flexibility in how to react to revocation information, or the lack thereof.

If no `PKIXBuilderParameters` is provided by the caller, then revocation checking is disabled. If the caller creates a `PKIXBuilderParameters` object and uses the `setRevocationEnabled` method to enable revocation checking, then stapled OCSP responses will be evaluated. This is also the case if the `com.sun.net.ssl.checkRevocation` property is set to `true`.

Hardware Acceleration and Smartcard Support

The Java Cryptography Architecture (JCA) is a set of packages that provides a framework and implementations for encryption, key generation and key agreement, and message authentication code (MAC) algorithms. (See Java Cryptography Architecture (JCA) Reference Guide.) The SunJSSE provider uses JCA exclusively for all of its cryptographic operations and can automatically take advantage of JCE features and enhancements, including JCA's support for RSA PKCS#11. This support enables the SunJSSE provider to use hardware cryptographic accelerators for significant performance improvements and to use smartcards as keystores for greater flexibility in key and trust management. .

Use of hardware cryptographic accelerators is automatic if JCA has been configured to use the Oracle PKCS#11 provider, which in turn has been configured to use the underlying accelerator hardware. The provider must be configured before any other JCA providers in the provider list. For details on how to configure the Oracle PKCS#11 provider, see PKCS#11 Reference Guide.

Configuring JSSE to Use Smartcards as Keystores and Truststores

Support for PKCS#11 in JCA also enables access to smartcards as a keystore. For details on how to configure the type and location of the keystores to be used by JSSE, see Customizing JSSE. To use a smartcard as a keystore or truststore, set the `javax.net.ssl.keyStoreType` and `javax.net.ssl.trustStoreType` system properties, respectively, to `pkcs11`, and set the `javax.net.ssl.keyStore` and `javax.net.ssl.trustStore` system properties, respectively, to `NONE`. To specify the use of a specific provider, use the `javax.net.ssl.keyStoreProvider` and `javax.net.ssl.trustStoreProvider` system properties (for example, set them to `SunPKCS11-joe`). By using these properties, you can configure an application that previously depended on these properties to access a file-based keystore to use a smartcard keystore with no changes to the application.

Some applications request the use of keystores programmatically. These applications can continue to use the existing APIs to instantiate a `Keystore` and pass it to its key manager and trust manager. If the `Keystore` instance refers to a PKCS#11 keystore backed by a Smartcard, then the JSSE application will have access to the keys on the smartcard.

Multiple and Dynamic Keystores

Smartcards (and other removable tokens) have additional requirements for an `X509KeyManager`. Different smartcards can be present in a smartcard reader during the lifetime of a Java application, and they can be protected using different passwords.

The `KeyStore.Builder` class abstracts the construction and initialization of a `KeyStore` object. It supports the use of `CallbackHandler` for password prompting, and its subclasses can be used to support additional features as desired by an application. For example, it is possible to implement a `Builder` that allows individual `KeyStore` entries to be protected with different passwords. The `KeyStoreBuilderParameters` class then can be used to initialize a `KeyManagerFactory` using one or more of these `Builder` objects.

A `X509KeyManager` implementation in the SunJSSE provider called NewSunX509 supports these parameters. If multiple certificates are available, it attempts to pick a certificate with the appropriate key usage and prefers valid to expired certificates.

Example 8-17 illustrates how to tell JSSE to use both a PKCS#11 keystore (which might in turn use a smartcard) and a PKCS#12 file-based keystore.

Example 8-17 Sample Code to Use PKCS#11 and PKCS#12 File-based Keystore

```
import javax.net.ssl.*;
import java.security.KeyStore.*;
// ...

// Specify keystore builder parameters for PKCS#11 keystores
Builder scBuilder = Builder.newInstance("PKCS11", null,
    new CallbackHandlerProtection(myGuiCallbackHandler));

// Specify keystore builder parameters for a specific PKCS#12 keystore
Builder fsBuilder = Builder.newInstance("PKCS12", null,
    new File(pkcsFileName), new PasswordProtection(pkcsKsPassword));

// Wrap them as key manager parameters
ManagerFactoryParameters ksParams = new KeyStoreBuilderParameters(
    Arrays.asList(new Builder[] { scBuilder, fsBuilder }) );

// Create KeyManagerFactory
KeyManagerFactory factory = KeyManagerFactory.getInstance("NewSunX509");

// Pass builder parameters to factory
factory.init(ksParams);

// Use factory
SSLContext ctx = SSLContext.getInstance("TLS");
ctx.init(factory.getKeyManagers(), null, null);
```

Additional Keystore Formats (PKCS12)

The PKCS#12 (Personal Information Exchange Syntax Standard) specifies a portable format for storage and/or transport of a user's private keys, certificates, miscellaneous secrets, and other items. The SunJSSE provider supplies a complete implementation of the PKCS12 `java.security.KeyStore` format for reading and writing PKCS12 files. This format is also supported by other toolkits and applications for importing and exporting keys and certificates, such as Mozilla Firefox, Microsoft Internet Explorer, and OpenSSL. For example, these implementations can export client certificates and keys into a file using the .p12 file name extension.

With the SunJSSE provider, you can access PKCS12 keys through the `KeyStore` API with a keystore type of PKCS12. In addition, you can list the installed keys and associated certificates by using the `keytool` command with the `-storetype` option set to `pkcs12`. See keytool in *Java Platform, Standard Edition Tools Reference*.

Server Name Indication (SNI) Extension

The SNI extension is a feature that extends the TLS/DTLS protocol to indicate what server name the client is attempting to connect to during handshaking. Servers can use server name indication information to decide if specific `SSLSocket` or `SSLEngine` instances should accept a connection. For example, when multiple virtual or name-based servers are hosted on a single underlying network address, the server application can use SNI information to determine whether this server is the exact server that the client wants to access. Instances of this class can be used by a server to verify the acceptable server names of a particular type, such as host names. See section 3 of TLS Extensions (RFC 6066).

Developers of client applications can explicitly set the server name indication using the `SSLParameters.setServerNames(List<SNIServerName> serverNames)` method. See Example 8-18.

Developers of server applications can use the `SNIMatcher` class to decide how to recognize server name indication. Example 8-19 and Example 8-20 illustrate this functionality:

Example 8-18 Sample Code to Set the Server Name Indication

The following code sample illustrates how to set the server name indication using the method `SSLParameters.setServerNames(List<SNIServerName> serverNames)`:

```
SSLSocketFactory factory = ...
SSLSocket sslSocket = factory.createSocket("172.16.10.6", 443);
// SSLEngine sslEngine = sslContext.createSSLEngine("172.16.10.6", 443);

SNIHostName serverName = new SNIHostName("www.example.com");
List<SNIServerName> serverNames = new ArrayList<>(1);
serverNames.add(serverName);

SSLParameters params = sslSocket.getSSLParameters();
params.setServerNames(serverNames);
sslSocket.setSSLParameters(params);
// sslEngine.setSSLParameters(params);
```

Example 8-19 Sample Code Using SSLSocket Class to Recognize SNI

The following code sample illustrates how the server applications can use the `SNIMatcher` class to decide how to recognize server name indication:

```
SSLSocket sslSocket = sslServerSocket.accept();

SNIMatcher matcher = SNIHostName.createSNIMatcher("www\\.example\\.(com|
org)");
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);

SSLParameters params = sslSocket.getSSLParameters();
params.setSNIMatchers(matchers);
sslSocket.setSSLParameters(params);
```

Example 8-20 Sample Code Using SSLServerSocket Class to Recognize SNI

The following code sample illustrates how the server applications can use the `SNIMatcher` class to decide how to recognize server name indication:

```
SSLServerSocket sslServerSocket = ...;

SNIMatcher matcher = SNIHostName.createSNIMatcher("www\\.example\\.(com|
org)");
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);

SSLParameters params = sslServerSocket.getSSLParameters();
params.setSNIMatchers(matchers);
sslServerSocket.setSSLParameters(params);

SSLSocket sslSocket = sslServerSocket.accept();
```

The following list provides examples for the behavior of the `SNIMatcher` when receiving various server name indication requests in the ClientHello message:

- Matcher configured to `www\\.example\\.com`:
  - If the requested host name is `www.example.com`, then it will be accepted and a confirmation will be sent in the ServerHello message.
  - If the requested host name is `www.example.org`, then it will be rejected with an `unrecognized_name` fatal error.
  - If there is no requested host name or it is empty, then the request will be accepted but no confirmation will be sent in the ServerHello message.
- Matcher configured to `www\\.invalid\\.com`:
  - If the requested host name is `www.example.com`, then it will be rejected with an `unrecognized_name` fatal error.
  - If the requested host name is `www.example.org`, then it will be accepted and a confirmation will be sent in the ServerHello message.

  - If there is no requested host name or it is empty, then the request will be accepted but no confirmation will be sent in the ServerHello message.
- Matcher is not configured:

  Any requested host name will be accepted but no confirmation will be sent in the ServerHello message.

For descriptions of new classes that implement the SNI extension, see:

- StandardConstants Class
- SNIServerName Class
- SNIMatcher Class
- SNIHostName Class

For examples, see Using the Server Name Indication (SNI) Extension.

TLS Application Layer Protocol Negotiation

Negotiate an application protocol for a TLS connection with Application Layer Protocol Negotiation (ALPN).

What is ALPN?

Some applications might want or need to negotiate a shared application level value before a TLS handshake has completed. For example, HTTP/2 uses the Application Layer Protocol Negotiation mechanism to help establish which HTTP version ("h2", "spdy/3", "http/1.1") can or will be used on a particular TCP or UDP port. ALPN (RFC 7301) does this without adding network round-trips between the client and the server. In the case of HTTP/2 the protocol must be established before the connection is negotiated, as client and server need to know what version of HTTP to use before they start communicating. Without ALPN it would not be possible to have application protocols HTTP/1 and HTTP/2 on the same port.

The client uses the ALPN extension at the beginning of the TLS handshake to send a list of supported application protocols to the server as part of the `ClientHello`. The server reads the list of supported application protocols in the `ClientHello`, and determines which of the supported protocols it prefers. It then sends a `ServerHello` message back to the client with the negotiation result. The message may contain either the name of the protocol that has been chosen or that no protocol has been chosen.

The application protocol negotiation can thus be accomplished within the TLS handshake, without adding network round-trips, and allows the server to associate a different certificate with each application protocol, if desired.

Unlike many other TLS extensions, this extension does not establish properties of the session, only of the connection. That's why you'll find the negotiated values in the `SSLSocket`/`SSLEngine`, not the `SSLSession`. When session resumption or session tickets are used (see TLS Session Resumption without Server-Side State), the previously negotiated values are irrelevant, and only the values in the new handshake messages are considered.

Setting up ALPN on the Client

Set the Application Layer Protocol Negotiation (ALPN) values supported by the client. During the handshake with the server, the server will read the client's list of application protocols and will determine which is most suitable.

For the client, use the `SSLParameters.setApplicationProtocols(String[])` method, followed by the `setSSLParameters` method of either `SSLSocket` or `SSLEngine` to set up the application protocols to send to the server.

Example 8-21 Sample Code for Setting and Getting ALPN Values in a Java Client

For example, here are the steps to set ALPN values of `"three"` and `"two"`, on the client.

To run the code the property `javax.net.ssl.trustStore` must be set to a valid root certificate. (This can be done on the command line).

```
import java.io.*;
import java.util.*;
import javax.net.ssl.*;
public class SSLClient {
    public static void main(String[] args) throws Exception {

        // Code for creating a client side SSLSocket
        SSLSocketFactory sslsf = (SSLSocketFactory)
SSLSocketFactory.getDefault();
        SSLSocket sslSocket = (SSLSocket) sslsf.createSocket("localhost",
9999);

        // Get an SSLParameters object from the SSLSocket
        SSLParameters sslp = sslSocket.getSSLParameters();

        // Populate SSLParameters with the ALPN values
        // On the client side the order doesn't matter as
        // when connecting to a JDK server, the server's list takes
priority
        String[] clientAPs = {"three", "two"};
        sslp.setApplicationProtocols(clientAPs);

        // Populate the SSLSocket object with the SSLParameters object
        // containing the ALPN values
        sslSocket.setSSLParameters(sslp);

        sslSocket.startHandshake();

        // After the handshake, get the application protocol that has been
negotiated
        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol client side: \"" + ap +
"\"");

        // Do simple write/read
        InputStream sslIS = sslSocket.getInputStream();
```

```
        OutputStream sslOS = sslSocket.getOutputStream();
        sslOS.write(280);
        sslOS.flush();
        sslIS.read();
        sslSocket.close();
    }
}
```

When this code is run and sends a `ClientHello` to a Java server that has set the ALPN values `one`, `two`, and `three`, the output will be:

```
Application Protocol client side: two
```

It is also possible to check the results of the negotiation during handshaking. See Determining Negotiated ALPN Value during Handshaking.

Setting up Default ALPN on the Server

Use the default ALPN mechanism to determine a suitable application protocol by setting ALPN values on the server.

To use the default mechanism for ALPN on the server, populate an `SSLParameters` object with the ALPN values you wish to set, and then use this `SSLParameters` object to populate either the `SSLSocket` object or the `SSLEngine` object with these parameters as you have done when you set up ALPN on the client (see the section Setting up ALPN on the Client). The first value of the ALPN values set on the server that matches any of the ALPN values contained in the `ClientHello` will be chosen and returned to the client as part of the `ServerHello`.

Example 8-22 Sample Code for Default ALPN Value Negotiation on the Server

Here is the code for a Java server that uses the default approach for protocol negotiation. To run the code the property `javax.net.ssl.keyStore` must be set to a valid keystore. (This can be done on the command line, see Creating a Keystore to Use with JSSE).

```
import java.util.*;
import javax.net.ssl.*;
public class SSLServer {
    public static void main(String[] args) throws Exception {

        // Code for creating a server side SSLSocket
        SSLServerSocketFactory sslssf =
            (SSLServerSocketFactory) SSLServerSocketFactory.getDefault();
        SSLServerSocket sslServerSocket =
            (SSLServerSocket) sslssf.createServerSocket(9999);
        SSLSocket sslSocket = (SSLSocket) sslServerSocket.accept();

        // Get an SSLParameters object from the SSLSocket
        SSLParameters sslp = sslSocket.getSSLParameters();

        // Populate SSLParameters with the ALPN values
        // As this is server side, put them in order of preference
        String[] serverAPs ={ "one", "two", "three" };
```

```
        sslp.setApplicationProtocols(serverAPs);

        // If necessary at any time, get the ALPN values set on the
        // SSLParameters object with:
        // String serverAPs = sslp.setApplicationProtocols();

        // Populate the SSLSocket object with the ALPN values
        sslSocket.setSSLParameters(sslp);

        sslSocket.startHandshake();

        // After the handshake, get the application protocol that
        // has been negotiated

        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol server side: \"" + ap +
"\"");

        // Continue with the work of the server
        InputStream sslIS = sslSocket.getInputStream();
        OutputStream sslOS = sslSocket.getOutputStream();
        sslIS.read();
        sslOS.write(85);
        sslOS.flush();
        sslSocket.close();
    }
}
```

When this code is run and a Java client sends a `ClientHello` with ALPN values `three` and `two`, the output is:

```
Application Protocol server side: two
```

It is also possible to check the results of the negotiation during handshaking. See Determining Negotiated ALPN Value during Handshaking.

Setting up Custom ALPN on the Server

Use the custom ALPN mechanism to determine a suitable application protocol by setting up a callback method.

If you do not want to use the server's default negotiation protocol, you can use the `setHandshakeApplicationProtocolSelector` method of `SSLEngine` or `SSLSocket` to register a `BiFunction` (lambda) callback that can examine the handshake state so far, and then make your selection based on the client's list of application protocols and any other relevant information. For example, you may consider using the cipher suite suggested, or the Server Name Indication (SNI) or any other data you can obtain in making the choice. If custom negotiation is used, the values set by the `setApplicationProtocols` method (default negotiation) will be ignored.

Example 8-23 Sample Code for Custom ALPN Value Negotiation on the Server

Here is the code for a Java server that uses the custom mechanism for protocol negotiation. To run the code the property `javax.net.ssl.keyStore` must be set to a

valid certificate. (This can be done on the command line, see Creating a Keystore to Use with JSSE).

```
import java.util.*;
import javax.net.ssl.*;
public class SSLServer {
    public static void main(String[] args) throws Exception {

        // Code for creating a server side SSLSocket
        SSLServerSocketFactory sslssf =
            (SSLServerSocketFactory) SSLServerSocketFactory.getDefault();
        SSLServerSocket sslServerSocket =
            (SSLServerSocket) sslssf.createServerSocket(9999);
        SSLSocket sslSocket = (SSLSocket) sslServerSocket.accept();

        // Code to set up a callback function
        // Pass in the current SSLSocket to be inspected and client AP
values
        sslSocket.setHandshakeApplicationProtocolSelector(
            (serverSocket, clientProtocols) -> {
                SSLSession handshakeSession =
serverSocket.getHandshakeSession();
                // callback function called with current SSLSocket and
client AP values
                // plus any other useful information to help determine
appropriate
                // application protocol. Here the protocol and ciphersuite
are also
                // passed to the callback function.
                return chooseApplicationProtocol(
                    serverSocket,
                    clientProtocols,
                    handshakeSession.getProtocol(),
                    handshakeSession.getCipherSuite());
         });

        sslSocket.startHandshake();

        // After the handshake, get the application protocol that has been
        // returned from the callback method.

        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol server side: \"" + ap +
"\"");

        // Continue with the work of the server
        InputStream sslIS = sslSocket.getInputStream();
        OutputStream sslOS = sslSocket.getOutputStream();
        sslIS.read();
        sslOS.write(85);
        sslOS.flush();
        sslSocket.close();
    }

    // The callback method. Note how the parameters match the call within
```

```
    // the setHandshakeApplicationProtocolSelector method above.
    public static String chooseApplicationProtocol(SSLSocket serverSocket,
            List<String> clientProtocols, String protocol, String
cipherSuite ) {
        // For example, check the cipher suite and return an application
protocol
        // value based on that.
        if (cipherSuite.equals("<--a_particular_ciphersuite-->")) {
            return "three";
        } else {
            return "";
        }
    }
}
```

If the cipher suite matches the one you specify in the condition statement when this code is run , then the value `three` will be returned. Otherwise an empty string will be returned.

Note that the `BiFunction` object's return value is a `String`, which will be the application protocol name, or null to indicate that none of the advertised names are acceptable. If the return value is an empty `String` then application protocol indications will not be used. If the return value is null (no value chosen) or is a value that was not advertised by the peer, the underlying protocol will determine what action to take. (For example, the server code will send a "no_application_protocol" alert and terminate the connection.)

After handshaking completes on both client and server, you can check the result of the negotiation by calling the `getApplicationProtocol` method on either the `SSLSocket` object or the `SSLEngine` object.

Determining Negotiated ALPN Value during Handshaking

To determine the ALPN value that has been negotiated during the handshaking, create a custom `KeyManager` or `TrustManager` class, and include in this custom class a call to the `getHandshakeApplicationProtocol` method.

There are some use cases where the selected ALPN and SNI values will affect the choices made by a `KeyManager` or `TrustManager`. For example, an application might want to select different certificate/private key sets depending on the attributes of the server and the chosen ALPN/SNI/ciphersuite values.

The sample code given illustrates how to call the `getHandshakeApplicationProtocol` method from within a custom `X509ExtendedKeyManager` that you create and register as the `KeyManager` object.

Example 8-24 Sample Code for a Custom KeyManager

This example shows the entire code for a custom `KeyManager` that extends `X509ExtendedKeyManager`. Most methods simply return the value returned from the `KeyManager` class that is being wrapped by this `MyX509ExtendedKeyManager` class. However the `chooseServerAlias` method calls the

`getHandshakeApplicationProtocol` on the `SSLSocket` object and therefore can determine the current negotiated ALPN value.

```
import java.net.Socket;
import java.security.*;
import javax.net.ssl.*;

public class MyX509ExtendedKeyManager extends X509ExtendedKeyManager {

    // X509ExtendedKeyManager is an abstract class so your new class
    // needs to implement all the abstract methods in this class.
    // The easiest way to do this is to wrap an existing KeyManager
    // and call its methods for each of the methods you need to
implement.

    X509ExtendedKeyManager akm;

    public MyX509ExtendedKeyManager(X509ExtendedKeyManager akm) {
        this.akm = akm;
    }

    @Override
    public String[] getClientAliases(String keyType, Principal[] issuers) {
        return akm.getClientAliases(keyType, issuers);
    }

    @Override
    public String chooseClientAlias(String[] keyType, Principal[] issuers,
        Socket socket) {
        return akm.chooseClientAlias(keyType, issuers, socket);
    }

    @Override
    public String chooseServerAlias(String keyType, Principal[] issuers,
        Socket socket) {

        // This method has access to a Socket, so it is possible to call
the
        // getHandshakeApplicationProtocol method here. Note the cast from
        // a Socket to an SSLSocket
        String ap = ((SSLSocket) socket).getHandshakeApplicationProtocol();
        System.out.println("In chooseServerAlias, ap is: " + ap);
        return akm.chooseServerAlias(keyType, issuers, socket);
    }

    @Override
    public String[] getServerAliases(String keyType, Principal[] issuers) {
        return akm.getServerAliases(keyType, issuers);
    }

    @Override
    public X509Certificate[] getCertificateChain(String alias) {
        return akm.getCertificateChain(alias);
```

```
    }

    @Override
    public PrivateKey getPrivateKey(String alias) {
        return akm.getPrivateKey(alias);
    }
}
```

When this code is registered as the `KeyManager` for a Java server and a Java client sends a `ClientHello` with ALPN values, the output will be:

```
In chooseServerAlias, ap is: <negotiated value>
```

Example 8-25 Sample Code for Using a Custom KeyManager in a Java Server

This example shows a simple Java server that uses the default ALPN negotiation strategy and the custom `KeyManager`, `MyX509ExtendedKeyManager`, shown in the prior code sample.

```
import java.io.*;
import java.util.*;
import javax.net.ssl.*;
import java.security.KeyStore;

public class SSLServerHandshake {

    public static void main(String[] args) throws Exception {
        SSLContext ctx = SSLContext.getInstance("TLS");

        // You need to explicitly create a create a custom KeyManager

        // Keystores
        KeyStore keyKS = KeyStore.getInstance("PKCS12");
        keyKS.load(new FileInputStream("serverCert.p12"),
            "password".toCharArray());

        // Generate KeyManager
        KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
        kmf.init(keyKS, "password".toCharArray());
        KeyManager[] kms = kmf.getKeyManagers();

        // Code to substitute MyX509ExtendedKeyManager
        if (!(kms[0] instanceof X509ExtendedKeyManager)) {
            throw new Exception("kms[0] not X509ExtendedKeyManager");
        }

        // Create a new KeyManager array and set the first index
        // of the array to an instance of MyX509ExtendedKeyManager.
        // Notice how creating this object is done by passing in the
        // existing default X509ExtendedKeyManager
        kms = new KeyManager[] {
            new MyX509ExtendedKeyManager((X509ExtendedKeyManager) kms[0])};

        // Initialize SSLContext using the new KeyManager
```

```
        ctx.init(kms, null, null);

        // Instead of using SSLServerSocketFactory.getDefault(),
        // get a SSLServerSocketFactory based on the SSLContext
        SSLServerSocketFactory sslssf = ctx.getServerSocketFactory();
        SSLServerSocket sslServerSocket =
            (SSLServerSocket) sslssf.createServerSocket(9999);
        SSLSocket sslSocket = (SSLSocket) sslServerSocket.accept();
        SSLParameters sslp = sslSocket.getSSLParameters();
        String[] serverAPs ={"one","two","three"};
        sslp.setApplicationProtocols(serverAPs);
        sslSocket.setSSLParameters(sslp);
        sslSocket.startHandshake();

        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol server side: \"" + ap +
"\"");

        InputStream sslIS = sslSocket.getInputStream();
        OutputStream sslOS = sslSocket.getOutputStream();
        sslIS.read();
        sslOS.write(85);
        sslOS.flush();

        sslSocket.close();
        sslServerSocket.close();
    }
}
```

With the custom `X509ExtendedKeyManager` in place, when `chooseServerAlias` is called during handshaking the `KeyManager` has the opportunity to examine the negotiated application protocol value. In the case of the example shown, this value is output to the console.

For example, when this code is run and a Java client sends a `ClientHello` with ALPN values `three` and `two`, the output will be:

```
Application Protocol server side: two
```

ALPN Related Classes and Methods

These classes and methods are used when working with Application Layer Protocol Negotiation (ALPN).

Classes and Methods to Use

`SSLEngine` and `SSLSocket` contain the same ALPN related methods and they have the same functionality.

Class	Method	Purpose
`SSLParameters`	`public String[] getApplicationProtocols();`	**Client-side and server-side**: use the method to return a `String` array containing each protocol set.

Class	Method	Purpose
`SSLParameters`	`public void setApplicationProtocols([] protocols);`	**Client-side**: use the method to set the protocols that can be chosen by the server. **Server-side**: use the method to set the protocols that the server can use. The String array should contain the protocols in order of preference.
`SSLEngine` `SSLSocket`	`public String getApplicationProtocol();`	**Client-side and server-side**: use the method *after* TLS protocol negotiation has completed to return a `String` containing the protocol that has been chosen for the connection.
`SSLEngine` `SSLSocket`	`public String getHandshakeApplicationProtocol();`	**Client-side and server-side**: use the method *during* handshaking to return a `String` containing the protocol that has been chosen for the connection. If this method is called before or after handshaking, it will return null. See Determining Negotiated ALPN Value during Handshaking for instructions on how to call this method.
`SSLEngine` `SSLSocket`	`public void setHandshakeApplicationProtocolSelector(BiFunction,String> selector)`	**Server-side**: use the method to register a callback function. The application protocol value can then be set in the callback based on any information available, for example the protocol or cipher suite. See Setting up Custom ALPN on the Server for instructions on how to use this method.

Troubleshooting JSSE

This section contains information for troubleshooting JSSE. First, it provides some common Configuration Problems and ways to solve them, and then it describes helpful Debugging Utilities.

Configuration Problems

Solutions to some common configuration problems.

SSLHandshakeException: No Available Authentication Scheme, Handshake Failure

Problem: The server throws this exception:

```
javax.net.ssl.SSLHandshakeException: No available authentication scheme
```

The client then receives a fatal alert:

```
javax.net.ssl.SSLHandshakeException: Received fatal alert:
handshake_failure
```

Cause: The server throws this `SSLHandshakeException` if TLSv1.3 is chosen as the protocol version and only DSA certificates are available in the server's keymanager. Verify this with the keytool command; change `testkeys.dsa` to the name of your keystore:

```
keytool -list -keystore testkeys.dsa -v

Enter keystore password:
Keystore type: PKCS12
Keystore provider: SUN

Your keystore contains 1 entry

Alias name: localhost
Creation date: Sep 19, 2018
Entry type: PrivateKeyEntry
Certificate chain length: 1
Certificate[1]: Owner: CN=localhost, OU=Widget, O=Ficticious, L=Sunnyvale,
ST=CA, C=US Issuer: CN=localhost, OU=Widget, O=Ficticious, L=Sunnyvale,
ST=CA, C=US

...deleted...

Signature algorithm name: SHA256withDSA
Subject Public Key Algorithm: 2048-bit DSA key

...deleted...
```

Solution: Update your certificates so that they contain RSA or EC public keys.

CertificateException While Handshaking

Problem: When negotiating an TLS/DTLS connection, the client or server throws a `CertificateException`.

Cause 1: This is generally caused by the remote side sending a certificate that is unknown to the local side.

Solution 1: The best way to debug this type of problem is to turn on debugging (see Debugging Utilities) and watch as certificates are loaded and when certificates are received via the network connection. Most likely, the received certificate is unknown to the trust mechanism because the wrong trust file was loaded.

Refer to the following sections:

- JSSE Classes and Interfaces
- TrustManager Interface
- KeyManager Interface

Cause 2: The system clock is not set correctly. In this case, the perceived time may be outside the validity period on one of the certificates, and unless the certificate can be replaced with a valid one from a truststore, the system must assume that the certificate is invalid, and therefore throw the exception.

Solution 2: Correct the system clock time.

Runtime Exception: SSL Service Not Available

Problem: When running a program that uses JSSE, an exception occurs indicating that an SSL service is not available. For example, an exception similar to one of the following is thrown:

```
Exception in thread "main" java.net.SocketException:
    no SSL Server Sockets

Exception in thread "main":
    SSL implementation not available
```

Cause: There was a problem with `SSLContext` initialization, for example, due to an incorrect password on a keystore or a corrupted keystore (a JDK vendor once shipped a keystore in an unknown format, and that caused this type of error).

Solution: Check initialization parameters. Ensure that any keystores specified are valid and that the passwords specified are correct. One way that you can check this is by trying to use `keytool` to examine the keystores and the relevant contents. See keytool in *Java Platform, Standard Edition Tools Reference*.

Runtime Exception: "No available certificate corresponding to the SSL cipher suites which are enabled"

Problem: When trying to run a simple SSL server program, the following exception is thrown:

```
Exception in thread "main" javax.net.ssl.SSLException:
    No available certificate corresponding to the SSL cipher suites
which are enabled...
```

Cause: Various cipher suites require certain types of key material. For example, if an RSA cipher suite is enabled, then an RSA `keyEntry` must be available in the keystore. If no such key is available, then this cipher suite cannot be used. This exception is thrown if there are no available key entries for all of the cipher suites enabled.

Solution: Create key entries for the various cipher suite types, or use an anonymous suite. Anonymous cipher suites are inherently dangerous because they are vulnerable to MITM (man-in-the-middle) attacks. See RFC 5246: The Transport Layer Security (TLS) Protocol, Version 1.2.

Refer to the following sections to learn how to pass the correct keystore and certificates:

- JSSE Classes and Interfaces
- Customizing the Default Keystores and Truststores, Store Types, and Store Passwords
- Additional Keystore Formats (PKCS12)

Runtime Exception: No Cipher Suites in Common

Problem 1: When handshaking, the client and/or server throw this exception.

Cause 1: Both sides of a TLS connection must agree on a common cipher suite. If the intersection of the client's cipher suite set with the server's cipher suite set is empty, then you will see this exception.

Solution 1: Configure the enabled cipher suites to include common cipher suites, and be sure to provide an appropriate `keyEntry` for asymmetric cipher suites. Also see Runtime Exception: "No available certificate corresponding to the SSL cipher suites which are enabled" in this section.)

Problem 2: When using Mozilla Firefox or Microsoft Internet Explorer to access files on a server that only has DSA-based certificates, a runtime exception occurs indicating that there are no cipher suites in common.

Cause 2: By default, `keyEntries` created with `keytool` use DSA public keys. If only DSA `keyEntries` exist in the keystore, then only DSA-based cipher suites can be used. By default, Firefox and Internet Explorer send only RSA-based cipher suites. Because the intersection of client and server cipher suite sets is empty, this exception is thrown.

Solution 2: To interact with Firefox or Internet Explorer, you should create certificates that use RSA-based keys. To do this, specify the `-keyalg` RSA option when using keytool. For example:

```
keytool -genkeypair -alias duke -keystore testkeys -keyalg rsa
```

Socket Disconnected After Sending ClientHello Message

Problem: A socket attempts to connect, sends a ClientHello message, and is immediately disconnected.

Cause: Some TLS servers will disconnect if a ClientHello message is received in a format they do not understand or with a protocol version number that they do not support.

Solution: Try adjusting the enabled protocols on the client side. This involves modifying or invoking some of the following system properties and methods:

- System property `https.protocols` for the `HttpsURLConnection` class
- System property `jdk.tls.client.protocols`
- `SSLContext.getInstance` method
- `SSLEngine.setEnabledProtocols` method
- `SSLSocket.setEnabledProtocols` method
- `SSLParameters.setProtocols` and `SSLEngine.setSSLParameters` methods
- `SSLParameters.setProtocols` and `SSLSocket.setSSLParameters` methods

For backwards compatibility, some TLS implementations (such as SunJSSE) can send TLS ClientHello messages encapsulated in the SSLv2 ClientHello format. The

SunJSSE provider supports this feature. If you want to use this feature, add the "SSLv2Hello" protocol to the enabled protocol list, if necessary. (See Protocols in the JDK Providers, which lists the protocols that are enabled by default for the SunJSSE provider.)

The TLS RFC standards require that implementations negotiate to the latest version both sides speak, but some non-conforming implementation simply hang up if presented with a version they don't understand. For example, some older server implementations that speak only SSLv3 will shutdown if TLSv1.2 is requested. In this situation, consider using a TLS version fallback scheme:

1. Fall back from TLSv1.2 to TLSv1.1 if the server does not understand TLSv1.2.
2. Fall back from TLSv1.1 to TLSv1.0 if the previous step does not work.

For example, if the enabled protocol list on the client is TLSv1, TLSv1.1, and TLSv1.2, a typical TLS version fallback scheme may look like:

1. Try to connect to server. If server rejects the TLS connection request immediately, go to step 2.
2. Try the version fallback scheme by removing the highest protocol version (for example, TLSv1.2 for the first failure) in the enabled protocol list.
3. Try to connect to the server again. If server rejects the connection, go to step 2 unless there is no version to which the server can fall back.
4. If the connection fails and SSLv2Hello is not on the enabled protocol list, restore the enable protocol list and enable SSLv2Hello. (For example, the enable protocol list should be SSLv2Hello, TLSv1, TLSv1.1, and TLSv1.2.) Start again from step 1.

> **Note:**
>
> A fallback to a previous version normally means security strength downgrading to a weaker protocol. It is not suggested to use a fallback scheme unless it is really necessary, and you clearly know that the server does not support a higher protocol version.

> **Note:**
>
> As part of disabling SSLv3, some servers have also disabled SSLv2Hello, which means communications with SSLv2Hello-active clients (JDK 6u95) will fail. Starting with JDK 7, SSLv2Hello default to disabled on clients, enabled on servers.

SunJSSE Cannot Find a JCA Provider That Supports a Required Algorithm and Causes a NoSuchAlgorithmException

Problem: A handshake is attempted and fails when it cannot find a required algorithm. Examples might include:

```
Exception in thread ...deleted...
    ...deleted...
```

```
    Caused by java.security.NoSuchAlgorithmException: Cannot find any
        provider supporting RSA/ECB/PKCS1Padding
```

or

```
Caused by java.security.NoSuchAlgorithmException: Cannot find any
    provider supporting AES/CBC/NoPadding
```

Cause: SunJSSE uses JCE for all its cryptographic algorithms. If the SunJCE provider has been deregistered from the `Provider` mechanism and an alternative implementation from JCE is not available, then this exception will be thrown.

Solution: Ensure that the SunJCE is available by checking that the provider is registered with the `Provider` interface. Try to run the following code in the context of your SSL connection:

```
import javax.crypto.*;

System.out.println("=====Where did you get AES=====");
Cipher c = Cipher.getInstance("AES/CBC/NoPadding");
System.out.println(c.getProvider());
```

Exception Thrown When Obtaining Application Resources from a Virtual Host Web Server that Requires an SNI Extension

Problem: If you receive an `Exception` when trying to obtain application resources from your web server over TLS, and your web server is implemented as a virtual host that requires a valid Server Name Indication (SNI) extension (such as Apache HTTP Server) to distinguish the virtual host, then the web server might not be configured correctly.

Cause: Because Java SE supports the SNI extension in the JSSE client, the requested host name of the virtual server is included in the first message sent from the client to the server during the TLS handshake. The server may deny the client's request for a connection if the requested host name (the server name indication) does not match the expected server name, which should be specified in the virtual host's configuration. This triggers an TLS handshake unrecognized name alert, which results in an `Exception` being thrown.

Solution: If the cause of the problem is `javax.net.ssl.SSLProtocolException: handshake alert: unrecognized_name`, it is likely that the virtual host configuration for SNI is incorrect. If you are using Apache HTTP Server, see Name-based Virtual Host Support about configuring virtual hosts. In particular, ensure that the `ServerName` directive is configured properly in a `<VirtualHost>` block.

See the following:

- SSL with Virtual Hosts Using SNI from Apache HTTP Server Wiki
- SSL/TLS Strong Encryption: FAQ from Apache HTTP Server Documentation
- RFC 3546, Transport Layer Security (TLS) Extensions
- Bug 7194590: SSL handshaking error caused by virtual server misconfiguration

IllegalArgumentException When RC4 Cipher Suites are Configured for DTLS

Problem: An `IllegalArgumentException` exception is thrown when RC4 cipher suite algorithm is specified in `SSLEngine.setEnabledCipherSuites(String[] suites)` method and the `SSLEngine` is a DTLS engine.

```
sslContext = SSLContext.getInstance("DTLS");

// Create the engine
SSLEngine engine = sslContext.createSSLengine(hostname, port);

String enabledSuites[] = { "SSL_RSA_WITH_RC4_128_SHA" };
engine.setEnabledCipherSuites(enabledSuites);
```

Cause: According to DTLS Version 1.0 and DTLS Version 1.2, RC4 cipher suites must not be used with DTLS.

Solution: Do not use RC4 based cipher suites for DTLS connections. See "JSSE Cipher Suite Names" in Java Security Standard Algorithm Names.

Debugging Utilities

The SunJSSE provider supports dynamic debug tracing. This is similar to the mechanism that debugs security library issues. The generic Java dynamic debug tracing support is accessed with the `java.security.debug` system property, whereas the JSSE-specific dynamic debug tracing support is accessed with the `javax.net.debug` system property.

> **Note:**
>
> Currently, the SunJSSE provider uses the debug utility. There is no guarantee that other providers use the debug utility. If other providers support the debug utility, then the implementation and output may be different. There is no guarantee the debug utility will continue to exist or be the same (for example, have the same options or output format) in future releases.

To view the options of the JSSE dynamic debug utility, use the following command-line option on the `java` command, where `MyApp` is an existing Java application:

```
java -Djavax.net.debug=help MyApp
```

> **Note:**
>
> - The `MyApp` application will not run after the debug help information is printed, as the help code causes the application to exit.
> - If you specify the value `help` with either dynamic debug utility when running a program that does not use any classes that the utility was designed to debug, you will not get the debugging options.

The current options are:

- `all`: Turn on all debugging
- `ssl`: Turn on SSL debugging

The following can be used with the `ssl` option to select what type of debug information to print:

- `defaultctx`: Print default SSL initialization
- `handshake`: Print each handshake message
- `keygen`: Print key generation data
- `keymanager`: Print key manager tracing
- `pluggability`: Print pluggability tracing
- `record`: Enable per-record tracing
- `respmgr`: Print status response manager tracing
- `session`: Print session activity
- `sessioncache`: Print session cache tracing
- `sslctx`: Print `SSLContext` tracing
- `trustmanager`: Print trust manager tracing

Messages generated from the `handshake` option can be widened with these options:

- `data`: Hex dump of each handshake message
- `verbose`: Verbose handshake message printing

Messages generated from the `record` option can be widened with these options:

- `plaintext`: Hex dump of record plaintext
- `packet`: Print raw SSL/TLS packets

To enable JSSE-specific dynamic debug tracing, set the value of the `javax.net.debug` system property (see How to Specify a java.security.Security Property) to either `all` or `ssl`. For the `ssl` option, to specify additional options, specify them after the `ssl` option. You do *not* have to have a separator between options, although a separator such as a colon (`:`) or a comma (`,`) helps readability. It does not matter what separators you use, and the ordering of the option keywords is also not important.

For an introduction to reading this debug information, see Debugging TLS Connections.

The following are examples of using the `javax.net.debug` system property:

- To view all debugging messages:

  ```
  java -Djavax.net.debug=all MyApp
  ```

- To view the hexadecimal dumps of each handshake message (the colons are optional):

  ```
  java -Djavax.net.debug=ssl:handshake:data MyApp
  ```

- To view the hexadecimal dumps of each handshake message, and to print trust manager tracing (the commas are optional):

  ```
  java -Djavax.net.debug=ssl,handshake,data,trustmanager MyApp
  ```

Debugging TLS Connections

Understanding TLS connection problems can sometimes be difficult, especially when it is not clear what messages are actually being sent and received. JSSE has a built-in debug facility and is activated by the system property `javax.net.debug`. To know more about `javax.net.debug` System property, see Debugging Utilities.

This section gives a brief overview of the debug output of the basic TLS 1.3 handshake. To know more about the TLS protocol, see RFC 8446: The Transport Layer Security (TLS) Protocol Version 1.3.

> **Note:**
>
> - Debug output information about all possible TLS handshake combinations and protocols is beyond the scope of this guide. Instead, refer to the relevant RFC for more detailed information about a particular version of TLS. See TLS and DTLS Protocols for a list of supported SSL/TLS/DTLS protocols and links to their respective RFCs.
> - The output is non-standard and may change from release to release.

This example uses the default JSSE X509KeyManager and X509TrustManager, which also prints debug information about the keys and trusted certificates used during a connection. It uses the `ClassFileServer` and `SSLSocketClientWithClientAuth` sample applications from JSSE Sample Code in the JDK 8 documentation. `ClassFileServer` is a simple HTTPS server that can require client authentication. `SSLSocketClientWithClientAuth` demonstrates how to use the `SSLSocket` class as a client to send an HTTP request and get a response from an HTTPS server. To make things simpler, both `ClassFileServer` and `SSLSocketClientWithClientAuth` are run from the same host.

Run ClassFileServer on localhost

The following command runs the `ClassFileServer` application on `localhost`, port 2002:

```
java \
  -Djavax.net.ssl.trustStore=/my_home_directory/jssesamples/samples/
samplecacerts \
  -Djavax.net.ssl.trustStorePassword=changeit \
  ClassFileServer 2002 \
  /my_home_directory/jssesamples/samples/ \
  TLS true
```

Run SSLSocketClientWithClientAuth on locahost

The following command runs the `SSLSocketClientWithClientAuth` application on `localhost`, port 2002. The application connects to the HTTPS server that you started with the previous command. It sends an HTTPS request to the server and receives the reply. Note that the command sets the value of the system property `javax.net.debug` to `all`, which turns on all debugging.

```
java -Djavax.net.debug=all -Djavax.net.ssl.trustStore=/my_home_directory/
jssesamples/samples/samplecacerts SSLSocketClientWithClientAuth localhost
2002 /index.html
```

Debug Output Format

Each line of the debug output contains the following information; each field is separated by a vertical bar (`|`):

- Logger name (`System.getLogger("javax.net.ssl")`)
- Debug level (`System.Logger.Level`)
- Thread ID (`Thread.currentThread().getId()`)
- Thread name (`Thread.currentThread().getName()`)
- Date and time
- Caller (location of the logging call)
- Message

Determine Client-Side and Server-Side Enabled Cipher Suites

The values of the system properties `jdk.tls.client.cipherSuites` and `jdk.tls.server.cipherSuites` are checked to determine the default enabled cipher suites; see Specifying Default Enabled Cipher Suites for more information about these system properties.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:46.990 EDT|
SSLContextImpl.java:427|System property jdk.tls.client.cipherSuites is set
to 'null'
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.026 EDT|
SSLContextImpl.java:427|System property jdk.tls.server.cipherSuites is set
```

```
to 'null'
...
```

The values of these system properties are null, so the default enabled cipher suites are those that the SunJSSE provider enables by default; see The SunJSSE Provider in JDK Providers Documentation.

The value of `jdk.tls.keyLimits` is checked to determine the limit of the amount of data an algorithm may encrypt with a specific set of keys; see Limiting Amount of Data Algorithms May Encrypt with a Set of Keys.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.124 EDT|SSLCipher.java:436|
jdk.net.keyLimits:  entry = AES/GCM/NoPadding KeyUpdate 2^37. AES/GCM/
NOPADDING:KEYUPDATE = 137438953472
...
```

The debug output lists unsupported and disabled cipher suites:

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.150 EDT|
SSLContextImpl.java:401|Ignore disabled cipher suite:
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
javax.net.ssl|ALL|01|main|2018-08-18 01:04:47.150 EDT|SSLContextImpl.java:
410|Ignore unsupported cipher suite: TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.152 EDT|
SSLContextImpl.java:401|Ignore disabled cipher suite:
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
...
```

Initialize X509KeyManager

The X509KeyManager is initialized. It discovers that there is one keyEntry in the supplied KeyStore for a subject called "duke". If this application wants to authenticate itself, then the X509KeyManager searches its list of keyEntries for an appropriate credential.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.224 EDT|
SunX509KeyManagerImpl.java:164|found key for : duke (
  "certificate" : {
    "version"            : "v1",
    "serial number"      : "3B 0A FA 66",
    "signature algorithm": "MD5withRSA",
    "issuer"             : "CN=Duke, OU=Java Software, O="Sun
Microsystems, Inc.", L=Cupertino, ST=CA, C=US",
    "not before"         : "2001-05-22 19:46:46.000 EDT",
    "not  after"         : "2011-05-22 19:46:46.000 EDT",
    "subject"            : "CN=Duke, OU=Java Software, O="Sun
Microsystems, Inc.", L=Cupertino, ST=CA, C=US",
    "subject public key" : "RSA"}
)
...
```

Initialize a TrustManager

A TrustManager is initialized and it finds in the truststore several certificates from various Certificate Authorities (CAs). It also finds a self-signed certificate with a distinguished name “localhost”. A server that presents valid credentials (certificates) that chain back to a trusted certificate in the truststore will itself be trusted.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.229 EDT|
TrustStoreManager.java:112|trustStore is: /my_home_directory/jssesamples/
samples/samplecacerts
trustStore type is: pkcs12
trustStore provider is:
the last modified time is: Tue Dec 11 06:43:38 EST 2012
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.229 EDT|
TrustStoreManager.java:311|Reload the trust store
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.264 EDT|
TrustStoreManager.java:318|Reload trust certs
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.266 EDT|
TrustStoreManager.java:323|Reloaded 32 trust certs
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.322 EDT|
X509TrustManagerImpl.java:79|adding as trusted certificates (
  "certificate" : {
    "version"            : "v1",
    "serial number"      : "00 9B 7E 06 49 A3 3E 62 B9 D5 EE 90 48 71 29
EF 57",
    "signature algorithm": "SHA1withRSA",
    "issuer"             : "CN=VeriSign Class 3 Public Primary
Certification Authority - G3, OU="(c) 1999 VeriSign, Inc. - For authorized
use only", OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US",
    "not before"         : "1999-09-30 20:00:00.000 EDT",
    "not  after"         : "2036-07-16 19:59:59.000 EDT",
    "subject"            : "CN=VeriSign Class 3 Public Primary
Certification Authority - G3, OU="(c) 1999 VeriSign, Inc. - For authorized
use only", OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US",
    "subject public key" : "RSA"},
  "certificate" : {
    "version"            : "v1",
    "serial number"      : "61 70 CB 49 8C 5F 98 45 29 E7 B0 A6 D9 50 5B
7A",
    "signature algorithm": "SHA1withRSA",
    "issuer"             : "CN=VeriSign Class 2 Public Primary
Certification Authority - G3, OU="(c) 1999 VeriSign, Inc. - For authorized
use only", OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US",
    "not before"         : "1999-09-30 20:00:00.000 EDT",
    "not  after"         : "2036-07-16 19:59:59.000 EDT",
    "subject"            : "CN=VeriSign Class 2 Public Primary
Certification Authority - G3, OU="(c) 1999 VeriSign, Inc. - For authorized
use only", OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US",
    "subject public key" : "RSA"},
...
  "certificate" : {
    "version"            : "v1",
    "serial number"      : "41 00 44 46",
    "signature algorithm": "MD5withRSA",
```

```
    "issuer"             : "CN=localhost, OU=Widget Development Group,
O="Ficticious Widgets, Inc.", L=Sunnyvale, ST=CA, C=US",
    "not before"         : "2004-07-22 18:48:38.000 EDT",
    "not  after"         : "2011-05-22 18:48:38.000 EDT",
    "subject"            : "CN=localhost, OU=Widget Development Group,
O="Ficticious Widgets, Inc.", L=Sunnyvale, ST=CA, C=US",
    "subject public key" : "RSA"},
...
```

Perform Additional Initialization

The example performs additional initialization code, then connects to the server.

```
javax.net.ssl|ALL|01|main|2018-08-18 01:04:47.326 EDT|SSLContextImpl.java:
115|trigger seeding of SecureRandom
javax.net.ssl|ALL|01|main|2018-08-18 01:04:47.524 EDT|SSLContextImpl.java:
119|done seeding of SecureRandom
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.562 EDT|
HandshakeContext.java:291|Ignore unsupported cipher suite:
TLS_AES_128_GCM_SHA256 for TLS12
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.563 EDT|
HandshakeContext.java:291|Ignore unsupported cipher suite:
TLS_AES_256_GCM_SHA384 for TLS12
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.577 EDT|
HandshakeContext.java:291|Ignore unsupported cipher suite:
TLS_AES_128_GCM_SHA256 for TLS11
...
```

The debug output also notifies you of disabled, unsupported, or unavailable extensions and signature algorithms:

```
javax.net.ssl|WARNING|01|main|2018-08-18 01:04:47.695 EDT|
ServerNameExtension.java:255|Unable to indicate server name
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.695 EDT|SSLExtensions.java:
235|Ignore, context unavailable extension: server_name
javax.net.ssl|WARNING|01|main|2018-08-18 01:04:47.703 EDT|
SignatureScheme.java:282|Signature algorithm, ed25519, is not supported by
the underlying providers
javax.net.ssl|WARNING|01|main|2018-08-18 01:04:47.704 EDT|
SignatureScheme.java:282|Signature algorithm, ed448, is not supported by
the underlying providers
javax.net.ssl|ALL|01|main|2018-08-18 01:04:47.724 EDT|SignatureScheme.java:
358|Ignore disabled signature sheme: rsa_md5
javax.net.ssl|INFO|01|main|2018-08-18 01:04:47.724 EDT|AlpnExtension.java:
161|No available application protocols
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.724 EDT|SSLExtensions.java:
235|Ignore, context unavailable extension:
application_layer_protocol_negotiation
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.725 EDT|SSLExtensions.java:
235|Ignore, context unavailable extension: cookie
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.763 EDT|SSLExtensions.java:
235|Ignore, context unavailable extension: renegotiation_info
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.763 EDT|
PreSharedKeyExtension.java:606|No session to resume.
```

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.763 EDT|SSLExtensions.java:
235|Ignore, context unavailable extension: pre_shared_key
...
```

Send ClientHello Message

The client sends a ClientHello message to the server. This message specifies the following:

- Client version: For TLS 1.3, this has a fixed value, TLSv1.2; TLS 1.3 uses the extension supported_versions and not this field to negotiate protocol version
- Random: A random value used to initialize the cryptographic algorithms
- Session ID: Previous versions of TLS use this ID to support a session resumption feature
- Cipher Suites: The list of cipher suites that the client requests; depending on the enabled cipher suites, there may be a broad mix of cipher suite names, some of which are only for TLSv1.3 while others are for TLSv1.2 and earlier
- Compression methods: For TLS 1.3, this must have the value 0
- Extensions:
  - status_request: The client requests OCSP; see Client-Driven OCSP and OCSP Stapling
  - supported_groups: Lists the named groups that the client supports for key exchange. These named groups include elliptic curve groups (ECDHE) and finite field groups (DHE). The ClientHello message must include this message if it's using ECDHE or DHE key exchange.
  - ec_point_formats: Lists the elliptical curve point formats that the client can parse; in this example, the client can parse uncompressed point formats only. Other formats include compressed and ansiX962_compressed_prime.
  - signature algorithms: Lists which signature algorithms may be used in CertificateVerify messages
  - signature_algorithms_cert: Lists which signature algorithms may be used in digital signatures
  - status_request_v2: Enables clients to specify and support several certificate status methods. Note that this extension is deprecated for TLS 1.3.
  - extended_master_secret: In TLS 1.2 and earlier, this extension requests that both sides digest larger parts of the handshake transcript into the master secret than the original version of the protocol did; see RFC 7627. The extension is included in TLS 1.3 handshakes in case a TLS 1.2 handshake is negotiated.
  - supported_versions: Lists which versions of TLS the client supports. In particular, if the client requests TLS 1.3, then the client version field has the value TLSv1.2 and this extension contains the value TLSv1.3; if the client requests TLS 1.2, then the client version field has the value TLSv1.2 and this extension either doesn't exist or contains the value TLSv1.2 but not the value TLSv1.3.
  - psk_key_exchange_modes: Lists which key exchange modes that may be used with pre-shared keys (PSKs); in this example, the client supports PSK

with (EC)DHE key establishment (psk_dhe_ke). In this mode, the client and server must supply values for the key_share extension.

– key_share: Lists cryptographic parameters for key exchange. It contains a field named client_shares that contains this list. Each item of this list contains two fields: group and key_exchange. This example contains key exchange information for the elliptical curve secp256r1.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.769 EDT|ClientHello.java:
633|Produced ClientHello handshake message (
"ClientHello": {
  "client version"      : "TLSv1.2",
  "random"              : "64 CF 68 A1 CF AB B1 6F 43 F6 DE 1B 49 49 DE 5A
42 9A 71 DD CB 9A E3 9F 32 00 E8 87 7A 00 DA C6",
  "session id"          : "02 0D BE 1B A4 5F F2 E8 B6 31 9D A4 EF F3 22 84
C3 58 0B 5C C0 57 0F A5 6D 8A 83 EB DC DA B1 B6",
  "cipher suites"       : "[TLS_AES_128_GCM_SHA256(0x1301),
TLS_AES_256_GCM_SHA384(0x1302),
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384(0xC02C),
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256(0xC02B),
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384(0xC030),
TLS_RSA_WITH_AES_256_GCM_SHA384(0x009D),
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384(0xC02E),
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384(0xC032),
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384(0x009F),
TLS_DHE_DSS_WITH_AES_256_GCM_SHA384(0x00A3),
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256(0xC02F),
TLS_RSA_WITH_AES_128_GCM_SHA256(0x009C),
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256(0xC02D),
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256(0xC031),
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256(0x009E),
TLS_DHE_DSS_WITH_AES_128_GCM_SHA256(0x00A2),
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384(0xC024),
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384(0xC028),
TLS_RSA_WITH_AES_256_CBC_SHA256(0x003D),
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384(0xC026),
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384(0xC02A),
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256(0x006B),
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256(0x006A),
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA(0xC00A),
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA(0xC014),
TLS_RSA_WITH_AES_256_CBC_SHA(0x0035),
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA(0xC005),
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA(0xC00F),
TLS_DHE_RSA_WITH_AES_256_CBC_SHA(0x0039),
TLS_DHE_DSS_WITH_AES_256_CBC_SHA(0x0038),
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256(0xC023),
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256(0xC027),
TLS_RSA_WITH_AES_128_CBC_SHA256(0x003C),
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256(0xC025),
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256(0xC029),
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256(0x0067),
TLS_DHE_DSS_WITH_AES_128_CBC_SHA256(0x0040),
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA(0xC009),
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA(0xC013),
```

```
TLS_RSA_WITH_AES_128_CBC_SHA(0x002F),
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA(0xC004),
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA(0xC00E),
TLS_DHE_RSA_WITH_AES_128_CBC_SHA(0x0033),
TLS_DHE_DSS_WITH_AES_128_CBC_SHA(0x0032),
TLS_EMPTY_RENEGOTIATION_INFO_SCSV(0x00FF)]",
  "compression methods" : "00",
  "extensions"          : [
    "status_request (5)": {
      "certificate status type": ocsp
      "OCSP status request": {
        "responder_id": <empty>
        "request extensions": {
          <empty>
        }
      }
    },
    "supported_groups (10)": {
      "versions": [secp256r1, secp384r1, secp521r1, sect283k1, sect283r1,
sect409k1, sect409r1, sect571k1, sect571r1, secp256k1, ffdhe2048,
ffdhe3072, ffdhe4096, ffdhe6144, ffdhe8192]
    },
    "ec_point_formats (11)": {
      "formats": [uncompressed]
    },
    "signature_algorithms (13)": {
      "signature schemes": [ecdsa_secp256r1_sha256,
ecdsa_secp384r1_sha384, ecdsa_secp512r1_sha512, rsa_pss_rsae_sha256,
rsa_pss_rsae_sha384, rsa_pss_rsae_sha512, rsa_pss_pss_sha256,
rsa_pss_pss_sha384, rsa_pss_pss_sha512, rsa_pkcs1_sha256,
rsa_pkcs1_sha384, rsa_pkcs1_sha512, dsa_sha256, ecdsa_sha1,
rsa_pkcs1_sha1, dsa_sha1]
    },
    "signature_algorithms_cert (50)": {
      "signature schemes": [ecdsa_secp256r1_sha256,
ecdsa_secp384r1_sha384, ecdsa_secp512r1_sha512, rsa_pss_rsae_sha256,
rsa_pss_rsae_sha384, rsa_pss_rsae_sha512, rsa_pss_pss_sha256,
rsa_pss_pss_sha384, rsa_pss_pss_sha512, rsa_pkcs1_sha256,
rsa_pkcs1_sha384, rsa_pkcs1_sha512, dsa_sha256, ecdsa_sha1,
rsa_pkcs1_sha1, dsa_sha1]
    },
    "status_request_v2 (17)": {
      "cert status request": {
        "certificate status type": ocsp_multi
        "OCSP status request": {
          "responder_id": <empty>
          "request extensions": {
            <empty>
          }
        }
      }
    },
    "extended_master_secret (23)": {
      <empty>
    },
```

```
    "supported_versions (43)": {
      "versions": [TLSv1.3, TLSv1.2, TLSv1.1, TLSv1]
    },
    "psk_key_exchange_modes (45)": {
      "ke_modes": [psk_dhe_ke]
    },
    "key_share (51)": {
      "client_shares": [
        {
          "named group": secp256r1
          "key_exchange": {
            0000: 04 1F 80 50 D9 C6 03 45   7B 59 0F A7 B6 9E AE
39  ...P...E.Y.....9
            0010: 37 BE B0 5B 09 D8 91 37   72 5D 2B 8E 01 0A 84 56  7..
[...7r]+....V
            0020: 99 0D 37 49 8F 92 61 A9   D6 54 E1 3B EE D1 E8 D2  ..
7I..a..T.;....
            0030: 92 22 F9 17 CE A7 F8 51   47 C9 1E 5C D6 59 0F
4F  ."....QG..\.Y.O
            0040: 55
          }
        },
      ]
    }
  ]
}
)
...
```

Show Actual Data Sent and Read

The debug output shows the actual data sent to the raw output object (in this case, an OutputStream):

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.770 EDT|
SSLSocketOutputRecord.java:217|WRITE: TLS13 handshake, length = 405
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.774 EDT|
SSLSocketOutputRecord.java:231|Raw write (
  0000: 16 03 03 01 95 01 00 01   91 03 03 64 CF 68 A1 CF  ...........d.h..
  0010: AB B1 6F 43 F6 DE 1B 49   49 DE 5A 42 9A 71 DD CB  ..oC...II.ZB.q..
  0020: 9A E3 9F 32 00 E8 87 7A   00 DA C6 20 02 0D BE 1B  ...2...z... ....
...
```

Then, the debug output shows the raw data read from the input device (InputStream) before any processing has been performed:

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.862 EDT|
SSLSocketInputRecord.java:215|READ: TLSv1.2 handshake, length = 155
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.863 EDT|
SSLSocketInputRecord.java:474|Raw read (
  0000: 02 00 00 97 03 03 66 24   0F F6 6D 4A 0C 5A A1 23  ......f$..mJ.Z.#
  0010: F6 5D 4B 87 B1 6E AC 13   BB 4D C1 A4 0F F0 2C EF  .]K..n...M....,.
  0020: D7 4F 03 11 19 B1 20 02   0D BE 1B A4 5F F2 E8 B6  .O.... ....._...
...
```

Whenever the client sends or reads a message, the debug output shows the raw data sent or read and how any messages (and their extensions) have been processed. The following sections omit these parts of the debug output.

Read ServerHello Message

At this point, TLS 1.3 has been negotiated. The server selects the TLS version and replies using a combination of the server version and the supported_versions extension. In this case, a TLSv1.3 protocol was indicated.

The ServerHello message specifies the following:

- Server version: For TLS 1.3, this must have the value TLSv1.2; TLS 1.3 uses the extension supported_versions and not this field to indicate the negotiated protocol version
- Random: Also used to initialize the cryptographic algorithms
- Session ID: For TLS 1.3, this has the same value as the corresponding field of the ClientHello message
- Cipher suite: The selected cipher suite; in this example, it is TLS_AES_128_GCM_SHA256
- Compression methods: For TLS 1.3, this must have the value 0
- Extensions
  - supported_versions: Specifies which TLS version the server uses. Note that for TLS 1.3, the server must use the value of the ClientHello message's supported_versions extension for version negotiation instead of the value of the client version field.
  - key_share: The named group and key values for a ECDHE key exchange

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.863 EDT|
SSLSocketInputRecord.java:251|READ: TLSv1.2 handshake, length = 155
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.867 EDT|ServerHello.java:
862|Consuming ServerHello handshake message (
"ServerHello": {
  "server version"      : "TLSv1.2",
  "random"              : "66 24 0F F6 6D 4A 0C 5A A1 23 F6 5D 4B 87 B1 6E
AC 13 BB 4D C1 A4 0F F0 2C EF D7 4F 03 11 19 B1",
  "session id"          : "02 0D BE 1B A4 5F F2 E8 B6 31 9D A4 EF F3 22 84
C3 58 0B 5C C0 57 0F A5 6D 8A 83 EB DC DA B1 B6",
  "cipher suite"        : "TLS_AES_128_GCM_SHA256(0x1301)",
  "compression methods" : "00",
  "extensions"          : [
    "supported_versions (43)": {
      "selected version": [TLSv1.3]
    },
    "key_share (51)": {
      "server_share": {
        "named group": secp256r1
        "key_exchange": {
          0000: 04 DE 5B 20 0E FD EB 6E   DA 70 C2 D0 FA 0D 4C 53  ..
[ ...n.p....LS
          0010: 6D E1 9E 67 77 65 36 AF   B5 EB E6 D2 88 92 9B EE
m..gwe6.........
```

```
          0020: E4 97 A3 B3 C1 FB D8 29   3B 92 87 D2 B3 9E 3D
AA  .......);.....=.
          0030: 14 99 1E 84 8F C2 E9 E3   E1 AC 9A 12 95 F0 26
B5  .............&.
          0040: 88
        }
      },
    }
  ]
}
)
...
```

The session is initialized:

```
javax.net.ssl|ALL|01|main|2018-08-18 01:04:47.873 EDT|SSLSessionImpl.java:
203|Session initialized:  Session(1534568687873|TLS_AES_128_GCM_SHA256)
...
```

Read EncryptedExtensions Message

At this point in the handshake, enough cryptographic information has been exchanged, and the remainder of the handshake will be performed encrypted.

The EncryptedExtensions message contains responses to ClientHello extensions that are not required to determine the cryptographic parameters, other than those that are specific to individual certificates; in this example, it returns the list of named groups that the client supports for key exchange.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.942 EDT|
EncryptedExtensions.java:171|Consuming EncryptedExtensions handshake
message (
"EncryptedExtensions": [
  "supported_groups (10)": {
    "versions": [secp256r1, secp384r1, secp521r1, sect283k1, sect283r1,
sect409k1, sect409r1, sect571k1, sect571r1, secp256k1, ffdhe2048,
ffdhe3072, ffdhe4096, ffdhe6144, ffdhe8192]
  }
]
)
...
```

Read Server's CertificateRequest Message

The server sends the CertificateRequest message if certificate-based client authentication is desired. This message contains the desired parameters for that certificate. It specifies the following:

- certificate_request_context: A string that identifies the certificate request; the value of this field is of zero length unless it's being used for post-handshake authentication
- Extensions: The following two extensions indicate which signature algorithms may be used in digital signatures:

- signature_algorithms: Originally appearing in TLS 1.2, applies to signatures in CertificateVerify messages
- signature_algorithms_cert: Applies to signatures in certificates

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.947 EDT|
CertificateRequest.java:864|Consuming CertificateRequest handshake message
(
"CertificateRequest": {
  "certificate_request_context": "",
  "extensions": [
    "signature_algorithms (13)": {
      "signature schemes": [ecdsa_secp256r1_sha256,
ecdsa_secp384r1_sha384, ecdsa_secp512r1_sha512, rsa_pss_rsae_sha256,
rsa_pss_rsae_sha384, rsa_pss_rsae_sha512, rsa_pss_pss_sha256,
rsa_pss_pss_sha384, rsa_pss_pss_sha512, rsa_pkcs1_sha256,
rsa_pkcs1_sha384, rsa_pkcs1_sha512, dsa_sha256, ecdsa_sha1,
rsa_pkcs1_sha1, dsa_sha1]
    },
    "signature_algorithms_cert (50)": {
      "signature schemes": [ecdsa_secp256r1_sha256,
ecdsa_secp384r1_sha384, ecdsa_secp512r1_sha512, rsa_pss_rsae_sha256,
rsa_pss_rsae_sha384, rsa_pss_rsae_sha512, rsa_pss_pss_sha256,
rsa_pss_pss_sha384, rsa_pss_pss_sha512, rsa_pkcs1_sha256,
rsa_pkcs1_sha384, rsa_pkcs1_sha512, dsa_sha256, ecdsa_sha1,
rsa_pkcs1_sha1, dsa_sha1]
    }
  ]
}
)
...
```

Read Server's Certificate Message

The Certificate message contains the authentication certificate and any other supporting certificates in the certificate chain. It specifies the following:

- certificate_request_context: For server authentication, this field is empty
- certificate_list: Contains a certificate chain signed by a signature algorithm advertised by the client. However, in this example, a self-signed certificate (a certificate whose subject and issue name are identical) was received. This same self-signed certificate was discovered earlier during initialization, so it will be trusted when the TrustManager is actually called to verify the received certificate.

There are many different ways of establishing trust, so if the default X509TrustManager is not doing the types of trust management you need, you can supply your own X509TrustManager to SSLContext.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:47.964 EDT|
CertificateMessage.java:1148|Consuming server Certificate handshake
message (
"Certificate": {
  "certificate_request_context": "",
  "certificate_list": [
  {
```

```
    "certificate" : {
      "version"            : "v1",
      "serial number"      : "41 00 44 46",
      "signature algorithm": "MD5withRSA",
      "issuer"             : "CN=localhost, OU=Widget Development Group,
O="Ficticious Widgets, Inc.", L=Sunnyvale, ST=CA, C=US",
      "not before"         : "2004-07-22 18:48:38.000 EDT",
      "not  after"         : "2011-05-22 18:48:38.000 EDT",
      "subject"            : "CN=localhost, OU=Widget Development Group,
O="Ficticious Widgets, Inc.", L=Sunnyvale, ST=CA, C=US",
      "subject public key" : "RSA"}
    "extensions": {
      <no extension>
    }
  },
]
}
)
...
```

The client recognizes this certificate and can trust it.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.165 EDT|
X509TrustManagerImpl.java:242|Found trusted certificate (
  "certificate" : {
    "version"            : "v1",
    "serial number"      : "41 00 44 46",
    "signature algorithm": "MD5withRSA",
    "issuer"             : "CN=localhost, OU=Widget Development Group,
O="Ficticious Widgets, Inc.", L=Sunnyvale, ST=CA, C=US",
    "not before"         : "2004-07-22 18:48:38.000 EDT",
    "not  after"         : "2011-05-22 18:48:38.000 EDT",
    "subject"            : "CN=localhost, OU=Widget Development Group,
O="Ficticious Widgets, Inc.", L=Sunnyvale, ST=CA, C=US",
    "subject public key" : "RSA"}
)
...
```

Read Server's CertificateVerify Message

The certificate sent by the server is verified by the CertificateVerify message. The message is used to provide explicit proof that the server has the private key corresponding to its certificate. This message specifies the following:

- Signature algorithm: The signature algorithm used; in this example, it is rsa_pss_rsae_sha256.
- Signature: The signature over the entire handshake using the private key corresponding to the public key in the Certificate message

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.194 EDT|
CertificateVerify.java:1128|Consuming CertificateVerify handshake message (
"CertificateVerify": {
  "signature algorithm": rsa_pss_rsae_sha256
  "signature": {
```

```
    0000: 0F 25 DD 62 03 6B 8C 8F   22 C7 8D 46 A2 A6 45 39  .
%.b.k.."..F..E9
    0010: 08 8D 51 1E 48 52 66 A4   F8 28 D3 FD 18 93 70 C6  ..Q.HRf..
(....p.
    0020: 32 74 C1 CC 0A C4 60 41   50 AF 7C DA 0C DB 92 F9
2t....`AP.......
    0030: 14 CB EF 15 7F 3E 52 16   F7 CC 8A 7C C9 1F 42
CA  .....>R.......B.
    0040: 90 8D FA B7 F2 3A 46 7E   F7 9F 43 CE C6 AA 15
59  .....:F...C....Y
    0050: EE AD 34 10 FF B7 BC FD   A2 F7 F3 1A FA 7F 26 61  ..
4...........&a
    0060: 80 2B 50 3A 8A 9E 5C 0E   4C A6 24 DA E6 3D 71 FA  .+P:..
\.L.$..=q.
    0070: AE 78 79 D2 DA 36 DE C1   A6 BC 18 46 04 CE 03 4E  .xy..
6.....F...N
  }
}
)
...
```

Read Server's Finished Message

The server sends a Finished message. This message contains a Message Authentication Code (MAC) over the entire handshake.

```
javax.net.ssl|DEBUG|01|main|2018-08-17 01:56:26.764 EDT|Finished.java:860|
Consuming server Finished handshake message (
"Finished": {
  "verify data": {
    0000: CA 7B 74 A6 79 36 ED 62   A7 0E 14 9D 9F D0 4A
0F  ..t.y6.b......J.
    0010: 02 4C 78 BB E2 89 A2 C6   E8 BD 28 CA E7 D9 DB 68  .Lx.......
(....h
  }'}
)
...
```

Send Certificate Message

The client sends a Certificate message because the server requested client authentication through a CertificateRequest message. The certificate message specifies similar information as the server's Certificate message. The client needs to send credentials back to the sever, so its X509KeyManager is consulted. The client looks for a match between the list of accepted issuers and the certificates that are in the KeyStore. In this case, there is a match: the client has the credentials for "duke". It's now up to the server's X509TrustManager to decide whether to accept these credentials.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.222 EDT|
CertificateMessage.java:1116|Produced client Certificate message (
"Certificate": {
  "certificate_request_context": "",
  "certificate_list": [
```

```
  {
    "certificate" : {
      "version"            : "v1",
      "serial number"      : "3B 0A FA 66",
      "signature algorithm": "MD5withRSA",
      "issuer"             : "CN=Duke, OU=Java Software, O="Sun
Microsystems, Inc.", L=Cupertino, ST=CA, C=US",
      "not before"         : "2001-05-22 19:46:46.000 EDT",
      "not  after"         : "2011-05-22 19:46:46.000 EDT",
      "subject"            : "CN=Duke, OU=Java Software, O="Sun
Microsystems, Inc.", L=Cupertino, ST=CA, C=US",
      "subject public key" : "RSA"}
    "extensions": {
      <no extension>
    }
  },
]
}
)
...
```

Send CertificateVerify Message

As with the CertificateVerify message sent by the server, the certificate sent by the client is verified by the CertificateVerify message. The message is used to provide explicit proof that the client has the private key corresponding to its certificate.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.268 EDT|
CertificateVerify.java:1097|Produced client CertificateVerify handshake
message (
"CertificateVerify": {
  "signature algorithm": rsa_pss_rsae_sha256
  "signature": {
    0000: 91 C2 F7 5D 8D 90 B4 82   E4 BA C6 23 08 E2 B4
DD  ...].......#....
    0010: 8D 95 8F 9F 31 4F 26 F3   97 3B FB 5B 10 4D AE F6  ....1O&..;.
[.M..
    0020: 71 78 FB 7B 3A 4F F6 1B   BF D2 E3 FB BE 53 F6 70
qx..:O.......S.p
    0030: 7E 73 83 F4 9A 5E 08 19   63 C1 97 4C 10 B1 C7
3F  .s...^..c..L...?
    0040: 4A 7D EF 4A 30 44 15 9F   D0 F2 8B C4 D1 45 69 B1
J..J0D.......Ei.
    0050: D9 DB 45 83 C4 11 91 B3   81 5E 69 F4 5C 2A CF 69  ..E......^i.
\*.i
    0060: D3 A6 7E 75 B4 C9 30 FB   5B AC BA 9F A3 C5 0C FD  ...u..0.
[.......
    0070: 9A 62 A4 DA 5A 80 6B 72   CD F5 A5 53 AD 14 74
1C  .b..Z.kr...S..t.
  }
}
)
```

Send Finished Message

The client then sends its Finished message to the sever:

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.271 EDT|Finished.java:658|
Produced client Finished handshake message (
"Finished": {
  "verify data": {
    0000: 93 04 B5 23 8F 48 3A CF   4A 85 35 9E 5F E0 1D 4C  ...#.H:.J.
5._..L
    0010: 9C 65 06 D4 E8 B4 ED 8F   01 6B 1E A2 DD 18 BD
78  .e.......k.....x
  }'}
)
...
```

The client and server have verified the Finished messages that they have received from their peers. Both sides may now send and receive application data over the connection.

Exchange Application Data, Client Sends GET Command

The server and client are ready to exchange application data. The client sends a "GET /index.html HTTP1.0" command.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.375 EDT|SSLCipher.java:
2019|Plaintext before ENCRYPTION (
  0000: 47 45 54 20 2F 69 6E 64   65 78 2E 68 74 6D 6C 20  GET /index.html
  0010: 48 54 54 50 2F 31 2E 30   0D 0A 0D 0A 17 00 00 00  HTTP/1.0........
  0020: 00 00 00 00 00 00 00 00   00 00 00 00 00           .............
)
...
```

Note that data over the wire is encrypted:

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.385 EDT|
SSLSocketOutputRecord.java:295|Raw write (
  0000: 17 03 03 00 3D 90 BF D1   81 E6 A3 E7 DA 50 A9 8B  ....=........P..
  0010: 18 F5 4B 30 AE 59 41 81   25 C4 9E 3E 70 29 5D C6  ..K0.YA.%..>p)].
  0020: 64 49 0B 4A 0E 93 E3 8F   DC 42 BA B5 21 42 38 88  dI.J.....B..!B8.
  0030: 62 4D 0C 86 FE 9A 8C B9   95 EF 89 93 61 3C 13 69  bM.........a<.i
  0040: 6C 45                                              lE
)
...
```

Read NewSessionTicket Message

After the server receives the client's Finished message, it can send a NewSessionTicket message anytime, which contains a PSK ticket that the client can use for speeding up future handshakes.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.517 EDT|
NewSessionTicket.java:330|Consuming NewSessionTicket message (
```

```
"NewSessionTicket": {
  "ticket_lifetime"      : "86,400",
  "ticket_age_add"       : "<omitted>",
  "ticket_nonce"         : "01",
  "ticket"               : "A5 30 8C B6 AD 95 79 E8 2A D1 95 C0 F0 2F 6F
AA 9E 97 58 AA 3D 19 82 2D 2C 47 C0 ED BF 64 48 AB",
  "extensions"           : [
    <no extension>
  ]
}
)
```

A duplicate SSLSession is created with the newly generated PSK information attached.

```
javax.net.ssl|ALL|01|main|2018-08-18 01:04:48.517 EDT|SSLSessionImpl.java:
203|Session initialized:  Session(1534568687873|TLS_AES_128_GCM_SHA256)
...
```

Exchange Application Data, Server Sends HTTPS Header and Data

The client receives application data from the server, first the HTTPS header, then the actual data.

```
javax.net.ssl|ALL|01|main|2018-08-18 01:04:48.517 EDT|SSLSessionImpl.java:
203|Session initialized:  Session(1534568687873|TLS_AES_128_GCM_SHA256)
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.617 EDT|
SSLSocketInputRecord.java:474|Raw read (
  0000: 17 03 03 00 63                                     ....c
)
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.618 EDT|
SSLSocketInputRecord.java:215|READ: TLSv1.2 application_data, length = 99
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.618 EDT|
SSLSocketInputRecord.java:474|Raw read (
  0000: 65 87 0E 1E 78 F7 AC C4   F7 C6 4D 55 91 6F 72 CC  e...x.....MU.or.
  0010: 18 2D 74 C3 B6 7B 2A F9   EB 2B F4 A8 C7 FD 09 FA  .-t...*..+......
  0020: 7E 36 9D F7 88 E7 44 DD   60 AF EB B0 F8 CF E1 64  .6....D.`......d
  0030: 0D 9B F4 B0 24 C2 BC B1   BF F7 F2 B6 CB E4 2E 39  ....$..........9
  0040: 78 B8 73 09 91 65 7A 0F   4C 49 DE 9A 7F 7B 42 86  x.s..ez.LI....B.
  0050: CA 33 87 DB 0D B2 E5 61   3C 70 6F F9 6A 15 A9 74  .3.....a<po.j..t
  0060: 64 E0 B0                                           d..
)
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.619 EDT|
SSLSocketInputRecord.java:251|READ: TLSv1.2 application_data, length = 99
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.621 EDT|SSLCipher.java:
1914|Plaintext after DECRYPTION (
  0000: 48 54 54 50 2F 31 2E 30   20 32 30 30 20 4F 4B 0D  HTTP/1.0 200 OK.
  0010: 0A 43 6F 6E 74 65 6E 74   2D 4C 65 6E 67 74 68 3A  .Content-Length:
  0020: 20 32 35 37 37 0D 0A 43   6F 6E 74 65 6E 74 2D 54   2577..Content-T
  0030: 79 70 65 3A 20 74 65 78   74 2F 68 74 6D 6C 0D 0A  ype: text/html..
  0040: 0D 0A                                              ..
)
...
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.626 EDT|
```

```
SSLSocketInputRecord.java:215|READ: TLSv1.2 application_data, length = 2610
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.628 EDT|
SSLSocketInputRecord.java:474|Raw read (
  0000: 69 8D F9 A3 E9 25 09 87   F0 E0 A1 63 12 9D 81 DF  i....%.....c....
  0010: 42 FC FA 7A 03 74 FD D5   ED 47 6C 5F 61 F2 BB 39  B..z.t...Gl_a..9
  0020: CF 64 0B B2 10 14 24 99   A3 66 8B D2 13 C9 66 FD  .d....$..f....f.
...
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.642 EDT|
SSLSocketInputRecord.java:251|READ: TLSv1.2 application_data, length = 2610
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.647 EDT|SSLCipher.java:
1914|Plaintext after DECRYPTION (
  0000: 3C 21 44 4F 43 54 59 50   45 20 68 74 6D 6C 20 50  <!DOCTYPE html P
  0010: 55 42 4C 49 43 20 22 2D   2F 2F 57 33 43 2F 2F 44  UBLIC "-//W3C//D
  0020: 54 44 20 58 48 54 4D 4C   20 31 2E 30 20 54 72 61  TD XHTML 1.0 Tra
  0030: 6E 73 69 74 69 6F 6E 61   6C 2F 2F 45 4E 22 0A 20  nsitional//EN".
...
```

Read Server's Alert Message

The server sends a close_notify alert, which notifies the client that it won't send anymore messages on this connection.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.658 EDT|Alert.java:232|
Received alert message (
"Alert": {
  "level"      : "warning",
  "description": "close_notify"
}
)
```

Close the Connection

The server closes the socket and then the TLS connection.

```
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.661 EDT|SSLSocketImpl.java:
1161|close the underlying socket
javax.net.ssl|DEBUG|01|main|2018-08-18 01:04:48.661 EDT|SSLSocketImpl.java:
921|close the ssl connection (passive)
javax.net.ssl|ALL|01|main|2018-08-18 01:04:48.661 EDT|SSLSocketImpl.java:
658|Closing input stream
javax.net.ssl|ALL|01|main|2018-08-18 01:04:48.661 EDT|SSLSocketImpl.java:
728|Closing output stream
```

Compatibility Risks and Known Issues

Enhancements to JSSE may introduce compatibility problems and other known issues, which are described in this section.

TLS 1.3 Not Directly Compatible with Previous Versions

TLS 1.3 is not directly compatible with previous versions. Although TLS 1.3 can be implemented with a backward-compatibility mode, there are still several compatibility risks to consider when upgrading to TLS 1.3:

- TLS 1.3 uses a half-close policy, while TLS 1.2 and earlier use a duplex-close policy. For applications that depend on the duplex-close policy, there may be compatibility issues when upgrading to TLS 1.3.
- The signature_algorithms_cert extension requires that pre-defined signature algorithms are used for certificate authentication. In practice, however, an application may use unsupported signature algorithms.
- The DSA signature algorithm is not supported in TLS 1.3. If a server is configured to only use DSA certificates, it cannot negotiate a TLS 1.3 connection.
- The supported cipher suites for TLS 1.3 are not the same as TLS 1.2 and earlier. If an application hardcodes cipher suites that are no longer supported, it may not be able to use TLS 1.3 without modifications to its code, for example TLS_AES_128_GCM_SHA256 (1.3 and later) versus TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (1.2 and earlier).
- The TLS 1.3 session resumption and key update behaviors are different from TLS 1.2 and earlier. The compatibility impact should be minimal, but it could be a risk if an application depends on the handshake details of the TLS protocols.

Code Examples

The following code examples are included in this section:

Topics

- Converting an Unsecure Socket to a Secure Socket
- Running the JSSE Sample Code
- Creating a Keystore to Use with JSSE
- Using the Server Name Indication (SNI) Extension

Converting an Unsecure Socket to a Secure Socket

Code examples that illustrate how to use JSSE to convert an unsecure socket connection to a secure socket connection. The code samples are excerpted from the book Java SE 6 Network Security by Marco Pistoia, et. al.

Example 8-26 shows sample code that can be used to set up communication between a client and a server using unsecure sockets. This code is then modified in Example 8-27 to use JSSE to set up secure socket communication.

Example 8-26 Socket Example Without SSL

The following examples demonstrates server-side and client-side code for setting up an unsecure socket connection.

In a Java program that acts as a server and communicates with a client using sockets, the socket communication is set up with code similar to the following:

```
import java.io.*;
import java.net.*;

. . .
```

```
int port = availablePortNumber;

ServerSocket s;

try {
    s = new ServerSocket(port);
    Socket c = s.accept();

    OutputStream out = c.getOutputStream();
    InputStream in = c.getInputStream();

    // Send messages to the client through
    // the OutputStream
    // Receive messages from the client
    // through the InputStream
} catch (IOException e) { }
```

The client code to set up communication with a server using sockets is similar to the following:

```
import java.io.*;
import java.net.*;

. . .

int port = availablePortNumber;
String host = "hostname";

try {
    s = new Socket(host, port);

    OutputStream out = s.getOutputStream();
    InputStream in = s.getInputStream();

    // Send messages to the server through
    // the OutputStream
    // Receive messages from the server
    // through the InputStream
} catch (IOException e) { }
```

Example 8-27 Socket Example with SSL

The following examples demonstrate server-side and client-side code for setting up a secure socket connection.

In a Java program that acts as a server and communicates with a client using secure sockets, the socket communication is set up with code similar to the following. Differences between this program and the one for communication using unsecure sockets are highlighted in bold.

```
import java.io.*;
import javax.net.ssl.*;
```

```
. . .

int port = availablePortNumber;

SSLServerSocket s;

try {
    SSLServerSocketFactory sslSrvFact =
        (SSLServerSocketFactory)SSLServerSocketFactory.getDefault();
    s = (SSLServerSocket)sslSrvFact.createServerSocket(port);

    SSLSocket c = (SSLSocket)s.accept();

    OutputStream out = c.getOutputStream();
    InputStream in = c.getInputStream();

    // Send messages to the client through
    // the OutputStream
    // Receive messages from the client
    // through the InputStream
}

catch (IOException e) {
}
```

The client code to set up communication with a server using secure sockets is similar to the following, where differences with the unsecure version are highlighted in bold:

```
import java.io.*;
import javax.net.ssl.*;

. . .

int port = availablePortNumber;
String host = "hostname";

try {
    SSLSocketFactory sslFact =
        (SSLSocketFactory)SSLSocketFactory.getDefault();
    SSLSocket s = (SSLSocket)sslFact.createSocket(host, port);

    OutputStream out = s.getOutputStream();
    InputStream in = s.getInputStream();

    // Send messages to the server through
    // the OutputStream
    // Receive messages from the server
    // through the InputStream
}

catch (IOException e) {
}
```

Running the JSSE Sample Code

The JSSE sample programs illustrate how to use JSSE.

- Sample Code Illustrating a Secure Socket Connection Between a Client and a Server
- Sample Code Illustrating HTTPS Connections
- Sample Code Illustrating a Secure RMI Connection
- Sample Code Illustrating the Use of an SSLEngine

When you use the sample code, be aware that the sample programs are designed to illustrate how to use JSSE. They are not designed to be robust applications.

> **Note:**
>
> Setting up secure communications involves complex algorithms. The sample programs provide no feedback during the setup process. When you run the programs, be patient: you may not see any output for a while. If you run the programs with the `javax.net.debug` system property set to `all`, you will see more feedback. For an introduction to reading this debug information, see Debugging TLS Connections.

Where to Find the Sample Code

JSSE Sample Code in the JDK 8 documentation lists all the sample code files and text files. That page also provides a link to a ZIP file that you can download to obtain all the sample code files.

The following sections describe the samples.

Sample Code Illustrating a Secure Socket Connection Between a Client and a Server

The sample programs in the `samples/sockets` directory illustrate how to set up a secure socket connection between a client and a server.

When running the sample client programs, you can communicate with an existing server, such as a web server, or you can communicate with the sample server program, `ClassFileServer`. You can run the sample client and the sample server programs on different machines connected to the same network, or you can run them both on one machine but from different terminal windows.

All the sample `SSLSocketClient*` programs in the samples/sockets/client directory (and `URLReader*` programs described in Sample Code Illustrating HTTPS Connections) can be run with the `ClassFileServer` sample server program. An example of how to do this is shown in Running SSLSocketClientWithClientAuth with ClassFileServer. You can make similar changes to run `URLReader`, `SSLSocketClient`, or `SSLSocketClientWithTunneling` with `ClassFileServer`.

If an authentication error occurs during communication between the client and the server (whether using a web server or `ClassFileServer`), it is most likely because the necessary keys are not in the truststore (trust key database). See Terms and Definitions. For example, the `ClassFileServer` uses a keystore called `testkeys`

containing the private key for `localhost` as needed during the SSL handshake. The `testkeys` keystore is included in the same samples/sockets/server directory as the `ClassFileServer` source. If the client cannot find a certificate for the corresponding public key of `localhost` in the truststore it consults, then an authentication error will occur. Be sure to use the `samplecacerts` truststore (which contains the public key and certificate of the `localhost`), as described in the next section.

Configuration Requirements

When running the sample programs that create a secure socket connection between a client and a server, you will need to make the appropriate certificates file (truststore) available. For both the client and the server programs, you should use the certificates file `samplecacerts` from the `samples` directory. Using this certificates file will allow the client to authenticate the server. The file contains all the common Certificate Authority (CA) certificates shipped with the JDK (in the cacerts file), plus a certificate for `localhost` needed by the client to authenticate `localhost` when communicating with the sample server `ClassFileServer`. The `ClassFileServer` uses a keystore containing the private key for `localhost` that corresponds to the public key in `samplecacerts`.

To make the `samplecacerts` file available to both the client and the server, you can either copy it to the file `java-home/lib/security/jssecacerts`, rename it to cacerts, and use it to replace the `java-home/lib/security/cacerts` file, or add the following option to the command line when running the `java` command for both the client and the server:

```
-Djavax.net.ssl.trustStore=path_to_samplecacerts_file
```

To know more about *java-home*, see Terms and Definitions.

The password for the `samplecacerts` truststore is `changeit`. You can substitute your own certificates in the samples by using the `keytool` utility.

If you use a browser, such as Mozilla Firefox or Microsoft Internet Explorer, to access the sample SSL server provided in the `ClassFileServer` example, then a dialog box may pop up with the message that it does not recognize the certificate. This is normal because the certificate used with the sample programs is self-signed and is for testing only. You can accept the certificate for the current session. After testing the SSL server, you should exit the browser, which deletes the test certificate from the browser's namespace.

For client authentication, a separate `duke` certificate is available in the appropriate directories. The public key and certificate is also stored in the `samplecacerts` file.

Running SSLSocketClient

The `SSLSocketClient.java` program in JSSE Sample Code in the JDK 8 documentation demonstrates how to create a client that uses an `SSLSocket` to send an HTTP request and to get a response from an HTTPS server. The output of this program is the HTML source for `https://www.verisign.com/index.html`.

You must not be behind a firewall to run this program as provided. If you run it from behind a firewall, you will get an `UnknownHostException` because JSSE cannot find a path through your firewall to `www.verisign.com`. To create an equivalent client that can run from behind a firewall, set up proxy tunneling as illustrated in the sample program `SSLSocketClientWithTunneling`.

Running SSLSocketClientWithTunneling

The `SSLSocketClientWithTunneling.java` program in JSSE Sample Code in the JDK 8 documentation illustrates how to do proxy tunneling to access a secure web server from behind a firewall. To run this program, you must set the following Java system properties to the appropriate values:

```
java -Dhttps.proxyHost=webproxy
-Dhttps.proxyPort=ProxyPortNumber
SSLSocketClientWithTunneling
```

> **Note:**
>
> Proxy specifications with the `-D` options are optional. Replace *webproxy* with the name of your proxy host and *ProxyPortNumber* with the appropriate port number.

The program will return the HTML source file from `https://www.verisign.com/index.html`.

Running SSLSocketClientWithClientAuth

The `SSLSocketClientWithClientAuth.java` program in JSSE Sample Code in the JDK 8 documentation shows how to set up a key manager to do client authentication if required by a server. This program also assumes that the client is not outside a firewall. You can modify the program to connect from inside a firewall by following the example in `SSLSocketClientWithTunneling`.

To run this program, you must specify three parameters: host, port, and requested file path. To mirror the previous examples, you can run this program without client authentication by setting the host to `www.verisign.com`, the port to `443`, and the requested file path to `https://www.verisign.com/`. The output when using these parameters is the HTML for the website `https://www.verisign.com/`.

To run `SSLSocketClientWithClientAuth` to do client authentication, you must access a server that requests client authentication. You can use the sample program `ClassFileServer` as this server. This is described in the following sections.

Running ClassFileServer

The program referred to herein as `ClassFileServer` is made up of two files: `ClassFileServer.java` and `ClassServer.java` in JSSE Sample Code in the JDK 8 documentation.

To execute them, run `ClassFileServer.class`, which requires the following parameters:

- `port` can be any available unused port number, for example, you can use the number `2001`.
- `docroot` indicates the directory on the server that contains the file you want to retrieve. For example, on Solaris, you can use /home/*userid*/ (where *userid* refers to your particular UID), whereas on Microsoft Windows systems, you can use c:\.

- `TLS` is an optional parameter that indicates that the server is to use SSL or TLS.
- `true` is an optional parameter that indicates that client authentication is required. This parameter is only consulted if the TLS parameter is set.

> **Note:**
>
> The `TLS` and `true` parameters are optional. If you omit them, indicating that an ordinary (not TLS) file server should be used, without authentication, then nothing happens. This is because one side (the client) is trying to negotiate with TLS, while the other (the server) is not, so they cannot communicate.

> **Note:**
>
> The server expects GET requests in the form `GET /path_to_file`.

Running SSLSocketClientWithClientAuth with ClassFileServer

You can use the sample programs `SSLSocketClientWithClientAuth.java` and `ClassFileServer` in JSSE Sample Code in the JDK 8 documentation to set up authenticated communication, where the client and server are authenticated to each other. You can run both sample programs on different machines connected to the same network, or you can run them both on one machine but from different terminal windows or command prompt windows. To set up both the client and the server, do the following:

1. Run the program `ClassFileServer` from one machine or terminal window. See Running ClassFileServer.
2. Run the program `SSLSocketClientWithClientAuth` on another machine or terminal window. `SSLSocketClientWithClientAuth` requires the following parameters:
   - `host` is the host name of the machine that you are using to run `ClassFileServer`.
   - `port` is the same port that you specified for `ClassFileServer`.
   - `requestedfilepath` indicates the path to the file that you want to retrieve from the server. You must give this parameter as `/filepath`. Forward slashes are required in the file path because it is used as part of a GET statement, which requires forward slashes regardless of what type of operating system you are running. The statement is formed as follows:

     ```
     "GET " + requestedfilepath + " HTTP/1.0"
     ```

> **Note:**
>
> You can modify the other `SSLClient*` applications' `GET` commands to connect to a local machine running `ClassFileServer`.

Sample Code Illustrating HTTPS Connections

There are two primary APIs for accessing secure communications through JSSE. One way is through a socket-level API that can be used for arbitrary secure communications, as illustrated by the `SSLSocketClient`, `SSLSocketClientWithTunneling`, and `SSLSocketClientWithClientAuth` (with and without `ClassFileServer`) sample programs.

A second, and often simpler, way is through the standard Java URL API. You can communicate securely with an SSL-enabled web server by using the HTTPS URL protocol or scheme using the `java.net.URL` class.

Support for HTTPS URL schemes is implemented in many of the common browsers, which allows access to secured communications without requiring the socket-level API provided with JSSE.

An example URL is `https://www.verisign.com`.

The trust and key management for the HTTPS URL implementation is environment-specific. The JSSE implementation provides an HTTPS URL implementation. To use a different HTTPS protocol implementation, set the `java.protocol.handler.pkgs`. See How to Specify a java.lang.System Property to the package name. See the `java.net.URL` class documentation for details.

The samples that you can download with JSSE include two sample programs that illustrate how to create an HTTPS connection. Both of these sample programs (`URLReader.java` and `URLReaderWithOptions.java`) are in the `samples/urls` directory.

Running URLReader

The `URLReader.java` program in JSSE Sample Code in the JDK 8 documentation illustrates using the URL class to access a secure site. The output of this program is the HTML source for `https://www.verisign.com/`. By default, the HTTPS protocol implementation included with JSSE is used. To use a different implementation, set the system property `java.protocol.handler.pkgs` value to be the name of the package containing the implementation.

If you are running the sample code behind a firewall, then you must set the `https.proxyHost` and `https.proxyPort` system properties. For example, to use the proxy host "webproxy" on port 8080, you can use the following options for the `java` command:

```
-Dhttps.proxyHost=webproxy
-Dhttps.proxyPort=8080
```

Alternatively, you can set the system properties within the source code with the `java.lang.System` method `setProperty()`. For example, instead of using the command-line options, you can include the following lines in your program:

```
System.setProperty("java.protocol.handler.pkgs",
"com.ABC.myhttpsprotocol");
System.setProperty("https.proxyHost", "webproxy");
System.setProperty("https.proxyPort", "8080");
```

Running URLReaderWithOptions

The `URLReaderWithOptions.java` program in JSSE Sample Code in the JDK 8 documentation is essentially the same as the `URLReader.java` program, except that it allows you to optionally input any or all of the following system properties as arguments to the program when you run it:

- `java.protocol.handler.pkgs`
- `https.proxyHost`
- `https.proxyPort`
- `https.cipherSuites`

To run `URLReaderWithOptions`, enter the following command:

```
java URLReaderWithOptions [-h proxyhost -p proxyport] [-k
protocolhandlerpkgs] [-c ciphersarray]
```

> **Note:**
>
> Multiple protocol handlers can be included in the `protocolhandlerpkgs` argument as a list with items separated by vertical bars. Multiple SSL cipher suite names can be included in the `ciphersarray` argument as a list with items separated by commas. The possible cipher suite names are the same as those returned by the `SSLSocket.getSupportedCipherSuites()` method. The suite names are taken from the SSL and TLS protocol specifications.

You need a `protocolhandlerpkgs` argument only if you want to use an HTTPS protocol handler implementation other than the default one provided by Oracle.

If you are running the sample code behind a firewall, then you must include arguments for the proxy host and the proxy port. Additionally, you can include a list of cipher suites to enable.

Here is an example of running `URLReaderWithOptions` and specifying the proxy host "webproxy" on port 8080:

```
java URLReaderWithOptions -h webproxy -p 8080
```

Sample Code Illustrating a Secure RMI Connection

The sample code in the `samples/rmi` directory illustrates how to create a secure Java Remote Method Invocation (RMI) connection. The sample code is basically a "Hello World" example modified to install and use a custom RMI socket factory.

Sample Code Illustrating the Use of an SSLEngine

`SSLEngine` gives application developers flexibility when choosing I/O and compute strategies. Rather than tie the SSL/TLS implementation to a specific I/O abstraction (such as single-threaded `SSLSockets`), `SSLEngine` removes the I/O and compute constraints from the SSL/TLS implementation.

As mentioned earlier, `SSLEngine` is an advanced API, and is not appropriate for casual use. Some introductory sample code is provided here that helps illustrate its use. The first demo removes most of the I/O and threading issues, and focuses on many of the SSLEngine methods. The second demo is a more realistic example showing how `SSLEngine` might be combined with Java NIO to create a rudimentary HTTP/HTTPS server.

Running SSLEngineSimpleDemo

The `SSLEngineSimpleDemo.java` program in JSSE Sample Code in the JDK 8 documentation is a very simple application that focuses on the operation of the `SSLEngine` while simplifying the I/O and threading issues. This application creates two `SSLEngine` objects that exchange SSL/TLS messages via common `ByteBuffer` objects. A single loop serially performs all of the engine operations and demonstrates how a secure connection is established (handshaking), how application data is transferred, and how the engine is closed.

The `SSLEngineResult` provides a great deal of information about the current state of the `SSLEngine`. This example does not examine all of the states. It simplifies the I/O and threading issues to the point that this is not a good example for a production environment; nonetheless, it is useful to demonstrate the overall function of the `SSLEngine`.

Running the NIO-Based Server

To fully exploit the flexibility provided by `SSLEngine`, you must first understand complementary APIs, such as I/O and threading models.

An I/O model that large-scale application developers find of use is the NIO `SocketChannel`. NIO was introduced in part to solve some of the scaling problem inherent in the `java.net.Socket` API. `SocketChannel` has many different modes of operation including:

- Blocking
- Nonblocking
- Nonblocking with selectors

Sample code for a basic HTTP server is provided that not only demonstrates many of the new NIO APIs, but also shows how `SSLEngine` can be employed to create a secure HTTPS server. The server is not production quality, but does show many of these new APIs in action.

Inside the samples directory is a `README.txt` file that introduces the server, explains how to build and configure the server, and provides a brief overview of the code layout. The files of most interest for `SSLEngine` users are `ChannelIO.java` and `ChannelIOSecure.java`.

> **Note:**
>
> The server example discussed in this section is included in the JDK. You can find the code bundled in the *`jdk-home`*`/samples/nio/server` directory.

Creating a Keystore to Use with JSSE

The procedure as to how you can use the `keytool` utility to create a simple PKCS12 keystore suitable for use with JSSE.

First you make a `keyEntry` (with public and private keys) in the keystore, and then you make a corresponding `trustedCertEntry` (public keys only) in a truststore. For client authentication, you follow a similar process for the client's certificates.

> **Note:**
>
> Storing trust anchors and secret keys in PKCS12 is supported since JDK 8.

> **Note:**
>
> It is beyond the scope of this example to explain each step in detail. See keytool.

User input is shown in bold.

1. Create a new keystore and self-signed certificate with corresponding public and private keys.

```
% keytool -genkeypair -alias duke -keyalg RSA -validity 7 -keystore
keystore

Enter keystore password:  <password>
What is your first and last name?
[Unknown]:  Duke
What is the name of your organizational unit?
[Unknown]:  Java Software
What is the name of your organization?
[Unknown]:  Oracle, Inc.
What is the name of your City or Locality?
[Unknown]:  Palo Alto
What is the name of your State or Province?
[Unknown]:  CA
What is the two-letter country code for this unit?
[Unknown]:  US
Is CN=Duke, OU=Java Software, O="Oracle, Inc.",
L=Palo Alto, ST=CA, C=US correct?
[no]:  yes
```

2. Examine the keystore. Notice that the entry type is `PrivatekeyEntry`, which means that this entry has a private key associated with it).

```
% keytool -list -v -keystore keystore
```

```
    Enter keystore password:  <password>

    Keystore type: PKCS12
    Keystore provider: SUN

    Your keystore contains 1 entry

    Alias name: duke
    Creation date: Jul 25, 2016
    Entry type: PrivateKeyEntry
    Certificate chain length: 1
    Certificate[1]:
    Owner: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto,
ST=CA, C=US
    Issuer: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto,
ST=CA, C=US
    Serial number: 210cccfc
    Valid from: Mon Jul 25 10:33:27 IST 2016 until: Mon Aug 01 10:33:27
IST 2016
    Certificate fingerprints:
         SHA1: 80:E5:8A:47:7E:4F:5A:70:83:97:DD:F4:DA:29:3D:15:6B:2A:
45:1F
         SHA256: ED:3C:70:68:4E:86:35:9C:63:CC:B9:59:35:58:94:1F:
7E:B8:B0:EE:D2:
    4B:9D:80:31:67:8A:D4:B4:7A:B5:12
    Signature algorithm name: SHA256withRSA
    Subject Public Key Algorithm: RSA (2048)
    Version: 3

    Extensions:

   #1: ObjectId: 2.5.29.14 Criticality=false
   SubjectKeyIdentifier [
   KeyIdentifier [
   0000: 7F C9 95 48 42 8D 68 91   BA 1E E6 5C 2C 6B FF 75  ...HB.h....
\,k.u
   0010: 5F 19 78 43                                        _.xC
   ]
   ]
```

3. Export and examine the self-signed certificate.

```
    % keytool -export -alias duke -keystore keystore -rfc -file duke.cer
    Enter keystore password:  <password>
    Certificate stored in file <duke.cer>
    % cat duke.cer
    -----BEGIN CERTIFICATE-----
    MIIDdzCCAl+gAwIBAgIEIQzM/DANBgkqhkiG9w0BAQsFADBsMQswCQYDVQQGEwJV
    UzELMAkGA1UECBMCQ0ExEjAQBgNVBAcTCVBhbG8gQWx0bzEVMBMGA1UEChMMT3Jh
    Y2xlLCBJbmMuMRYwFAYDVQQLEw1KYXZhIFNvZnR3YXJlMQ0wCwYDVQQDEwREdWtl
    MB4XDTE2MDcyNTA1MDMyN1oXDTE2MDgwMTA1MDMyN1owbDELMAkGA1UEBhMCVVMx
    CzAJBgNVBAgTAkNBMRIwEAYDVQQHEwlQYWxvIEFsdG8xFTATBgNVBAoTDE9yYWNs
    ZSwgSW5jLjEWMBQGA1UECxMNSmF2YSBTb2Z0d2FyZTENMAsGA1UEAxMERHVrZTCC
    ASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBAJ7+Yeu6HDZgWwkGlG4iKH9w
```

```
vGKrxXVR57FaFyheMevrgj1ovVnQVFhfdMvjPkjWmpqLg6rfTqU4bKbtoMWV6+Rn
uQrCw2w9xNC93hX9PxRa20UKrSRDKnUSvi1wjlaxfj0KUKuMwbbY9S8x/naYGeTL
lwbHiiMvkoFkP2kzhVgeqHjIwSz4HRN8vWHCwgIDFWX/ZlS+LbvB4TSZkS0ZcQUV
vJWTocOd8RB90W3bkibWkWq166XYGE1Nq1L4WIhrVJwbav6ual69yJsEpVcshVkx
E1WKzJg7dGb03to4agbReb6+aoCUwb2vNUudNWasSrxoEFArVFGD/ZkPT0esfqEC
AwEAAaMhMB8wHQYDVR0OBBYEFH/JlUhCjWiRuh7mXCxr/3VfGXhDMA0GCSqGSIb3
DQEBCwUAA4IBAQAmcTm2ahsIJLayajsvm8yPzQsHA7kIwWfPPHCoHmNbynG67oHB
fleaNvrgm/raTT3TrqQkg0525qI6Cqaoyy8JA2fAp3i+hmyoGHaIlo14bKazaiPS
RCCqk0J8vwY3CY9nVal1XlHJMEcYV7X1sxKbuAKFoAJ29E/p6ie0JdHtQe31M7X9
FNLYzt8EpJYUtWo13B9Oufz/Guuex9PQ7aC93rbO32MxtnnCGMxQHlaHLLPygc/x
cffGz5Xe5s+NEm78CY7thgN+drI7icBYmv4navsnr2OQaD3AfnJ4WYSQyyUUCPxN
zuk+B0fbLn7PCCcQspmqfgzIpgbEM9M1/yav
-----END CERTIFICATE-----
```

Alternatively, you could generate a Certificate Signing Request (CSR) with `-certreq` and send that to a Certificate Authority (CA) for signing, but that is beyond the scope of this example.

4. Import the certificate into a new truststore.

```
% keytool -import -alias dukecert -file duke.cer -keystore
truststore
Enter keystore password:  <password>
Re-enter new password:
Owner: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto,
ST=CA, C=US
Issuer: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto,
ST=CA, C=US
Serial number: 210cccfc
Valid from: Mon Jul 25 10:33:27 IST 2016 until: Mon Aug 01 10:33:27
IST 2016
Certificate fingerprints:
     SHA1: 80:E5:8A:47:7E:4F:5A:70:83:97:DD:F4:DA:29:3D:15:6B:2A:
45:1F
     SHA256: ED:3C:70:68:4E:86:35:9C:63:CC:B9:59:35:58:94:1F:
7E:B8:B0:EE:D2:
4B:9D:80:31:67:8A:D4:B4:7A:B5:12
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: RSA (2048)
Version: 3

Extensions:

#1: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 7F C9 95 48 42 8D 68 91   BA 1E E6 5C 2C 6B FF 75  ...HB.h....
\,k.u
0010: 5F 19 78 43                                        _.xC
]
]

Trust this certificate? [no]:  yes
Certificate was added to keystore
```

5. Examine the truststore. Note that the entry type is `trustedCertEntry`, which means that a private key is not available for this entry. It also means that this file is not suitable as a keystore of the `KeyManager`.

```
% keytool -list -v -keystore truststore
Enter keystore password:  <password>

Keystore type: PKCS12
Keystore provider: SUN

Your keystore contains 1 entry

Alias name: dukecert
Creation date: Jul 25, 2016
Entry type: trustedCertEntry

Owner: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto,
ST=CA, C=US
Issuer: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto,
ST=CA, C=US
Serial number: 210cccfc
Valid from: Mon Jul 25 10:33:27 IST 2016 until: Mon Aug 01 10:33:27
IST 2016
Certificate fingerprints:
     SHA1: 80:E5:8A:47:7E:4F:5A:70:83:97:DD:F4:DA:29:3D:15:6B:2A:
45:1F
     SHA256: ED:3C:70:68:4E:86:35:9C:63:CC:B9:59:35:58:94:1F:
7E:B8:B0:EE:D2:
4B:9D:80:31:67:8A:D4:B4:7A:B5:12
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: RSA (2048)
Version: 3

Extensions:

#1: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 7F C9 95 48 42 8D 68 91   BA 1E E6 5C 2C 6B FF 75  ...HB.h....
\,k.u
0010: 5F 19 78 43                                        _.xC
]
]



*******************************************
*******************************************
```

6. Now run your applications with the appropriate keystores. Because this example assumes that the default `X509KeyManager` and `X509TrustManager` are used, you select the keystores using the system properties described in Customizing JSSE.

```
% java -Djavax.net.ssl.keyStore=keystore -
Djavax.net.ssl.keyStorePassword=password Server

% java -Djavax.net.ssl.trustStore=truststore -
Djavax.net.ssl.trustStorePassword=trustword Client
```

> **Note:**
>
> This example authenticated the server only. For client authentication, provide a similar keystore for the client's keys and an appropriate truststore for the server.

Using the Server Name Indication (SNI) Extension

These examples illustrate how you can use the Server Name Indication (SNI) Extension for client-side and server-side applications, and how it can be applied to a virtual infrastructure.

For all examples in this section, to apply the parameters after you set them, call the `setSSLParameters(SSLParameters)` method on the corresponding `SSLSocket`, `SSLEngine`, or `SSLServerSocket` object.

Typical Client-Side Usage Examples

The following is a list of use cases that require understanding of the SNI extension for developing a client application:

- Case 1. The client wants to access `www.example.com`.

  Set the host name explicitly:

  ```
  SNIHostName serverName = new SNIHostName("www.example.com");
  sslParameters.setServerNames(Collections.singletonList(serverName));
  ```

  The client should always specify the host name explicitly.

- Case 2. The client does not want to use SNI because the server does not support it.

  Disable SNI with an empty server name list:

  ```
  sslParameters.setServerNames(Collections.emptyList());
  ```

- Case 3. The client wants to access URL `https://www.example.com`.

Oracle providers will set the host name in the SNI extension by default, but third-party providers may not support the default server name indication. To keep your application provider-independent, always set the host name explicitly.

- Case 4. The client wants to switch a socket from server mode to client mode.

  First switch the mode with the following method: `sslSocket.setUseClientMode(true)`. Then reset the server name indication parameters on the socket.

Typical Server-Side Usage Examples

The following is a list of use cases that require understanding of the SNI extension for developing a server application:

- Case 1. The server wants to accept all server name indication types.

  If you do not have any code dealing with the SNI extension, then the server ignores all server name indication types.

- Case 2. The server wants to deny all server name indications of type `host_name`.

  Set an invalid server name pattern for `host_name`:

```
SNIMatcher matcher = SNIHostName.createSNIMatcher("");
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
sslParameters.setSNIMatchers(matchers);
```

  Another way is to create an `SNIMatcher` subclass with a `matches()` method that always returns `false`:

```
class DenialSNIMatcher extends SNIMatcher {
    DenialSNIMatcher() {
        super(StandardConstants.SNI_HOST_NAME);
    }

    @Override
    public boolean matches(SNIServerName serverName) {
        return false;
    }
}

SNIMatcher matcher = new DenialSNIMatcher();
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
sslParameters.setSNIMatchers(matchers);
```

- Case 3. The server wants to accept connections to any host names in the `example.com` domain.

  Set the recognizable server name for `host_name` as a pattern that includes all `*.example.com` addresses:

```
SNIMatcher matcher = SNIHostName.createSNIMatcher("(.*\\.)*example\
\.com");
```

```
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
sslParameters.setSNIMatchers(matchers);
```

- Case 4. The server wants to switch a socket from client mode to server mode.

  First switch the mode with the following method: `sslSocket.setUseClientMode(false)`. Then reset the server name indication parameters on the socket.

Working with Virtual Infrastructures

This section describes how to use the Server Name Indication (SNI) extension from within a virtual infrastructure. It illustrates how to create a parser for ClientHello messages from a socket, provides examples of virtual server dispatchers using `SSLSocket` and `SSLEngine`, describes what happens when the SNI extension is not available, and demonstrates how to create a failover `SSLContext`.

Preparing the ClientHello Parser

Applications must implement an API to parse the ClientHello messages from a socket. The following examples illustrate the `SSLCapabilities` and `SSLExplorer` classes that can perform these functions.

SSLSocketClient.java encapsulates the TLS/DTLS security capabilities during handshaking (that is, the list of cipher suites to be accepted in an TLS/DTLS handshake, the record version, the hello version, and the server name indication). It can be retrieved by exploring the network data of an TLS/DTLS connection via the `SSLExplorer.explore()` method.

SSLExplorer.java explores the initial ClientHello message from a TLS client, but it does not initiate handshaking or consume network data. The `SSLExplorer.explore()` method parses the ClientHello message, and retrieves the security parameters into `SSLCapabilities`. The method must be called before handshaking occurs on any TLS connections.

Virtual Server Dispatcher Based on SSLSocket

This section describes the procedure for using a virtual server dispatcher based on `SSLSocket`.

1. **Register the server name handler.**

   At this step, the application may create different `SSLContext` objects for different server name indications, or link a certain server name indication to a specified virtual machine or distributed system.

   For example, if the server name is `www.example.org`, then the registered server name handler may be for a local virtual hosting web service. The local virtual hosting web service will use the specified `SSLContext`. If the server name is `www.example.com`, then the registered server name handler may be for a virtual machine hosting on `10.0.0.36`. The handler may map this connection to the virtual machine.

2. **Create a `ServerSocket` and accept the new connection.**

```
ServerSocket serverSocket = new ServerSocket(serverPort);
Socket socket = serverSocket.accept();
```

3. **Read and buffer bytes from the socket input stream, and then explore the buffered bytes.**

```
InputStream ins = socket.getInputStream();
byte[] buffer = new byte[0xFF];
int position = 0;
SSLCapabilities capabilities = null;

// Read the header of TLS record
while (position < SSLExplorer.RECORD_HEADER_SIZE) {
    int count = SSLExplorer.RECORD_HEADER_SIZE - position;
    int n = ins.read(buffer, position, count);
    if (n < 0) {
        throw new Exception("unexpected end of stream!");
    }
    position += n;
}

// Get the required size to explore the SSL capabilities
int recordLength = SSLExplorer.getRequiredSize(buffer, 0, position);
if (buffer.length < recordLength) {
    buffer = Arrays.copyOf(buffer, recordLength);
}

while (position < recordLength) {
    int count = recordLength - position;
    int n = ins.read(buffer, position, count);
    if (n < 0) {
        throw new Exception("unexpected end of stream!");
    }
    position += n;
}

// Explore
capabilities = SSLExplorer.explore(buffer, 0, recordLength);
if (capabilities != null) {
    System.out.println("Record version: " +
capabilities.getRecordVersion());
    System.out.println("Hello version: " +
capabilities.getHelloVersion());
}
```

4. **Get the requested server name from the explored capabilities.**

```
List<SNIServerName> serverNames = capabilities.getServerNames();
```

5. **Look for the registered server name handler for this server name indication.**

   If the service of the host name is resident in a virtual machine or another distributed system, then the application must forward the connection to the

destination. The application will need to read and write the raw internet data, rather then the SSL application from the socket stream.

```
Socket destinationSocket = new Socket(serverName, 443);
// Forward buffered bytes and network data from the current socket to
the destinationSocket.
```

If the service of the host name is resident in the same process, and the host name service can use the `SSLSocket` directly, then the application will need to set the `SSLSocket` instance to the server:

```
// Get service context from registered handler
// or create the context
SSLContext serviceContext = ...

SSLSocketFactory serviceSocketFac =
serviceContext.getSSLSocketFactory();

// wrap the buffered bytes
ByteArrayInputStream bais = new ByteArrayInputStream(buffer, 0,
position);
SSLSocket serviceSocket =
(SSLSocket)serviceSocketFac.createSocket(socket, bais, true);

// Now the service can use serviceSocket as usual.
```

Virtual Server Dispatcher Based on SSLEngine

This section describes the procedure for using a virtual server dispatcher based on `SSLEngine`.

1. **Register the server name handler.**

   At this step, the application may create different `SSLContext` objects for different server name indications, or link a certain server name indication to a specified virtual machine or distributed system.

   For example, if the server name is `www.example.org`, then the registered server name handler may be for a local virtual hosting web service. The local virtual hosting web service will use the specified `SSLContext`. If the server name is `www.example.com`, then the registered server name handler may be for a virtual machine hosting on `10.0.0.36`. The handler may map this connection to the virtual machine.

2. **Create a `ServerSocket` or `ServerSocketChannel` and accept the new connection.**

   ```
   ServerSocketChannel serverSocketChannel = ServerSocketChannel.open();
   serverSocketChannel.bind(...);
   ...
   SocketChannel socketChannel = serverSocketChannel.accept();
   ```

3. **Read and buffer bytes from the socket input stream, and then explore the buffered bytes.**

```
ByteBuffer buffer = ByteBuffer.allocate(0xFF);
SSLCapabilities capabilities = null;
while (true) {
    // ensure the capacity
    if (buffer.remaining() == 0) {
        ByteBuffer oldBuffer = buffer;
        buffer = ByteBuffer.allocate(buffer.capacity() + 0xFF);
        buffer.put(oldBuffer);
    }

    int n = sc.read(buffer);
    if (n < 0) {
        throw new Exception("unexpected end of stream!");
    }

    int position = buffer.position();
    buffer.flip();
    capabilities = explorer.explore(buffer);
    buffer.rewind();
    buffer.position(position);
    buffer.limit(buffer.capacity());
    if (capabilities != null) {
        System.out.println("Record version: " +
            capabilities.getRecordVersion());
        System.out.println("Hello version: " +
            capabilities.getHelloVersion());
        break;
    }
}

buffer.flip();  // reset the buffer position and limitation
```

4. **Get the requested server name from the explored capabilities.**

```
List<SNIServerName> serverNames = capabilities.getServerNames();
```

5. **Look for the registered server name handler for this server name indication.**

   If the service of the host name is resident in a virtual machine or another distributed system, then the application must forward the connection to the destination. The application will need to read and write the raw internet data, rather then the SSL application from the socket stream.

```
Socket destinationSocket = new Socket(serverName, 443);
// Forward buffered bytes and network data from the current socket to
the destinationSocket.
```

If the service of the host name is resident in the same process, and the host name service can use the `SSLEngine` directly, then the application will simply feed the net data to the `SSLEngine` instance:

```
// Get service context from registered handler
// or create the context
SSLContext serviceContext = ...

SSLEngine serviceEngine = serviceContext.createSSLEngine();
// Now the service can use the buffered bytes and other byte buffer as
usual.
```

No SNI Extension Available

If there is no server name indication in a ClientHello message, then there is no way to select the proper service according to SNI. For such cases, the application may need to specify a default service, so that the connection can be delegated to it if there is no server name indication.

Failover SSLContext

The `SSLExplorer.explore()` method does not check the validity of TLS/DTLS contents. If the record format does not comply with TLS/DTLS specification, or the `explore()` method is invoked after handshaking has started, then the method may throw an `IOException` and be unable to produce network data. In such cases, handle the exception thrown by `SSLExplorer.explore()` by using a failover `SSLContext`, which is not used to negotiate a TLS/DTLS connection, but to close the connection with the proper alert message. The following example illustrates a failover `SSLContext`. You can find an example of the `DenialSNIMatcher` class in Case 2 in Typical Server-Side Usage Examples.

```
byte[] buffer = ...       // buffered network data
boolean failed = true;    // SSLExplorer.explore() throws an exception

SSLContext context = SSLContext.getInstance("TLS");
// the failover SSLContext

context.init(null, null, null);
SSLSocketFactory sslsf = context.getSocketFactory();
ByteArrayInputStream bais = new ByteArrayInputStream(buffer, 0, position);
SSLSocket sslSocket = (SSLSocket)sslsf.createSocket(socket, bais, true);

SNIMatcher matcher = new DenialSNIMatcher();
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
SSLParameters params = sslSocket.getSSLParameters();
params.setSNIMatchers(matchers);    // no recognizable server name
sslSocket.setSSLParameters(params);

try {
    InputStream sslIS = sslSocket.getInputStream();
    sslIS.read();
} catch (Exception e) {
    System.out.println("Server exception " + e);
} finally {
```

```
    sslSocket.close();
}
```

Standard Names

The JDK Security API requires and uses a set of standard names for algorithms, certificates and keystore types. See the Java Security Standard Algorithm Names specification. Find specific provider information in JDK Providers Documentation.

Provider Pluggability

JSSE is fully pluggable and does not restrict the use of third-party JSSE providers in any way.

JAXP Security Processing

JAXP security processing instructs JAXP components such as parsers, transformers, and so on to behave in a secure fashion. Note that when a Security Manager is present, JAXP security processing is turned on automatically; otherwise, JAXP security processing is disabled by default.

JAXP Security Processing Default Limitations

The following table describes which XML-related factory classes are disabled and which processing limits are set if JAXP security processing is enabled.

Table 8-8 Default Limitations Set by JAXP Security Processing on XML-Related Factory Classes

XML-Related Factory Class	Enabled?	Processing Limits
`DocumentBuilderFactory`	true	`entityExpansionLimit` = 64,000 `elementAttributeLimit` = 10,000 `maxOccurLimit` = 5,000
`SAXParserFactory`	true	`entityExpansionLimit` = 64,000 `elementAttributeLimit` = 10,000 `maxOccurLimit` = 5,000
`SchemaFactory`	true	`maxOccurLimit` = 5,000
`TransformerFactory`	false	Extension functions disabled
`XPathFactory`	false	Extension functions disabled

See Processing Limits in *The Java Tutorials* for more information about `entityExpansionLimit`, `elementAttributeLimit`, `maxOccurLimit` and other JAXP processing limits.

The following sections describes the processing limits in this table in detail and how you can change them.

Limiting Entity Expansion

Limit the the number of entity expansions by either setting the system property `entityExpansionLimit` or the parser property `http://apache.org/xml/properties/entity-expansion-limit`. Both properties accept `java.lang.Integer` values. The parser throws a fatal error once it has reached the entity expansion limit. By default, `entityExpansionLimit` is set to 64,000.

The following command-line example sets the entity expansion limit to 10,000:

```
java -DentityExpansionLimit=10000 MyApp
```

The following code example sets the entity expansion limit to 10,000:

```
System.setProperty("entityExpansionLimit","10000");
```

The following code example sets the parser property `http://apache.org/xml/properties/entity-expansion-limit` to 10,000:

```
DocumentBuilderFactory dfactory = DocumentBuilderFactory.newInstance();
dfactory.setAttribute(
    "http://apache.org/xml/properties/entity-expansion-limit",
    new Integer("10000"));
DocumentBuilder docBuilder = dbFactory.newDocumentBuilder();
```

Limiting Number of Element Attributes

Limit the number of attributes in an element by either setting the system property `elementAttributeLimit` or by setting the parser property `http://apache.org/xml/properties/elementAttributeLimit`. Both properties accept `Integer` values. By default, `elementAttributeLimit` is set to 10,000. When the parser property `http://apache.org/xml/properties/elementAttributeLimit` is set, it overrides the system property. The parser throws a fatal error if the number of attributes in a element exceeds the limit.

The following command-line example sets the element attribute limit to 20:

```
java -DelementAttributeLimit=20 MyApp
```

The following code example sets the element attribute limit to 20:

```
System.setProperty("elementAttributeLimit","20");
```

The following code example sets the parser property `http://apache.org/xml/properties/entity-expansion-limit` to 20:

```
DocumentBuilderFactory dfactory = DocumentBuilderFactory.newInstance();
dfactory.setAttribute(
    "http://apache.org/xml/properties/elementAttributeLimit",
    new Integer(20));
DocumentBuilder docBuilder = dbFactory.newDocumentBuilder();
```

Limit Number of Nodes Created by Constructs That Contain maxOccurs

In constructs like `xsd:sequence`, the validating parser may use space (memory) proportional to the value of the `maxOccurs` occurrence indicator. This may cause the VM to run out of memory, or simply run for a very long time. To prevent potential attacks that exploit this behavior, enable secure processing on a factory as follows:

```
factory.setFeature(XMLConstants.FEATURE_SECURE_PROCESSING, Boolean.TRUE);
```

Note that for `xsd:element` and `xsd:any`, the validating parser uses a constant amount of space, which is independent of the value of the `maxOccurs` occurrence indicator.

The default value of `maxOccursLimit` is 5,000. This system property limits the number of content model nodes that may be created when building a grammar for a W3C XML Schema that contains `maxOccurs` occurrence indicators with values other than "unbounded".

Disabling XPath and XSLT Extension Functions

By default, XPath and XSLT extension functions are disabled when JAXP secure processing is enabled. The following code enables JAXP secure processing and disables XPath and XSLT extension functions for `XPathFactory`:

```
    XPathFactory xpf = xPathFactory.newInstance();
    xpf.setFeature(XMLConstants.FEATURE_SECURE_PROCESSING, true);
```

The following code enables JAXP secure processing and disables XSLT extension functions for `TransformerFactory`:

```
    TransformerFactory tf = TransformerFactory.newInstance();
    tf.setFeature(XMLConstants.FEATURE_SECURE_PROCESSING, true);
```

Security Issue Posed by Nested Entity Definitions

While XML does not allow recursive entity definitions, it does permit nested entity definitions, which produces the potential for Denial of Service attacks on a server which accepts XML data from external sources. For example, a SOAP document like the following that has very deeply nested entity definitions can consume 100% of CPU time and large amounts of memory in entity expansions:

```
<?xml version="1.0" encoding ="UTF-8"?>
<!DOCTYPE foobar[
    <!ENTITY x100 "foobar">
    <!ENTITY  x99 "&x100;&x100;">
    <!ENTITY  x98 "&x99;&x99;">
    ...
    <!ENTITY   x2 "&x3;&x3;">
    <!ENTITY   x1 "&x2;&x2;">
]>
<SOAP-ENV:Envelope xmlns:SOAP-ENV=...>
    <SOAP-ENV:Body>
        <ns1:aaa xmlns:ns1="urn:aaa" SOAP-ENV:encodingStyle="...">
            <foobar xsi:type="xsd:string">&x1;</foobar>
```

```
        </ns1:aaa>
    </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

You don't have worry about this issue if your system doesn't take in external XML data, but a system that does should turn on the secure processing feature and reset the limits as described in Limiting Entity Expansion and Limiting Number of Element Attributes.

Disallowing DTDs

When you set the `http://apache.org/xml/features/disallow-doctype-decl` parser property to true, a fatal error is then thrown if the incoming XML document contains a DOCTYPE declaration. (The default value for this property is false.) This property is typically useful for SOAP based applications where a SOAP message must not contain a Document Type Declaration.

Secure Processing Using StAX

The class `XMLInputFactory` includes the property `javax.xml.stream.supportDTD` that requests processors that do not support DTDs. StAX includes a similar property, `XMLInputFactory.SUPPORT_DTD`, that you can use to disable DTD processing:

```
XMLInputFactory xif = XMLInputFactory.newInstance();
xif.setProperty(XMLInputFactory.SUPPORT_DTD, Boolean.FALSE);
```

Resolving External Resources

The following system properties restrict how XML parsers resolve external resources:

- `javax.xml.XMLConstants.ACCESS_EXTERNAL_DTD`
- `javax.xml.XMLConstants.ACCESS_EXTERNAL_SCHEMA`
- `javax.xml.XMLConstants.ACCESS_EXTERNAL_STYLESHEET`

See JAXP 1.5 and New Properties in *The Java Tutorials* for more information about these properties.

Turning off JAXP Secure Processing

Turn off JAXP secure processing by calling the `setFeature` method on factories. The following code example turns off JAXP secure processing for the SAX parser:

```
SAXParserFactory spf = SAXParserFactory.newInstance();
spf.setFeature(XMLConstants.FEATURE_SECURE_PROCESSING,false);
```

When you turn off JAXP secure processing for the DOM or SAX parser, you remove the default limiations specified by `entityExpansionLimit`, `elementAttributeLimit`, and `maxOccurs`.

9
Running the JSSE Sample Code

The following JSSE sample programs illustrate how to use JSSE:

- Sample Code Illustrating a Secure Socket Connection Between a Client and a Server
- Sample Code Illustrating HTTPS Connections
- Sample Code Illustrating a Secure RMI Connection
- Sample Code Illustrating the Use of an SSLEngine

When you use the sample code, be aware that the sample programs are designed to illustrate how to use JSSE. They are not designed to be robust applications.

> **Note:**
>
> Setting up secure communications involves complex algorithms. The sample programs provide no feedback during the setup process. When you run the programs, be patient: you may not see any output for a while. If you run the programs with the `javax.net.debug` system property set to `all`, you will see more feedback. For an introduction to reading this debug information, see Debugging TLS Connections.

Sample Truststores

This bundle includes sample keystore files that are used with the sample code. They are stored in the "JKS" keystore format, which is the default format used by Oracle' JDK implementation. (If another keystore format is desired, the JDK will need to be configured to recognize the new default format.)

JSSE uses the following certificate keystore files to authenticate the clients and servers.

- `*/testkeys`

  These files are used by the code samples as the source of public/private key and certificate material. In the client program directories, the `testkeys` files contains the certificate entry for the Java mascot `Duke`. In the server program directories (`./sockets/server` and `rmi`), the file contains a certificate entry for the server `localhost`.

  The sample code expects the `testkeys` file to be in the current working directory.

> **Note:**
>
> These are very simple certificates and are not appropriate for a production environment, but they should be sufficient for running the samples here.
>
> The password for these keystores is: `passphrase`

- `samplecacerts`

  This truststore file is very similar to the stock JDK `cacerts` file, in that it contains trust certificates from several vendors. It also contains the trusted certificates from `Duke` and `localhost` above.

  The password for this keystore is the same as the JDK `cacert`'s initial password: `changeit`

  Please see your provider's documentation for how to configure the location of your trusted certificate file.

  > **Note:**
  >
  > Users of the JDK can specify the location of the truststore by using one of the following methods:

  1. System properties:

     ```
     java -Djavax.net.ssl.trustStore=samplecacerts \
          -Djavax.net.ssl.trustStorePassword=changeit Application
     ```

  2. Install the file into:

     ```
     <java-home>/lib/security/jssecacerts
     ```

  3. Install the file into:

     ```
     <java-home>/lib/security/cacerts
     ```

  If you choose (2) or (3), be sure to replace this file with a production `cacerts` file before deployment.

The utility `keytool` can be used to generate alternate certificates and keystore files.

> **Note:**
>
> Ensure that you verify your `cacerts` file. Since you trust the CAs in the `cacerts` file as entities for signing and issuing certificates to other entities, you must manage the `cacerts` file carefully. The `cacerts` file should contain only certificates of the entities and CAs you trust. It is your responsibility to verify the trusted root CA certificates bundled in the `cacerts` file and make your own trust decisions. To remove an untrusted CA certificate from the `cacerts` file, use the delete option of the `keytool` command. You can find the `cacerts` file in the jre installation directory. Contact your system administrator if you do not have permission to edit this file.

Sample Code Illustrating a Secure Socket Connection Between a Client and a Server

These samples illustrate how to set up a secure socket connection between a client and a server.

When running the sample client programs, you can communicate with an existing server, such as a web server, or you can communicate with the sample server program, `ClassFileServer`. You can run the sample client and the sample server programs on different machines connected to the same network, or you can run them both on one machine but from different terminal windows.

All the sample `SSLSocketClient*` programs in the samples/sockets/client directory (and `URLReader*` programs described in Sample Code Illustrating HTTPS Connections) can be run with the `ClassFileServer` sample server program. An example of how to do this is shown in Running SSLSocketClientWithClientAuth with ClassFileServer. You can make similar changes to run `URLReader`, `SSLSocketClient`, or `SSLSocketClientWithTunneling` with `ClassFileServer`.

If an authentication error occurs during communication between the client and the server (whether using a web server or `ClassFileServer`), it is most likely because the necessary keys are not in the truststore (trust key database). See Terms and Definitions. For example, the `ClassFileServer` uses a keystore called `testkeys` containing the private key for `localhost` as needed during the SSL handshake. The `testkeys` keystore is included in the same samples/sockets/server directory as the `ClassFileServer` source. If the client cannot find a certificate for the corresponding public key of `localhost` in the truststore it consults, then an authentication error will occur. Be sure to use the `samplecacerts` truststore (which contains the public key and certificate of the `localhost`), as described in Configuration Requirements for SSL Socket Samples.

Configuration Requirements for SSL Socket Samples

When running the sample programs that create a secure socket connection between a client and a server, you will need to make the appropriate certificates file (truststore) available. For both the client and the server programs, you should use the certificates file `samplecacerts` from the `samples` directory. Using this certificates file will allow the

client to authenticate the server. The file contains all the common Certificate Authority (CA) certificates shipped with the JDK (in the cacerts file), plus a certificate for `localhost` needed by the client to authenticate `localhost` when communicating with the sample server `ClassFileServer`. The `ClassFileServer` uses a keystore containing the private key for `localhost` that corresponds to the public key in `samplecacerts`.

To make the `samplecacerts` file available to both the client and the server, you can either copy it to the file `java-home/lib/security/jssecacerts`, rename it to `cacerts`, and use it to replace the `java-home/lib/security/cacerts` file, or add the following option to the command line when running the `java` command for both the client and the server:

```
-Djavax.net.ssl.trustStore=path_to_samplecacerts_file
```

The password for the `samplecacerts` truststore is `changeit`. You can substitute your own certificates in the samples by using the `keytool` utility.

If you use a browser to access the sample SSL server provided in the `ClassFileServer` example, then a dialog box may pop up with the message that it does not recognize the certificate. This is normal because the certificate used with the sample programs is self-signed and is for testing only. You can accept the certificate for the current session. After testing the SSL server, you should exit the browser, which deletes the test certificate from the browser's namespace.

For client authentication, a separate `duke` certificate is available in the appropriate directories. The public key and certificate is also stored in the `samplecacerts` file.

Running SSLSocketClient

The example `SSLSocketClient.java` demonstrates how to create a client that uses an `SSLSocket` to send an HTTP request and to get a response from an HTTPS server. By default, this example connects to `www.verisign.com`, but it can easily be adapted to connect to `ClassFileServer` (see Running ClassFileServer). The output of this program is the HTML source for `https://www.verisign.com/index.html`.

You must not be behind a firewall to run this program as provided. If you run it from behind a firewall, you will get an `UnknownHostException` because JSSE cannot find a path through your firewall to `www.verisign.com`. To create an equivalent client that can run from behind a firewall, set up proxy tunneling as illustrated in the sample program `SSLSocketClientWithTunneling`.

> **Note:**
>
> The GET request must be slightly modified so that a file is specified.

Usage

```
java SSLSocketClient
```

SSLSocketClient.java

```
import java.net.*;
import java.io.*;
import javax.net.ssl.*;

/*
 * This example demostrates how to use a SSLSocket as client to
 * send a HTTP request and get response from an HTTPS server.
 * It assumes that the client is not behind a firewall
 */

public class SSLSocketClient {

    public static void main(String[] args) throws Exception {
        try {
            SSLSocketFactory factory =
                (SSLSocketFactory)SSLSocketFactory.getDefault();
            SSLSocket socket =
                (SSLSocket)factory.createSocket("www.verisign.com", 443);

            /*
             * send http request
             *
             * Before any application data is sent or received, the
             * SSL socket will do SSL handshaking first to set up
             * the security attributes.
             *
             * SSL handshaking can be initiated by either flushing data
             * down the pipe, or by starting the handshaking by hand.
             *
             * Handshaking is started manually in this example because
             * PrintWriter catches all IOExceptions (including
             * SSLExceptions), sets an internal error flag, and then
             * returns without rethrowing the exception.
             *
             * Unfortunately, this means any error messages are lost,
             * which caused lots of confusion for others using this
             * code.  The only way to tell there was an error is to call
             * PrintWriter.checkError().
             */
            socket.startHandshake();

            PrintWriter out = new PrintWriter(
                                  new BufferedWriter(
                                  new OutputStreamWriter(
                                  socket.getOutputStream())));

            out.println("GET / HTTP/1.0");
            out.println();
            out.flush();

            /*
             * Make sure there were no surprises
```

```
             */
            if (out.checkError())
                System.out.println(
                    "SSLSocketClient:  java.io.PrintWriter error");

            /* read response */
            BufferedReader in = new BufferedReader(
                                    new InputStreamReader(
                                    socket.getInputStream()));

            String inputLine;
            while ((inputLine = in.readLine()) != null)
                System.out.println(inputLine);

            in.close();
            out.close();
            socket.close();

        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

Running SSLSocketClientWithTunnelling

The example `SSLSocketClientWithTunneling.java` illustrates how to do proxy tunneling to access a secure web server from behind a firewall.

The program will return the HTML source file from `https://www.verisign.com/index.html`.

Usage

```
java -Dhttps.proxyHost=webproxy -Dhttps.proxyPort=ProxyPortNumber
SSLSocketClientWithTunneling
```

Replace *`webproxy`* with the name of your proxy host and *`ProxyPortNumber`* with the appropriate port number.

The system properties `https.proxyHost` and `https.proxyPort` are used to make a socket connection to the proxy host, and then the `SSLSocket` is layered on top of that `Socket`.

SSLSocketClientWithTunneling.java

```
import java.net.*;
import java.io.*;
import javax.net.ssl.*;

/*
 * This example illustrates how to do proxy Tunneling to access a
 * secure web server from behind a firewall.
 *
```

```
 * Please set the following Java system properties
 * to the appropriate values:
 *
 *   https.proxyHost = <secure proxy server hostname>
 *   https.proxyPort = <secure proxy server port>
 */

public class SSLSocketClientWithTunneling {

    public static void main(String[] args) throws Exception {
        new SSLSocketClientWithTunneling().doIt("www.verisign.com", 443);
    }

    String tunnelHost;
    int tunnelPort;

    public void doIt(String host, int port) {
        try {

            /*
             * Let's setup the SSLContext first, as there's a lot of
             * computations to be done.  If the socket were created
             * before the SSLContext, the server/proxy might timeout
             * waiting for the client to actually send something.
             */
            SSLSocketFactory factory =
                (SSLSocketFactory)SSLSocketFactory.getDefault();

            /*
             * Set up a socket to do tunneling through the proxy.
             * Start it off as a regular socket, then layer SSL
             * over the top of it.
             */
            tunnelHost = System.getProperty("https.proxyHost");
            tunnelPort = Integer.getInteger("https.proxyPort").intValue();

            Socket tunnel = new Socket(tunnelHost, tunnelPort);
            doTunnelHandshake(tunnel, host, port);

            /*
             * Ok, let's overlay the tunnel socket with SSL.
             */
            SSLSocket socket =
                (SSLSocket)factory.createSocket(tunnel, host, port, true);

            /*
             * register a callback for handshaking completion event
             */
            socket.addHandshakeCompletedListener(
                new HandshakeCompletedListener() {
                    public void handshakeCompleted(
                            HandshakeCompletedEvent event) {
                        System.out.println("Handshake finished!");
                        System.out.println(
                            "\t CipherSuite:" + event.getCipherSuite());
```

```
                        System.out.println(
                            "\t SessionId " + event.getSession());
                        System.out.println(
                            "\t PeerHost " +
event.getSession().getPeerHost());
                    }
                }
            );

            /*
             * send http request
             *
             * See SSLSocketClient.java for more information about why
             * there is a forced handshake here when using PrintWriters.
             */
            socket.startHandshake();

            PrintWriter out = new PrintWriter(
                                  new BufferedWriter(
                                  new OutputStreamWriter(
                                  socket.getOutputStream()));

            out.println("GET / HTTP/1.0");
            out.println();
            out.flush();

            /*
             * Make sure there were no surprises
             */
            if (out.checkError())
                System.out.println(
                    "SSLSocketClient:  java.io.PrintWriter error");

            /* read response */
            BufferedReader in = new BufferedReader(
                                    new InputStreamReader(
                                    socket.getInputStream()));

            String inputLine;

            while ((inputLine = in.readLine()) != null)
                System.out.println(inputLine);

            in.close();
            out.close();
            socket.close();
            tunnel.close();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }

    /*
     * Tell our tunnel where we want to CONNECT, and look for the
     * right reply.  Throw IOException if anything goes wrong.
```

```
     */
    private void doTunnelHandshake(Socket tunnel, String host, int port)
    throws IOException
    {
        OutputStream out = tunnel.getOutputStream();
        String msg = "CONNECT " + host + ":" + port + " HTTP/1.0\n"
                     + "User-Agent: "
                     +
sun.net.www.protocol.http.HttpURLConnection.userAgent
                     + "\r\n\r\n";
        byte b[];
        try {
            /*
             * We really do want ASCII7 -- the http protocol doesn't change
             * with locale.
             */
            b = msg.getBytes("ASCII7");
        } catch (UnsupportedEncodingException ignored) {
            /*
             * If ASCII7 isn't there, something serious is wrong, but
             * Paranoia Is Good (tm)
             */
            b = msg.getBytes();
        }
        out.write(b);
        out.flush();

        /*
         * We need to store the reply so we can create a detailed
         * error message to the user.
         */
        byte            reply[] = new byte[200];
        int             replyLen = 0;
        int             newlinesSeen = 0;
        boolean         headerDone = false;     /* Done on first newline */

        InputStream     in = tunnel.getInputStream();
        boolean         error = false;

        while (newlinesSeen < 2) {
            int i = in.read();
            if (i < 0) {
                throw new IOException("Unexpected EOF from proxy");
            }
            if (i == '\n') {
                headerDone = true;
                ++newlinesSeen;
            } else if (i != '\r') {
                newlinesSeen = 0;
                if (!headerDone && replyLen < reply.length) {
                    reply[replyLen++] = (byte) i;
                }
            }
        }
```

```
        /*
         * Converting the byte array to a string is slightly wasteful
         * in the case where the connection was successful, but it's
         * insignificant compared to the network overhead.
         */
        String replyStr;
        try {
            replyStr = new String(reply, 0, replyLen, "ASCII7");
        } catch (UnsupportedEncodingException ignored) {
            replyStr = new String(reply, 0, replyLen);
        }

        /* We asked for HTTP/1.0, so we should get that back */
        if (!replyStr.startsWith("HTTP/1.0 200")) {
            throw new IOException("Unable to tunnel through "
                    + tunnelHost + ":" + tunnelPort
                    + ".  Proxy returns \"" + replyStr + "\"");
        }

        /* tunneling Handshake was successful! */
    }
}
```

Running SSLSocketClientWithClientAuth

The example `SSLSocketClientWithClientAuth.java` is similar to Running SSLSocketClient, but this shows how to set up a key manager to do client authentication if required by a server. This program also assumes that the client is not outside a firewall. You can modify the program to connect from inside a firewall by following the example in Running SSLSocketClientWithTunnelling.

To run this program, you must specify three parameters: host, port, and requested file path. To mirror the previous examples, you can run this program without client authentication by setting the host to `www.verisign.com`, the port to `443`, and the requested file path to `https://www.verisign.com/`. The output when using these parameters is the HTML for the website `https://www.verisign.com/`.

To run `SSLSocketClientWithClientAuth` to do client authentication, you must access a server that requests client authentication. You can use the sample program `ClassFileServer` as this server. See Running SSLSocketClientWithClientAuth with ClassFileServer.

Usage

```
java SSLSocketClientWithClientAuth host port requestedfilepath
```

> **Note:**
>
> If you are connecting to the `ClassFileServer` application above, be sure that it can find the credentials for the `duke` user. See Sample Truststores.

SSLSocketClientWithClientAuth.java

```
import java.net.*;
import java.io.*;
import javax.net.ssl.*;
import javax.security.cert.X509Certificate;
import java.security.KeyStore;

/*
 * This example shows how to set up a key manager to do client
 * authentication if required by server.
 *
 * This program assumes that the client is not inside a firewall.
 * The application can be modified to connect to a server outside
 * the firewall by following SSLSocketClientWithTunneling.java.
 */
public class SSLSocketClientWithClientAuth {

    public static void main(String[] args) throws Exception {
        String host = null;
        int port = -1;
        String path = null;
        for (int i = 0; i < args.length; i++)
            System.out.println(args[i]);

        if (args.length < 3) {
            System.out.println(
                "USAGE: java SSLSocketClientWithClientAuth " +
                "host port requestedfilepath");
            System.exit(-1);
        }

        try {
            host = args[0];
            port = Integer.parseInt(args[1]);
            path = args[2];
        } catch (IllegalArgumentException e) {
             System.out.println("USAGE: java SSLSocketClientWithClientAuth
" +
                 "host port requestedfilepath");
             System.exit(-1);
        }

        try {

            /*
             * Set up a key manager for client authentication
             * if asked by the server.  Use the implementation's
             * default TrustStore and secureRandom routines.
             */
            SSLSocketFactory factory = null;
            try {
                SSLContext ctx;
                KeyManagerFactory kmf;
```

```
    KeyStore ks;
    char[] passphrase = "passphrase".toCharArray();

    ctx = SSLContext.getInstance("TLS");
    kmf = KeyManagerFactory.getInstance("SunX509");
    ks = KeyStore.getInstance("JKS");

    ks.load(new FileInputStream("testkeys"), passphrase);

    kmf.init(ks, passphrase);
    ctx.init(kmf.getKeyManagers(), null, null);

    factory = ctx.getSocketFactory();
} catch (Exception e) {
    throw new IOException(e.getMessage());
}

SSLSocket socket = (SSLSocket)factory.createSocket(host, port);

/*
 * send http request
 *
 * See SSLSocketClient.java for more information about why
 * there is a forced handshake here when using PrintWriters.
 */
socket.startHandshake();

PrintWriter out = new PrintWriter(
                      new BufferedWriter(
                      new OutputStreamWriter(
                      socket.getOutputStream())));
out.println("GET " + path + " HTTP/1.0");
out.println();
out.flush();

/*
 * Make sure there were no surprises
 */
if (out.checkError())
    System.out.println(
        "SSLSocketClient: java.io.PrintWriter error");

/* read response */
BufferedReader in = new BufferedReader(
                        new InputStreamReader(
                        socket.getInputStream()));

String inputLine;

while ((inputLine = in.readLine()) != null)
    System.out.println(inputLine);

in.close();
out.close();
socket.close();
```

```
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

Running ClassFileServer

This example, which consists of `ClassFileServer.java` and `ClassServer.java`, demonstrates the implementation of a mini-webserver, which can service simple HTTP or HTTPS requests (only the GET method is supported). By default, the server does not use SSL/TLS. However, a command-line option enables SSL/TLS. Requests must be of the form:

```
GET /<filename>
```

Usage

```
java ClassFileServer port docroot [TLS [true]]
```

- *`port`*: The port on which the server resides. It can be any available unused port number, for example, you can use the number `2001`.
- *`docroot`*: The root of the local directory hierarchy. It indicates the directory on the server that contains the file you want to retrieve. For example, on Linux, you can use `/home/userid/` (where *`userid`* refers to your particular UID), whereas on Windows systems, you can use `C:\`.
- `TLS`: An optional flag which enables SSL/TLS services. If you omit the `TLS` and `true` parameters, which indicates that an ordinary (not TLS) file server should be used, without authentication, then nothing happens. This is because one side (the client) is trying to negotiate with TLS, while the other (the server) is not, so they cannot communicate.
- `true`: An optional flag which requires that clients authenticate themselves. This option requires that SSL/TLS support be enabled.

The secure server comes preinstalled with a certificate for `localhost`. If the server is on the same host as the client, URLs of the form `https://localhost:port/file` should pass hostname verification. If you choose to run on separate hosts, you should create a new host certificate for the https hostname being used, otherwise there will be hostname mismatch problems. (Note: in Java, this can be corrected in the `HttpsURLConnection` class by providing a custom `HostnameVerifier` implementation, or in a browser by accepting the dialog box that describes the hostname mismatch.)

If you are using the TLS variant (HTTPS), remember to specify the https protocol:

```
https://hostname:2001/dir1/file1
```

> **Note:**
>
> If you use a browser, you will see a dialog popup with the message that the application doesn't recognize the `localhost` certificate. This is normal because the self-signed certificate being presented to the browser is not initially trusted. If desired, you could import the `localhost` certificate into the browser's truststore.

ClassFileServer.java

```
import java.io.*;
import java.net.*;
import java.security.KeyStore;
import javax.net.*;
import javax.net.ssl.*;
import javax.security.cert.X509Certificate;

/* ClassFileServer.java -- a simple file server that can server
 * Http get request in both clear and secure channel
 *
 * The ClassFileServer implements a ClassServer that
 * reads files from the file system. See the
 * doc for the "Main" method for how to run this
 * server.
 */

public class ClassFileServer extends ClassServer {

    private String docroot;

    private static int DefaultServerPort = 2001;

    /**
     * Constructs a ClassFileServer.
     *
     * @param path the path where the server locates files
     */
    public ClassFileServer(ServerSocket ss, String docroot) throws
IOException
    {
        super(ss);
        this.docroot = docroot;
    }

    /**
     * Returns an array of bytes containing the bytes for
     * the file represented by the argument <b>path</b>.
     *
     * @return the bytes for the file
     * @exception FileNotFoundException if the file corresponding
     * to <b>path</b> could not be loaded.
     */
    public byte[] getBytes(String path)
```

```
    throws IOException
{
    System.out.println("reading: " + path);
    File f = new File(docroot + File.separator + path);
    int length = (int)(f.length());
    if (length == 0) {
        throw new IOException("File length is zero: " + path);
    } else {
        FileInputStream fin = new FileInputStream(f);
        DataInputStream in = new DataInputStream(fin);

        byte[] bytecodes = new byte[length];
        in.readFully(bytecodes);
        return bytecodes;
    }
}

/**
 * Main method to create the class server that reads
 * files. This takes two command line arguments, the
 * port on which the server accepts requests and the
 * root of the path. To start up the server: <br><br>
 *
 * <code>   java ClassFileServer <port> <path>
 * </code><br><br>
 *
 * <code>   new ClassFileServer(port, docroot);
 * </code>
 */
public static void main(String args[])
{
    System.out.println(
        "USAGE: java ClassFileServer port docroot [TLS [true]]");
    System.out.println("");
    System.out.println(
        "If the third argument is TLS, it will start as\n" +
        "a TLS/SSL file server, otherwise, it will be\n" +
        "an ordinary file server. \n" +
        "If the fourth argument is true,it will require\n" +
        "client authentication as well.");

    int port = DefaultServerPort;
    String docroot = "";

    if (args.length >= 1) {
        port = Integer.parseInt(args[0]);
    }

    if (args.length >= 2) {
        docroot = args[1];
    }
    String type = "PlainSocket";
    if (args.length >= 3) {
        type = args[2];
    }
```

```
        try {
            ServerSocketFactory ssf =
                ClassFileServer.getServerSocketFactory(type);
            ServerSocket ss = ssf.createServerSocket(port);
            if (args.length >= 4 && args[3].equals("true")) {
                ((SSLServerSocket)ss).setNeedClientAuth(true);
            }
            new ClassFileServer(ss, docroot);
        } catch (IOException e) {
            System.out.println("Unable to start ClassServer: " +
                               e.getMessage());
            e.printStackTrace();
        }
    }

    private static ServerSocketFactory getServerSocketFactory(String type)
{
        if (type.equals("TLS")) {
            SSLServerSocketFactory ssf = null;
            try {
                // set up key manager to do server authentication
                SSLContext ctx;
                KeyManagerFactory kmf;
                KeyStore ks;
                char[] passphrase = "passphrase".toCharArray();

                ctx = SSLContext.getInstance("TLS");
                kmf = KeyManagerFactory.getInstance("SunX509");
                ks = KeyStore.getInstance("JKS");

                ks.load(new FileInputStream("testkeys"), passphrase);
                kmf.init(ks, passphrase);
                ctx.init(kmf.getKeyManagers(), null, null);

                ssf = ctx.getServerSocketFactory();
                return ssf;
            } catch (Exception e) {
                e.printStackTrace();
            }
        } else {
            return ServerSocketFactory.getDefault();
        }
        return null;
    }
}
```

ClassServer.java

```
import java.io.*;
import java.net.*;
import javax.net.*;

/*
 * ClassServer.java -- a simple file server that can serve
```

```
 * Http get request in both clear and secure channel
 */

public abstract class ClassServer implements Runnable {

    private ServerSocket server = null;
    /**
     * Constructs a ClassServer based on <b>ss</b> and
     * obtains a file's bytecodes using the method <b>getBytes</b>.
     *
     */
    protected ClassServer(ServerSocket ss)
    {
        server = ss;
        newListener();
    }

    /**
     * Returns an array of bytes containing the bytes for
     * the file represented by the argument <b>path</b>.
     *
     * @return the bytes for the file
     * @exception FileNotFoundException if the file corresponding
     * to <b>path</b> could not be loaded.
     * @exception IOException if error occurs reading the class
     */
    public abstract byte[] getBytes(String path)
        throws IOException, FileNotFoundException;

    /**
     * The "listen" thread that accepts a connection to the
     * server, parses the header to obtain the file name
     * and sends back the bytes for the file (or error
     * if the file is not found or the response was malformed).
     */
    public void run()
    {
        Socket socket;

        // accept a connection
        try {
            socket = server.accept();
        } catch (IOException e) {
            System.out.println("Class Server died: " + e.getMessage());
            e.printStackTrace();
            return;
        }

        // create a new thread to accept the next connection
        newListener();

        try {
            OutputStream rawOut = socket.getOutputStream();

            PrintWriter out = new PrintWriter(
```

```
                                new BufferedWriter(
                                new OutputStreamWriter(
                                rawOut)));
            try {
                // get path to class file from header
                BufferedReader in =
                    new BufferedReader(
                        new InputStreamReader(socket.getInputStream()));
                String path = getPath(in);
                // retrieve bytecodes
                byte[] bytecodes = getBytes(path);
                // send bytecodes in response (assumes HTTP/1.0 or later)
                try {
                    out.print("HTTP/1.0 200 OK\r\n");
                    out.print("Content-Length: " + bytecodes.length +
                                    "\r\n");
                    out.print("Content-Type: text/html\r\n\r\n");
                    out.flush();
                    rawOut.write(bytecodes);
                    rawOut.flush();
                } catch (IOException ie) {
                    ie.printStackTrace();
                    return;
                }

            } catch (Exception e) {
                e.printStackTrace();
                // write out error response
                out.println("HTTP/1.0 400 " + e.getMessage() + "\r\n");
                out.println("Content-Type: text/html\r\n\r\n");
                out.flush();
            }

        } catch (IOException ex) {
            // eat exception (could log error to log file, but
            // write out to stdout for now).
            System.out.println("error writing response: " +
ex.getMessage());
            ex.printStackTrace();

        } finally {
            try {
                socket.close();
            } catch (IOException e) {
            }
        }
    }

    /**
     * Create a new thread to listen.
     */
    private void newListener()
    {
        (new Thread(this)).start();
    }
```

```
    /**
     * Returns the path to the file obtained from
     * parsing the HTML header.
     */
    private static String getPath(BufferedReader in)
        throws IOException
    {
        String line = in.readLine();
        String path = "";
        // extract class from GET line
        if (line.startsWith("GET /")) {
            line = line.substring(5, line.length()-1).trim();
            int index = line.indexOf(' ');
            if (index != -1) {
                path = line.substring(0, index);
            }
        }

        // eat the rest of header
        do {
            line = in.readLine();
        } while ((line.length() != 0) &&
                 (line.charAt(0) != '\r') && (line.charAt(0) != '\n'));

        if (path.length() != 0) {
            return path;
        } else {
            throw new IOException("Malformed Header");
        }
    }
}
```

Running SSLSocketClientWithClientAuth with ClassFileServer

You can use the sample programs `SSLSocketClientWithClientAuth.java` and `ClassFileServer.;ava` to set up authenticated communication, where the client and server are authenticated to each other. You can run both sample programs on different machines connected to the same network, or you can run them both on one machine but from different terminal windows or command prompt windows. To set up both the client and the server, do the following:

1. Run the program `ClassFileServer` from one machine or terminal window. See Running ClassFileServer.
2. Run the program `SSLSocketClientWithClientAuth` on another machine or terminal window. `SSLSocketClientWithClientAuth` requires the following parameters:
   - `host` is the host name of the machine that you are using to run `ClassFileServer`.
   - `port` is the same port that you specified for `ClassFileServer`.

- `requestedfilepath` indicates the path to the file that you want to retrieve from the server. You must give this parameter as `/filepath`. Forward slashes are required in the file path because it is used as part of a GET statement, which requires forward slashes regardless of what type of operating system you are running. The statement is formed as follows:

  ```
  "GET " + requestedfilepath + " HTTP/1.0"
  ```

> **Note:**
>
> You can modify the other `SSLClient*` applications' `GET` commands to connect to a local machine running `ClassFileServer`.

Sample Code Illustrating HTTPS Connections

There are two primary APIs for accessing secure communications through JSSE. One way is through a socket-level API that can be used for arbitrary secure communications, as illustrated by the `SSLSocketClient.java`, `SSLSocketClientWithTunneling.java`, and `SSLSocketClientWithClientAuth.java` examples (with and without the examples described in Running ClassFileServer).

A second, and often simpler, way is through the standard Java URL API. You can communicate securely with an SSL-enabled web server by using the HTTPS URL protocol or scheme using the `java.net.URL` class.

Support for HTTPS URL schemes is implemented in many of the common browsers, which allows access to secured communications without requiring the socket-level API provided with JSSE. An example URL is `https://www.verisign.com`.

The trust and key management for the HTTPS URL implementation is environment-specific. The JSSE implementation provides an HTTPS URL implementation. To use a different HTTPS protocol implementation, set the `java.protocol.handler.pkgs` system property; see How to Specify a java.lang.System Property.

Running URLReader

The example `URLReader.java` illustrates using a URL to access resources on a secure site. By default, this example connects to `www.verisign.com`, but it can be adapted to connect to `ClassFileServer.java`. To do so, the URL will need to be modified to point to the correct address. You may also need to update the server's certificate or provide a custom `HostNameVerifier` (see `HttpsURLConnection`) if the hostname in the server's certificate doesn't match the URL's hostname.

> **Note:**
>
> If you are behind a firewall, you may need to set the `https.proxyHost` and `https.proxyPort` system properties to correctly specify the proxy.

Usage

```
java URLReader
```

URLReader.java

```
import java.net.*;
import java.io.*;

/*
 * This example illustrates using a URL to access resources
 * on a secure site.
 *
 * If you are running inside a firewall, please also set the following
 * Java system properties to the appropriate value:
 *
 *   https.proxyHost = <secure proxy server hostname>
 *   https.proxyPort = <secure proxy server port>
 *
 */

public class URLReader {
    public static void main(String[] args) throws Exception {
        URL verisign = new URL("https://www.verisign.com/");
        BufferedReader in = new BufferedReader(
                                new InputStreamReader(
                                verisign.openStream()));

        String inputLine;

        while ((inputLine = in.readLine()) != null)
            System.out.println(inputLine);

        in.close();
    }
}
```

Running URLReaderWithOptions

The example is very similar to `URLReaader.java`, but it enables you to set the system properties through main method arguments rather than as `-D` options to the Java runtime environment.

Usage

`java URLReaderWithOptions [-h` *`proxyhost`*`] [-p` *`proxyport`*`] [-k` *`protocolhandlerpkgs`*`] [-c` *`ciphersarray`*`]`

- *`proxyHost`*: secure proxy server hostname (`https.proxyHost`)
- *`proxyPort`*: secure proxy server port (`https.proxyPort`)

- *protocolhandlerpkgs*: a pipe-separated (|) list of protocol handlers (`java.protocol.handler.pkgs`)
- *ciphersarray*: enabled cipher suites as a comma-separated list (`https.cipherSuites`)

> **Note:**
>
> Multiple protocol handlers can be included in the *protocolhandlerpkgs* argument as a list with items separated by vertical bars. Multiple SSL cipher suite names can be included in the *ciphersarray* argument as a list with items separated by commas. The possible cipher suite names are the same as those returned by the `SSLSocket.getSupportedCipherSuites()` method. The suite names are taken from the SSL and TLS protocol specifications.

You need a *protocolhandlerpkgs* argument only if you want to use an HTTPS protocol handler implementation other than the default one provided by Oracle.

If you are running the sample code behind a firewall, then you must include arguments for the proxy host and the proxy port. Additionally, you can include a list of cipher suites to enable.

Here is an example of running `URLReaderWithOptions` and specifying the proxy host "webproxy" on port 8080:

```
java URLReaderWithOptions -h webproxy -p 8080
```

URLReaderWithOptions.java

```
import java.net.*;
import java.io.*;

/*
 * Using a URL to access resources on a secure site.
 *
 * You can optionally set the following command line options:
 *
 *     -h <secure proxy server hostname>
 *     -p <secure proxy server port>
 *     -k <| separated list of protocol handlers>
 *     -c <enabled cipher suites as a comma separated list>
 *
 */

public class URLReaderWithOptions {
    public static void main(String[] args) throws Exception {

        System.out.println("USAGE: java URLReaderWithOptions " +
            "[-h proxyhost] [-p proxyport] [-k protocolhandlerpkgs] " +
            "[-c ciphersarray]");

        // initialize system properties
        char option = 'd';
```

```
        for (int i = 0; i < args.length; i++) {
            System.out.println(option+": "+args[i]);
            switch(option) {
            case 'h':
                System.setProperty("https.proxyHost", args[i]);
                option = 'd';
                break;
            case 'p':
                System.setProperty("https.proxyPort", args[i]);
                option = 'd';
                break;
            case 'k':
                System.setProperty("java.protocol.handler.pkgs", args[i]);
                option = 'd';
                break;
            case 'c':
                System.setProperty("https.cipherSuites", args[i]);
                option = 'd';
                break;
            default:
                // get the next option
                if (args[i].startsWith("-")) {
                    option = args[i].charAt(1);
                }
            }
        }

        URL verisign = new URL("https://www.verisign.com/");
        BufferedReader in = new BufferedReader(
                                new InputStreamReader(
                                verisign.openStream()));

        String inputLine;

        while ((inputLine = in.readLine()) != null)
            System.out.println(inputLine);

        in.close();
    }
}
```

Sample Code Illustrating a Secure RMI Connection

The example `HelloClient.java` illustrates how to create a secure Java Remote Method Invocation (RMI) connection. The sample code is basically a "Hello World" example modified to install and use a custom RMI socket factory. It uses RMI over an SSL transport layer using JSSE. The server runs `HelloImpl.java`, which sets up an internal RMI registry (rather than using the `rmiregistry` command). The client runs `HelloClient` and communicates over a secured connection.

Usage

Setting up this sample can be a little tricky; here are the necessary steps:

`% javac *.java`

`% rmic HelloImpl`

`% java -Djava.security.policy=policy HelloImpl` (run in another window)

`% java HelloClient` (run in another window)

For the server, the RMI security manager will be installed, and the supplied policy file grants permission to accept connections from any host. Obviously, giving all permissions should not be done in a production environment. You will need to give it the appropriate restrictive network privileges, such as:

```
permission java.net.SocketPermission "hostname:1024-", "accept,resolve";
```

In addition, this example can be easily updated to run with the standard SSL/TLS-based RMI socket factories. To do this, modify the `HelloImpl.java` file to use:

```
javax.rmi.ssl.SslRMIClientSocketFactory
javax.rmi.ssl.SslRMIServerSocketFactory
```

instead of:

```
RMISSLClientSocketFactory
RMISSLServerSocketFactory
```

These classes use `SSLSocketFactory.getDefault()` and `SSLServerSocketFactory.getDefault()`, so you will need to configure the system properly to locate your key and trust material.

> **Note:**
>
> If you use the standard SSL/TLS-based RMI socket factories, then you can specify the key stores with system properties:
>
> ```
> -Djavax.net.ssl.keyStore=testkeys
> -Djavax.net.ssl.keyStorePassword=passphrase
> ```

HelloClient.java

```
import java.net.InetAddress;
import java.rmi.RemoteException;
import java.rmi.registry.LocateRegistry;
import java.rmi.registry.Registry;

public class HelloClient {

    private static final int PORT = 2019;

    public static void main(String args[]) {
```

```
        try {
            // Make reference to SSL-based registry
            Registry registry = LocateRegistry.getRegistry(
                InetAddress.getLocalHost().getHostName(), PORT,
                new RMISSLClientSocketFactory());

            // "obj" is the identifier that we'll use to refer
            // to the remote object that implements the "Hello"
            // interface
            Hello obj = (Hello) registry.lookup("HelloServer");

            String message = "blank";
            message = obj.sayHello();
            System.out.println(message+"\n");
        } catch (Exception e) {
            System.out.println("HelloClient exception: " + e.getMessage());
            e.printStackTrace();
        }
    }
}
```

Hello.java

```
import java.rmi.Remote;
import java.rmi.RemoteException;

public interface Hello extends Remote {
    String sayHello() throws RemoteException;
}
```

HelloImpl.java

```
import java.io.*;
import java.net.InetAddress;
import java.rmi.RemoteException;
import java.rmi.RMISecurityManager;
import java.rmi.registry.LocateRegistry;
import java.rmi.registry.Registry;
import java.rmi.server.UnicastRemoteObject;

public class HelloImpl extends UnicastRemoteObject implements Hello {

    private static final int PORT = 2019;

    public HelloImpl() throws Exception {
        super(PORT,
              new RMISSLClientSocketFactory(),
              new RMISSLServerSocketFactory());
    }

    public String sayHello() {
        return "Hello World!";
    }
```

```
    public static void main(String args[]) {

        // Create and install a security manager
        if (System.getSecurityManager() == null) {
            System.setSecurityManager(new RMISecurityManager());
        }

        try {
            // Create SSL-based registry
            Registry registry = LocateRegistry.createRegistry(PORT,
                new RMISSLClientSocketFactory(),
                new RMISSLServerSocketFactory());

            HelloImpl obj = new HelloImpl();

            // Bind this object instance to the name "HelloServer"
            registry.bind("HelloServer", obj);

            System.out.println("HelloServer bound in registry");
        } catch (Exception e) {
            System.out.println("HelloImpl err: " + e.getMessage());
            e.printStackTrace();
        }
    }
}
```

RMISSLClientSocketFactory.java

```
import java.io.*;
import java.net.*;
import java.rmi.server.*;
import javax.net.ssl.*;

public class RMISSLClientSocketFactory
        implements RMIClientSocketFactory, Serializable {

    public Socket createSocket(String host, int port) throws IOException {
            SSLSocketFactory factory =
                (SSLSocketFactory)SSLSocketFactory.getDefault();
            SSLSocket socket = (SSLSocket)factory.createSocket(host, port);
            return socket;
    }

    public int hashCode() {
        return getClass().hashCode();
    }

    public boolean equals(Object obj) {
        if (obj == this) {
            return true;
        } else if (obj == null || getClass() != obj.getClass()) {
            return false;
        }
        return true;
```

```
    }
}
```

RMISSLServerSocketFactory.java

```
import java.io.*;
import java.net.*;
import java.rmi.server.*;
import javax.net.ssl.*;
import java.security.KeyStore;
import javax.net.ssl.*;

public class RMISSLServerSocketFactory implements RMIServerSocketFactory {

    /*
     * Create one SSLServerSocketFactory, so we can reuse sessions
     * created by previous sessions of this SSLContext.
     */
    private SSLServerSocketFactory ssf = null;

    public RMISSLServerSocketFactory() throws Exception {
        try {
            // set up key manager to do server authentication
            SSLContext ctx;
            KeyManagerFactory kmf;
            KeyStore ks;

            char[] passphrase = "passphrase".toCharArray();
            ks = KeyStore.getInstance("JKS");
            ks.load(new FileInputStream("testkeys"), passphrase);

            kmf = KeyManagerFactory.getInstance("SunX509");
            kmf.init(ks, passphrase);

            ctx = SSLContext.getInstance("TLS");
            ctx.init(kmf.getKeyManagers(), null, null);

            ssf = ctx.getServerSocketFactory();
        } catch (Exception e) {
            e.printStackTrace();
            throw e;
        }
    }

    public ServerSocket createServerSocket(int port) throws IOException {
            return ssf.createServerSocket(port);
    }

    public int hashCode() {
        return getClass().hashCode();
    }

    public boolean equals(Object obj) {
        if (obj == this) {
```

```
            return true;
        } else if (obj == null || getClass() != obj.getClass()) {
            return false;
        }
        return true;
    }
}
```

policy

```
// In this example, for simplicity, we will use a policy file that
// gives global permission to anyone from anywhere. Do not use this
// policy file in a production environment.

grant {
    permission java.net.SocketPermission "*", "accept,resolve";
};
```

Sample Code Illustrating the Use of an SSLEngine

`SSLEngine` gives application developers flexibility when choosing I/O and compute strategies. Rather than tie the SSL/TLS implementation to a specific I/O abstraction (such as single-threaded `SSLSockets`), `SSLEngine` removes the I/O and compute constraints from the SSL/TLS implementation.

As mentioned earlier, `SSLEngine` is an advanced API, and is not appropriate for casual use. An introductory example is provided that helps illustrate its use. It removes most of the I/O and threading issues, and focuses on many of the `SSLEngine` methods.

Running SSLEngineSimpleDemo

The example `SSLEngineSimpleDemo.java` is a very simple application that focuses on the operation of the `SSLEngine` while simplifying the I/O and threading issues. This application creates two `SSLEngine` objects that exchange SSL/TLS messages via common `ByteBuffer` objects. A single loop serially performs all of the engine operations and demonstrates how a secure connection is established (handshaking), how application data is transferred, and how the engine is closed.

The `SSLEngineResult` provides a great deal of information about the current state of the `SSLEngine`. This example does not examine all of the states. It simplifies the I/O and threading issues to the point that this is not a good example for a production environment; nonetheless, it is useful to demonstrate the overall function of the `SSLEngine`.

Usage

```
java SSLEngineSimpleDemo
```

SSLEngineSimpleDemo.java

```
/**
 * A SSLEngine usage example which simplifies the presentation
 * by removing the I/O and multi-threading concerns.
 *
 * The demo creates two SSLEngines, simulating a client and server.
 * The "transport" layer consists two ByteBuffers:  think of them
 * as directly connected pipes.
 *
 * Note, this is a *very* simple example: real code will be much more
 * involved.  For example, different threading and I/O models could be
 * used, transport mechanisms could close unexpectedly, and so on.
 *
 * When this application runs, notice that several messages
 * (wrap/unwrap) pass before any application data is consumed or
 * produced.  (For more information, please see the SSL/TLS
 * specifications.)  There may several steps for a successful handshake,
 * so it's typical to see the following series of operations:
 *
 *      client          server          message
 *      ======          ======          =======
 *      wrap()          ...             ClientHello
 *      ...             unwrap()        ClientHello
 *      ...             wrap()          ServerHello/Certificate
 *      unwrap()        ...             ServerHello/Certificate
 *      wrap()          ...             ClientKeyExchange
 *      wrap()          ...             ChangeCipherSpec
 *      wrap()          ...             Finished
 *      ...             unwrap()        ClientKeyExchange
 *      ...             unwrap()        ChangeCipherSpec
 *      ...             unwrap()        Finished
 *      ...             wrap()          ChangeCipherSpec
 *      ...             wrap()          Finished
 *      unwrap()        ...             ChangeCipherSpec
 *      unwrap()        ...             Finished
 */

import javax.net.ssl.*;
import javax.net.ssl.SSLEngineResult.*;
import java.io.*;
import java.security.*;
import java.nio.*;

public class SSLEngineSimpleDemo {

    /*
     * Enables logging of the SSLEngine operations.
     */
    private static boolean logging = true;

    /*
     * Enables the JSSE system debugging system property:
     *
```

```
     *     -Djavax.net.debug=all
     *
     * This gives a lot of low-level information about operations underway,
     * including specific handshake messages, and might be best examined
     * after gaining some familiarity with this application.
     */
    private static boolean debug = false;

    private SSLContext sslc;

    private SSLEngine clientEngine;     // client Engine
    private ByteBuffer clientOut;       // write side of clientEngine
    private ByteBuffer clientIn;        // read side of clientEngine

    private SSLEngine serverEngine;     // server Engine
    private ByteBuffer serverOut;       // write side of serverEngine
    private ByteBuffer serverIn;        // read side of serverEngine

    /*
     * For data transport, this example uses local ByteBuffers.  This
     * isn't really useful, but the purpose of this example is to show
     * SSLEngine concepts, not how to do network transport.
     */
    private ByteBuffer cTOs;            // "reliable" transport client-
>server
    private ByteBuffer sTOc;            // "reliable" transport server-
>client

    /*
     * The following is to set up the keystores.
     */
    private static String keyStoreFile = "testkeys";
    private static String trustStoreFile = "testkeys";
    private static String passwd = "passphrase";

    /*
     * Main entry point for this demo.
     */
    public static void main(String args[]) throws Exception {
        if (debug) {
            System.setProperty("javax.net.debug", "all");
        }

        SSLEngineSimpleDemo demo = new SSLEngineSimpleDemo();
        demo.runDemo();

        System.out.println("Demo Completed.");
    }

    /*
     * Create an initialized SSLContext to use for this demo.
     */
    public SSLEngineSimpleDemo() throws Exception {

        KeyStore ks = KeyStore.getInstance("JKS");
```

```
        KeyStore ts = KeyStore.getInstance("JKS");

        char[] passphrase = "passphrase".toCharArray();

        ks.load(new FileInputStream(keyStoreFile), passphrase);
        ts.load(new FileInputStream(trustStoreFile), passphrase);

        KeyManagerFactory kmf = KeyManagerFactory.getInstance("SunX509");
        kmf.init(ks, passphrase);

        TrustManagerFactory tmf =
TrustManagerFactory.getInstance("SunX509");
        tmf.init(ts);

        SSLContext sslCtx = SSLContext.getInstance("TLS");

        sslCtx.init(kmf.getKeyManagers(), tmf.getTrustManagers(), null);

        sslc = sslCtx;
    }

    /*
     * Run the demo.
     *
     * Sit in a tight loop, both engines calling wrap/unwrap regardless
     * of whether data is available or not.  We do this until both engines
     * report back they are closed.
     *
     * The main loop handles all of the I/O phases of the SSLEngine's
     * lifetime:
     *
     *     initial handshaking
     *     application data transfer
     *     engine closing
     *
     * One could easily separate these phases into separate
     * sections of code.
     */
    private void runDemo() throws Exception {
        boolean dataDone = false;

        createSSLEngines();
        createBuffers();

        SSLEngineResult clientResult;   // results from client's last
operation
        SSLEngineResult serverResult;   // results from server's last
operation

        /*
         * Examining the SSLEngineResults could be much more involved,
         * and may alter the overall flow of the application.
         *
         * For example, if we received a BUFFER_OVERFLOW when trying
         * to write to the output pipe, we could reallocate a larger
```

```
             * pipe, but instead we wait for the peer to drain it.
             */
            while (!isEngineClosed(clientEngine) ||
                    !isEngineClosed(serverEngine)) {

                log("================");

                clientResult = clientEngine.wrap(clientOut, cTOs);
                log("client wrap: ", clientResult);
                runDelegatedTasks(clientResult, clientEngine);

                serverResult = serverEngine.wrap(serverOut, sTOc);
                log("server wrap: ", serverResult);
                runDelegatedTasks(serverResult, serverEngine);

                cTOs.flip();
                sTOc.flip();

                log("----");

                clientResult = clientEngine.unwrap(sTOc, clientIn);
                log("client unwrap: ", clientResult);
                runDelegatedTasks(clientResult, clientEngine);

                serverResult = serverEngine.unwrap(cTOs, serverIn);
                log("server unwrap: ", serverResult);
                runDelegatedTasks(serverResult, serverEngine);

                cTOs.compact();
                sTOc.compact();

                /*
                 * After we've transfered all application data between the
client
                 * and server, we close the clientEngine's outbound stream.
                 * This generates a close_notify handshake message, which the
                 * server engine receives and responds by closing itself.
                 *
                 * In normal operation, each SSLEngine should call
                 * closeOutbound().  To protect against truncation attacks,
                 * SSLEngine.closeInbound() should be called whenever it has
                 * determined that no more input data will ever be
                 * available (say a closed input stream).
                 */
                if (!dataDone && (clientOut.limit() == serverIn.position()) &&
                        (serverOut.limit() == clientIn.position())) {

                    /*
                     * A sanity check to ensure we got what was sent.
                     */
                    checkTransfer(serverOut, clientIn);
                    checkTransfer(clientOut, serverIn);

                    log("\tClosing clientEngine's *OUTBOUND*...");
                    clientEngine.closeOutbound();
```

```
            // serverEngine.closeOutbound();
            dataDone = true;
        }
    }
}

/*
 * Using the SSLContext created during object creation,
 * create/configure the SSLEngines we'll use for this demo.
 */
private void createSSLEngines() throws Exception {
    /*
     * Configure the serverEngine to act as a server in the SSL/TLS
     * handshake.  Also, require SSL client authentication.
     */
    serverEngine = sslc.createSSLEngine();
    serverEngine.setUseClientMode(false);
    serverEngine.setNeedClientAuth(true);

    /*
     * Similar to above, but using client mode instead.
     */
    clientEngine = sslc.createSSLEngine("client", 80);
    clientEngine.setUseClientMode(true);
}

/*
 * Create and size the buffers appropriately.
 */
private void createBuffers() {

    /*
     * We'll assume the buffer sizes are the same
     * between client and server.
     */
    SSLSession session = clientEngine.getSession();
    int appBufferMax = session.getApplicationBufferSize();
    int netBufferMax = session.getPacketBufferSize();

    /*
     * We'll make the input buffers a bit bigger than the max needed
     * size, so that unwrap()s following a successful data transfer
     * won't generate BUFFER_OVERFLOWS.
     *
     * We'll use a mix of direct and indirect ByteBuffers for
     * tutorial purposes only.  In reality, only use direct
     * ByteBuffers when they give a clear performance enhancement.
     */
    clientIn = ByteBuffer.allocate(appBufferMax + 50);
    serverIn = ByteBuffer.allocate(appBufferMax + 50);

    cTOs = ByteBuffer.allocateDirect(netBufferMax);
    sTOc = ByteBuffer.allocateDirect(netBufferMax);

    clientOut = ByteBuffer.wrap("Hi Server, I'm Client".getBytes());
```

```
    serverOut = ByteBuffer.wrap("Hello Client, I'm Server".getBytes());
}

/*
 * If the result indicates that we have outstanding tasks to do,
 * go ahead and run them in this thread.
 */
private static void runDelegatedTasks(SSLEngineResult result,
        SSLEngine engine) throws Exception {

    if (result.getHandshakeStatus() == HandshakeStatus.NEED_TASK) {
        Runnable runnable;
        while ((runnable = engine.getDelegatedTask()) != null) {
            log("\trunning delegated task...");
            runnable.run();
        }
        HandshakeStatus hsStatus = engine.getHandshakeStatus();
        if (hsStatus == HandshakeStatus.NEED_TASK) {
            throw new Exception(
                "handshake shouldn't need additional tasks");
        }
        log("\tnew HandshakeStatus: " + hsStatus);
    }
}

private static boolean isEngineClosed(SSLEngine engine) {
    return (engine.isOutboundDone() && engine.isInboundDone());
}

/*
 * Simple check to make sure everything came across as expected.
 */
private static void checkTransfer(ByteBuffer a, ByteBuffer b)
        throws Exception {
    a.flip();
    b.flip();

    if (!a.equals(b)) {
        throw new Exception("Data didn't transfer cleanly");
    } else {
        log("\tData transferred cleanly");
    }

    a.position(a.limit());
    b.position(b.limit());
    a.limit(a.capacity());
    b.limit(b.capacity());
}

/*
 * Logging code
 */
private static boolean resultOnce = true;

private static void log(String str, SSLEngineResult result) {
```

```
        if (!logging) {
            return;
        }
        if (resultOnce) {
            resultOnce = false;
            System.out.println("The format of the SSLEngineResult is: \n" +
                "\t\"getStatus() / getHandshakeStatus()\" +\n" +
                "\t\"bytesConsumed() / bytesProduced()\"\n");
        }
        HandshakeStatus hsStatus = result.getHandshakeStatus();
        log(str +
            result.getStatus() + "/" + hsStatus + ", " +
            result.bytesConsumed() + "/" + result.bytesProduced() +
            " bytes");
        if (hsStatus == HandshakeStatus.FINISHED) {
            log("\t...ready for application data");
        }
    }

    private static void log(String str) {
        if (logging) {
            System.out.println(str);
        }
    }
}
```

Troubleshooting JSSE Sample Code

One of the most common problems people have in using JSSE is when the JSSE receives a certificate that is unknown to the mechanism that makes trust decisions. If an unknown certificate is received, the trust mechanism will throw an exception saying that the certificate is untrusted. Make sure that the correct trust store is being used, and that the JSSE is installed and configured correctly.

If you are using the "localhost" or "duke" credentials, be sure that you have correctly specified the location of the samplecacerts file, otherwise your application will not work. (See Sample Truststores for more information.)

The SSL debug mechanism can be used to investigate such trust problems. See the implementation documentation for more information about this subject.

10

Java PKI Programmer's Guide

The Java Certification Path API consists of classes and interfaces for handling certification paths, which are also called certification chains. If a certification path meets certain validation rules, it may be used to securely establish the mapping of a public key to a subject.

Topics

PKI Programmer's Guide Overview

Core Classes and Interfaces

Implementing a Service Provider

Appendix A: Standard Names

Appendix B: CertPath Implementation in SUN Provider

Appendix C: OCSP Support

Appendix D: CertPath Implementation in JdkLDAP Provider

Appendix E: Disabling Cryptographic Algorithms

PKI Programmer's Guide Overview

The Java Certification Path API defines interfaces and abstract classes for creating, building, and validating certification paths. Implementations may be plugged in using a provider-based interface.

This API is based on the Cryptographic Service Providers architecture, described in the *Java Cryptography Architecture Reference Guide*, and includes algorithm-specific classes for building and validating X.509 certification paths according to the PKIX standards. The PKIX standards were developed by the IETF PKIX working group.

This API was originally specified using the Java Community Process program as Java Specification Request (JSR) 000055. The API was included in the Java SDK, starting with Java SE Development Kit (JDK) 1.4. See JSR 55: Certification Path API.

Who Should Read This Document

This document is intended for two types of experienced developers:

1. Those who want to design secure applications that build or validate certification paths.
2. Those who want to write a service provider implementation for building or validating certification paths.

This document assumes that you have already read Cryptographic Service Providers.

Introduction to Public Key Certificates

Users of public key applications and systems must be confident that a subject's public key is genuine, i.e., that the associated private key is owned by the subject. Public key certificates are used to establish this trust.

A **public key (or identity) certificate** is a binding of a public key to an identity, which is digitally signed by the private key of another entity, often called a **Certification Authority** (CA). For the remainder of this section, the term CA is used to refer to an entity that signs a certificate.

If the user does not have a trusted copy of the public key of the CA that signed the subject's public key certificate, then another public key certificate vouching for the signing CA is required. This logic can be applied recursively, until a chain of certificates (or **a certification path**) is discovered from a **trust anchor** or a **most-trusted CA** to the target subject (commonly referred to as the **end-entity**). The **most-trusted CA** is usually specified by a certificate issued to a CA that the user directly trusts. In general, a certification path is an ordered list of certificates, usually comprised of the end-entity's public key certificate and zero or more additional certificates. A certification path typically has one or more encodings, allowing it to be safely transmitted across networks and to different operating system architectures.

The following figure illustrates a certification path from a most-trusted CA's public key (CA 1) to the target subject (Alice). The certification path establishes trust in Alice's public key through an intermediate CA named CA2.

Figure 10-1 Certification Path from CA's Public Key (CA 1) to the Target Subject

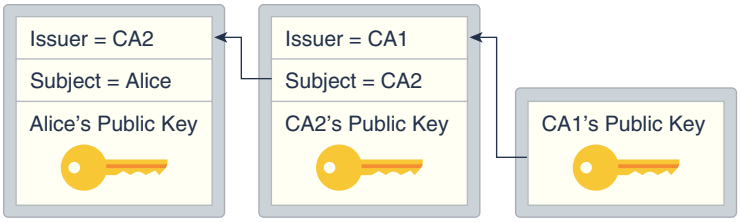

A certification path must be validated before it can be relied on to establish trust in a subject's public key. Validation can consist of various checks on the certificates contained in the certification path, such as verifying the signatures and checking that each certificate has not been revoked. The PKIX standards define an algorithm for validating certification paths consisting of X.509 certificates.

Often a user may not have a certification path from a most-trusted CA to the subject. Providing services to build or discover certification paths is an important feature of public key enabled systems. RFC 2587 defines an LDAP (Lightweight Directory Access Protocol) schema definition that facilitates the discovery of X.509 certification paths using the LDAP directory service protocol.

Building and validating certification paths is an important part of many standard security protocols such as SSL/TLS/DTLS, S/MIME, and IPsec. The Java Certification Path API provides a set of classes and interfaces for developers who need to integrate this functionality into their applications. This API benefits two types of developers: those who need to write service provider implementations for a specific certification

path building or validation algorithm; and those who need to access standard algorithms for creating, building, and validating certification paths in an implementation-independent manner.

X.509 Certificates and Certificate Revocation Lists (CRLs)

A public-key certificate is a digitally signed statement from one entity saying that the public key and some other information of another entity has some specific value.

The following table defines some of the key terms:

Public Keys
These are numbers associated with a particular entity, and are intended to be known to everyone who needs to have trusted interactions with that entity. Public keys are used to verify signatures.

Digitally Signed
If some data is *digitally signed* it has been stored with the "identity" of an entity, and a signature that proves that entity knows about the data. The data is rendered unforgeable by signing with the entitys' private key.

Identity
A known way of addressing an entity. In some systems the identity is the public key, in others it can be anything from a UNIX UID to an Email address to an X.509 Distinguished Name.

Signature
A signature is computed over some data using the private key of an entity (the signer).

Private Keys
These are numbers, each of which is supposed to be known only to the particular entity whose private key it is (that is, it's supposed to be kept secret). Private and public keys exist in pairs in all public key cryptography systems (also referred to as "public key crypto systems"). In a typical public key crypto system, such as DSA, a private key corresponds to exactly one public key. Private keys are used to compute signatures.

Entity
An entity is a person, organization, program, computer, business, bank, or something else you are trusting to some degree.

Basically, public key cryptography requires access to users' public keys. In a large-scale networked environment it is impossible to guarantee that prior relationships between communicating entities have been established or that a trusted repository exists with all used public keys. Certificates were invented as a solution to this public key distribution problem. Now a *Certification Authority* (CA) can act as a *Trusted Third Party*. CAs are entities (e.g., businesses) that are trusted to sign (issue) certificates for other entities. It is assumed that CAs will only create valid and reliable certificates as they are bound by legal agreements. There are many public Certification Authorities, such as VeriSign, Thawte, Entrust, and so on. You can also run your own Certification Authority using products such as the Netscape/Microsoft Certificate Servers or the Entrust CA product for your organization.

What Applications use Certificates?

Probably the most widely visible application of X.509 certificates today is in web browsers (such as Mozilla Firefox and Microsoft Internet Explorer) that support the TLS protocol. TLS (Transport Layer Security) is a security protocol that provides privacy and authentication for your network traffic. These browsers can only use this protocol with web servers that support TLS.

Other technologies that rely on X.509 certificates include:

- Various code-signing schemes, such as signed Java ARchives, and Microsoft Authenticode.
- Various secure E-Mail standards, such as PEM and S/MIME.

How do I Get a Certificate?

There are two basic techniques used to get certificates:

- You can create one yourself (using the right tools, such as keytool).
- You can ask a Certification Authority to issue you one (either directly or using a tool such as **keytool** to generate the request).

The main inputs to the certificate creation process are:

- Matched *public and private keys*, generated using some special tools (such as keytool), or a browser. **Only the public key is ever shown to anyone else.** The private key is used to sign data; if someone knows your private key, they can masquerade as you ... perhaps forging legal documents attributed to you!
- You need to provide *information about the entity being certified* (e.g., you). This normally includes information such as your name and organizational address. If you ask a CA to issue a certificate for you, you will normally need to provide proof to show correctness of the information.

If you are asking a CA to issue you a certificate, you provide your public key and some information about you. You'll use a tool (such as keytool or a browser that supports Certificate Signing Request generation). to digitally sign this information, and send it to the CA. The CA will then generate the certificate and return it.

If you're generating the certificate yourself, you'll take that same information, add a little more (dates during which the certificate is valid, a serial number), and just create the certificate using some tool (such as keytool). Not everyone will accept self-signed certificates; one part of the value provided by a CA is to serve as a neutral and trusted introduction service, based in part on their verification requirements, which are openly published in their Certification Service Practices (CSP).

What's Inside an X.509 Certificate?

The X.509 standard defines what information can go into a certificate, and describes how to write it down (the data format). All X.509 certificates have the following data, in addition to the signature:

Version
This identifies which version of the X.509 standard applies to this certificate, which affects what information can be specified in it. Thus far, three versions are defined.

Serial Number
The entity that created the certificate is responsible for assigning it a serial number to distinguish it from other certificates it issues. This information is used in numerous ways, for example when a certificate is revoked its serial number is placed in a Certificate Revocation List (CRL).

Signature Algorithm Identifier
This identifies the algorithm used by the CA to sign the certificate.

Issuer Name
The X.500 name of the entity that signed the certificate. This is normally a CA. Using this certificate implies trusting the entity that signed this certificate. (Note that in some cases, such as *root or top-level* CA certificates, the issuer signs its own certificate.)

Validity Period
Each certificate is valid only for a limited amount of time. This period is described by a start date and time and an end date and time, and can be as short as a few seconds or almost as long as a century. The validity period chosen depends on a number of factors, such as the strength of the private key used to sign the certificate or the amount one is willing to pay for a certificate. This is the expected period that entities can rely on the public value, if the associated private key has not been compromised.

Subject Name
The name of the entity whose public key the certificate identifies. This name uses the X.500 standard, so it is intended to be unique across the Internet. This is the Distinguished Name (DN) of the entity, for example,

```
CN=Java Duke, OU=Java Software Division, O=Sun Microsystems Inc, C=US
```

(These refer to the subject's Common Name, Organizational Unit, Organization, and Country.)

Subject Public Key Information
This is the public key of the entity being named, together with an algorithm identifier which specifies which public key crypto system this key belongs to and any associated key parameters.

X.509 Version 1 has been available since 1988, is widely deployed, and is the most generic.

X.509 Version 2 introduced the concept of subject and issuer unique identifiers to handle the possibility of reuse of subject and/or issuer names over time. Most certificate profile documents strongly recommend that names not be reused, and that certificates should not make use of unique identifiers. Version 2 certificates are not widely used.

X.509 Version 3 is the most recent (1996) and supports the notion of extensions, whereby anyone can define an extension and include it in the certificate. Some common extensions in use today are: *KeyUsage* (limits the use of the keys to particular purposes such as "signing-only") and *AlternativeNames* (allows other identities to also be associated with this public key, e.g. DNS names, Email addresses, IP addresses). Extensions can be marked *critical* to indicate that the extension should be checked and enforced/used. For example, if a certificate has the KeyUsage extension marked critical and set to "keyCertSign" then if this certificate is presented during SSL communication, it should be rejected, as the certificate extension indicates

that the associated private key should only be used for signing certificates and not for SSL use.

All the data in a certificate is encoded using two related standards called ASN.1/DER. *Abstract Syntax Notation 1* describes data. The *Distinguished Encoding Rules* describe a single way to store and transfer that data.

What Java API Can Be Used to Access and Manage Certificates?

The Certificate API, found in the `java.security.cert` package, includes the following:

- **CertificateFactory** class defines the functionality of a certificate factory, which is used to generate certificate, certificate revocation list (CRL), and certification path objects from their encoding.
- **Certificate** class is an abstract class for managing a variety of certificates. It is an abstraction for certificates that have different formats but important common uses. For example, different types of certificates, such as X.509 and PGP, share general certificate functionality (like encoding and verifying) and some types of information like public key.
- **CRL** class is an abstract class for managing a variety of Certificate Revocation Lists (CRLs).
- **X509Certificate** class is an abstract class for X.509 Certificates. It provides a standard way to access all the attributes of an X.509 certificate.
- **X509Extension** interface is an interface for an X.509 extension. The extensions defined for X.509 v3 certificates and v2 CRLs (Certificate Revocation Lists) provide mechanisms for associating additional attributes with users or public keys, such as for managing the certification hierarchy, and for managing CRL distribution.
- **X509CRL** class is an abstract class for an X.509 Certificate Revocation List (CRL). A CRL is a time-stamped list identifying revoked certificates. It is signed by a Certification Authority (CA) and made freely available in a public repository.
- **X509CRLEntry** class is an abstract class for a CRL entry.

What Java Tool Can Generate, Display, Import, and Export X.509 Certificates?

There is a tool named keytool that can be used to create public/private key pairs and X.509 v3 certificates, and to manage keystores. Keys and certificates are used to digitally sign your Java applications and applets (see jarsigner).

A *keystore* is a protected database that holds keys and certificates. Access to a keystore is guarded by a password (defined at the time the keystore is created, by the person who creates the keystore, and changeable only when providing the current password). In addition, each private key in a keystore can be guarded by its own password.

Using **keytool**, it is possible to display, import, and export X.509 v1, v2, and v3 certificates stored as files, and to generate new v3 certificates. For examples, see the "EXAMPLES" section for keytool in the *Java Platform, Standard Edition Tools Reference*.

Core Classes and Interfaces

The core classes of the Java Certification Path API consist of interfaces and classes that support certification path functionality in an algorithm and implementation-independent manner.

The API builds on and extends the existing `java.security.cert` package for handling certificates. The core classes can be broken up into 4 class categories: Basic, Validation, Building, and Storage:

- Basic Certification Path Classes
  - `CertPath`, `CertificateFactory`, and `CertPathParameters`
- Certification Path Validation Classes
  - `CertPathValidator`, `CertPathValidatorResult`, and `CertPathChecker`
- Certification Path Building Classes
  - `CertPathBuilder`, and `CertPathBuilderResult`
- Certificate/CRL Storage Classes
  - `CertStore`, `CertStoreParameters`, `CertSelector`, and `CRLSelector`

The Java Certification Path API also includes a set of algorithm-specific classes modeled for use with the PKIX certification path validation algorithm defined in RFC 5280: *Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. The PKIX Classes are:

- `TrustAnchor`
- `PKIXParameters`
- `PKIXCertPathValidatorResult`
- `PKIXBuilderParameters`
- `PKIXCertPathBuilderResult`
- `PKIXCertPathChecker`
- `PKIXRevocationChecker`

The complete reference documentation for the relevant Certification Path API classes can be found in `java.security.cert` .

Most of the classes and interfaces in the `CertPath` API are not thread-safe. However, there are some exceptions, which will be noted in this guide and in the API specification. Multiple threads that need to access a single non-thread-safe object concurrently should synchronize amongst themselves and provide the necessary locking. Multiple threads each manipulating separate objects need not synchronize.

Topics

Basic Certification Path Classes

Certification Path Validation Classes

Certification Path Building Classes

Certificate/CRL Storage Classes

PKIX Classes

Basic Certification Path Classes

The basic certification path classes provide fundamental functionality for encoding and representing certification paths. The key basic class in the Java Certification Path API is `CertPath`, which encapsulates the universal aspects shared by all types of certification paths. An application uses an instance of the `CertificateFactory` class to create a `CertPath` object.

Topics

The CertPath Class

The CertificateFactory Class

The CertPathParameters Interface

The CertPath Class

The `CertPath` class is an abstract class for certification paths. It defines the functionality shared by all certification path objects. Various certification path types can be implemented by subclassing the `CertPath` class, even though they may have different contents and ordering schemes.

All `CertPath` objects are serializable, immutable and thread-safe and share the following characteristics:

- A type

  This corresponds to the type of the certificates in the certification path, for example: X.509. The type of a `CertPath` is obtained using the method:

  ```
  public String getType()
  ```

  For standard certificate types, see CertificateFactory Types.

- A list of certificates

  The `getCertificates` method returns the list of certificates in the certification path:

  ```
  public abstract List<? extends Certificate> getCertificates()
  ```

  This method returns a `List` of zero or more `java.security.cert.Certificate` objects. The returned `List` and the `Certificates` contained within it are immutable, in order to protect the contents of the `CertPath` object. The ordering of the certificates returned depends on the type. By convention, the certificates in a `CertPath` object of type X.509 are ordered starting with the target certificate and ending with a certificate issued by the trust anchor. That is, the issuer of one certificate is the subject of the following one. The certificate representing the `TrustAnchor` should not be included in the certification path. Unvalidated X.509 `CertPaths` may not follow this convention. PKIX CertPathValidators will detect any

departure from these conventions that cause the certification path to be invalid and throw a `CertPathValidatorException`.

- One or more encodings

  Each `CertPath` object supports one or more encodings. These are external encoded forms for the certification path, used when a standard representation of the path is needed outside the Java Virtual Machine (as when transmitting the path over a network to some other party). Each path can be encoded in a default format, the bytes of which are returned using the method:

  ```
  public abstract byte[] getEncoded()
  ```

  Alternatively, the `getEncoded(String)` method returns a specific supported encoding by specifying the encoding format as a `String` (ex: "PKCS7"). For standard encoding formats, see CertPath Encodings.

  ```
  public abstract byte[] getEncoded(String encoding)
  ```

  Also, the `getEncodings` method returns an iterator over the supported encoding format `Strings` (the default encoding format is returned first):

  ```
  public abstract Iterator<String> getEncodings()
  ```

All `CertPath` objects are also `Serializable`. `CertPath` objects are resolved into an alternate CertPath.CertPathRep object during serialization. This allows a `CertPath` object to be serialized into an equivalent representation regardless of its underlying implementation.

`CertPath` objects are generated from an encoded byte array or list of `Certificates` using a CertificateFactory. Alternatively, a `CertPathBuilder` may be used to try to find a `CertPath` from a most-trusted CA to a particular subject. Once a `CertPath` object has been created, it may be validated by passing it to the `validate` method of `CertPathValidator`. Each of these concepts are explained in more detail in subsequent sections.

The CertificateFactory Class

The `CertificateFactory` class is an engine class that defines the functionality of a certificate factory. It is used to generate `Certificate`, `CRL`, and `CertPath` objects.

A `CertificateFactory` should not be confused with a `CertPathBuilder`. A CertPathBuilder (discussed later) is used to discover or find a certification path when one does not exist. In contrast, a `CertificateFactory` is used when a certification path has already been discovered and the caller needs to instantiate a CertPath object from its contents, which exist in a different form such as an encoded byte array or an array of `Certificates`.

Creating a CertificateFactory Object

See the CertificateFactory section in the *Java Cryptography Architecture Reference Guide* for the details of creating a `CertificateFactory` object.

Generating CertPath Objects

A `CertificateFactory` instance generates `CertPath` objects from a `List` of `Certificate` objects or from an `InputStream` that contains the encoded form of a `CertPath`. Just like a `CertPath`, each `CertificateFactory` supports a default encoding format for certification paths (ex: PKCS#7). To generate a `CertPath` object and initialize it with the data read from an input stream (in the default encoding format), use the `generateCertPath` method:

```
public final CertPath generateCertPath(InputStream inStream)
```

or from a particular encoding format:

```
    public final CertPath generateCertPath(InputStream inStream,
                                           String encoding)
```

To find out what encoding formats are supported, use the `getCertPathEncodings` method (the default encoding is returned first):

```
public final Iterator<String> getCertPathEncodings()
```

To generate a certification path object from a `List` of `Certificate` objects, use the following method:

```
public final CertPath generateCertPath(List<? extends Certificate>
certificates)
```

A `CertificateFactory` always returns `CertPath` objects that consist of `Certificates` that are of the same type as the factory. For example, a `CertificateFactory` of type X.509 returns `CertPath` objects consisting of certificates that are an instance of `java.security.cert.X509Certificate`.

The following code sample illustrates generating a certification path from a PKCS#7 encoded certificate reply stored in a file:

```
    // open an input stream to the file
    FileInputStream fis = new FileInputStream(filename);
    // instantiate a CertificateFactory for X.509
    CertificateFactory cf = CertificateFactory.getInstance("X.509");
    // extract the certification path from
    // the PKCS7 SignedData structure
    CertPath cp = cf.generateCertPath(fis, "PKCS7");
    // print each certificate in the path
    List<Certificate> certs = cp.getCertificates();
    for (Certificate cert : certs) {
        System.out.println(cert);
    }
```

Here's another code sample that fetches a certificate chain from a `KeyStore` and converts it to a `CertPath` using a `CertificateFactory`:

```
// instantiate a KeyStore with type JKS
KeyStore ks = KeyStore.getInstance("JKS");
// load the contents of the KeyStore
ks.load(new FileInputStream("./keystore"),
    "password".toCharArray());
// fetch certificate chain stored with alias "sean"
Certificate[] certArray = ks.getCertificateChain("sean");
// convert chain to a List
List certList = Arrays.asList(certArray);
// instantiate a CertificateFactory for X.509
CertificateFactory cf = CertificateFactory.getInstance("X.509");
// extract the certification path from
// the List of Certificates
CertPath cp = cf.generateCertPath(certList);
```

Note that there is an existing method in `CertificateFactory` named `generateCertificates` that parses a sequence of `Certificates`. For encodings consisting of multiple certificates, use `generateCertificates` when you want to parse a collection of possibly unrelated certificates. Otherwise, use `generateCertPath` when you want to generate a `CertPath` and subsequently validate it with a `CertPathValidator` (discussed later).

The CertPathParameters Interface

The `CertPathParameters` interface is a transparent representation of the set of parameters used with a particular certification path builder or validation algorithm.

Its main purpose is to group (and provide type safety for) all certification path parameter specifications. The `CertPathParameters` interface extends the `Cloneable` interface and defines a `clone()` method that does not throw an exception. All concrete implementations of this interface should implement and override the `Object.clone()` method, if necessary. This allows applications to clone any `CertPathParameters` object.

Objects implementing the `CertPathParameters` interface are passed as arguments to methods of the CertPathValidator and CertPathBuilder classes. Typically, a concrete implementation of the `CertPathParameters` interface will hold a set of input parameters specific to a particular certification path build or validation algorithm. For example, the `PKIXParameters` class is an implementation of the `CertPathParameters` interface that holds a set of input parameters for the PKIX certification path validation algorithm. One such parameter is the set of most-trusted CAs that the caller trusts for anchoring the validation process. This parameter among others is discussed in more detail in the section discussing the PKIXParameters class.

Certification Path Validation Classes

The Java Certification Path API includes classes and interfaces for validating certification paths. An application uses an instance of the `CertPathValidator` class to

validate a `CertPath` object. If successful, the result of the validation algorithm is returned in an object implementing the `CertPathValidatorResult` interface.

Topics

The CertPathValidator Class

The CertPathValidatorResult Interface

The CertPathValidator Class

The `CertPathValidator` class is an engine class used to validate a certification path.

Creating a CertPathValidator Object

As with all engine classes, the way to get a `CertPathValidator` object for a particular validation algorithm is to call one of the `getInstance` static factory methods on the `CertPathValidator` class:

```
public static CertPathValidator getInstance(String algorithm)
public static CertPathValidator getInstance(String algorithm,
                                            String provider)
public static CertPathValidator getInstance(String algorithm,
                                            Provider provider)
```

The `algorithm` parameter is the name of a certification path validation algorithm (for example, "PKIX"). Standard `CertPathValidator` algorithm names are listed in the Java Security Standard Algorithm Names.

Validating a Certification Path

Once a `CertPathValidator` object is created, paths can be validated by calling the `validate` method, passing it the certification path to be validated and a set of algorithm-specific parameters:

```
public final CertPathValidatorResult
        validate(CertPath certPath, CertPathParameters params)
        throws CertPathValidatorException,
               InvalidAlgorithmParameterException
```

If the validation algorithm is successful, the result is returned in an object implementing the CertPathValidatorResult interface. Otherwise, a `CertPathValidatorException` is thrown. The `CertPathValidatorException` contains methods that return the `CertPath`, and if relevant, the index of the certificate that caused the algorithm to fail and the root exception or cause of the failure.

Note that the `CertPath` and `CertPathParameters` passed to the `validate` method must be of a type that is supported by the validation algorithm. Otherwise, an `InvalidAlgorithmParameterException` is thrown. For example, a `CertPathValidator` instance that implements the PKIX algorithm validates `CertPath` objects of type X.509 and `CertPathParameters` that are an instance of `PKIXParameters`.

The CertPathValidatorResult Interface

The `CertPathValidatorResult` interface is a transparent representation of the successful result or output of a certification path validation algorithm.

The main purpose of this interface is to group and provide type safety for all validation results. Similar to the `CertPathParameters` interface, `CertPathValidatorResult` extends `Cloneable` and defines a `clone()` method that does not throw an exception. This allows applications to clone any `CertPathValidatorResult` object.

Objects implementing the `CertPathValidatorResult` interface are returned by the `validate` method of CertPathValidatorResult interface when successful. If not successful, a `CertPathValidatorException` is thrown with a description of the failure. Typically, a concrete implementation of the `CertPathValidatorResult` interface will hold a set of output parameters specific to a particular certification path validation algorithm. For example, the `PKIXCertPathValidatorResult` class is an implementation of the `CertPathValidatorResult` interface, which contains methods to get the output parameters of the PKIX certification path validation algorithm. One such parameter is the valid policy tree. This parameter among others is discussed in more detail in the section discussing the PKIXCertPathValidatorResult class.

The following code sample shows how to create a `CertPathValidator` and use it to validate a certification path. The sample assumes that the `CertPath` and `CertPathParameters` objects which are passed to the `validate` method have been previously created; a more complete example will be illustrated in the section describing the PKIX classes.

```
    // create CertPathValidator that implements the "PKIX" algorithm
    CertPathValidator cpv = null;
    try {
        cpv = CertPathValidator.getInstance("PKIX");
    } catch (NoSuchAlgorithmException nsae) {
        System.err.println(nsae);
        System.exit(1);
    }
    // validate certification path ("cp") with specified parameters
("params")
    try {
        CertPathValidatorResult cpvResult = cpv.validate(cp, params);
    } catch (InvalidAlgorithmParameterException iape) {
        System.err.println("validation failed: " + iape);
        System.exit(1);
    } catch (CertPathValidatorException cpve) {
        System.err.println("validation failed: " + cpve);
        System.err.println("index of certificate that caused exception: "
                + cpve.getIndex());
        System.exit(1);
    }
```

Certification Path Building Classes

The Java Certification Path API includes classes for building (or discovering) certification paths. An application uses an instance of the `CertPathBuilder` class to build a `CertPath` object. If successful, the result of the build is returned in an object implementing the `CertPathBuilderResult` interface.

Topics

The CertPathBuilder Class

The CertPathBuilderResult Interface

The CertPathBuilder Class

The `CertPathBuilder` class is an engine class used to build a certification path.

Creating a CertPathBuilder Object

As with all engine classes, the way to get a `CertPathBuilder` object for a particular build algorithm is to call one of the `getInstance` static factory method on the `CertPathBuilder` class:

```
public static CertPathBuilder getInstance(String algorithm)
public static CertPathBuilder getInstance(String algorithm,
                                          String provider)
public static CertPathBuilder getInstance(String algorithm,
                                          Provider provider)
```

The `algorithm` parameter is the name of a certification path builder algorithm (for example, "PKIX"). Standard `CertPathBuilder` algorithm names are listed in Java Security Standard Algorithm Names.

Building a Certification Path

Once a `CertPathBuilder` object is created, paths can be constructed by calling the `build` method, passing it an algorithm-specific parameter specification:

```
public final CertPathBuilderResult build(CertPathParameters params)
        throws CertPathBuilderException,
               InvalidAlgorithmParameterException
```

If the build algorithm is successful, the result is returned in an object implementing the CertPathBuilderResult interface. Otherwise, a `CertPathBuilderException` is thrown containing information about the failure; for example, the underlying exception (if any) and an error message.

Note that the `CertPathParameters` passed to the `build` method must be of a type that is supported by the build algorithm. Otherwise, an `InvalidAlgorithmParameterException` is thrown.

The CertPathBuilderResult Interface

The `CertPathBuilderResult` interface is a transparent representation of the result or output of a certification path builder algorithm.

This interface contains a method to return the certification path that has been successfully built:

```
public CertPath getCertPath()
```

The purpose of the `CertPathBuilderResult` interface is to group (and provide type safety for) all build results. Like the `CertPathValidatorResult` interface, `CertPathBuilderResult` extends `Cloneable` and defines a `clone()` method that does not throw an exception. This allows applications to clone any `CertPathBuilderResult` object.

Objects implementing the `CertPathBuilderResult` interface are returned by the `build` method of CertPathBuilder.

The following code sample shows how to create a `CertPathBuilder` and use it to build a certification path. The sample assumes that the `CertPathParameters` object which is passed to the `build` method has been previously created; a more complete example will be illustrated in the section describing the PKIX classes.

```
// create CertPathBuilder that implements the "PKIX" algorithm
CertPathBuilder cpb = null;
try {
    cpb = CertPathBuilder.getInstance("PKIX");
} catch (NoSuchAlgorithmException nsae) {
    System.err.println(nsae);
    System.exit(1);
}
// build certification path using specified parameters ("params")
try {
    CertPathBuilderResult cpbResult = cpb.build(params);
    CertPath cp = cpbResult.getCertPath();
    System.out.println("build passed, path contents: " + cp);
} catch (InvalidAlgorithmParameterException iape) {
    System.err.println("build failed: " + iape);
    System.exit(1);
} catch (CertPathBuilderException cpbe) {
    System.err.println("build failed: " + cpbe);
    System.exit(1);
}
```

Certificate/CRL Storage Classes

The Java Certification Path API includes the `CertStore` class for retrieving certificates and CRLs from a repository.

This class enables a caller to specify the repository a `CertPathValidator` or `CertPathBuilder` implementation should use to find certificates and CRLs. See the `addCertStores` method of the PKIXParameters class.

A `CertPathValidator` implementation may use the `CertStore` object that the caller specifies as a callback mechanism to fetch CRLs for performing revocation checks. Similarly, a `CertPathBuilder` may use the `CertStore` as a callback mechanism to fetch certificates and, if performing revocation checks, CRLs.

Topics

The CertStore Class

The CertStoreParameters Interface

The CertSelector and CRLSelector Interfaces

The CertStore Class

The `CertStore` class is an engine class used to provide the functionality of a certificate and certificate revocation list (CRL) repository.

This class can be used by CertPathBuilder and CertPathValidator implementations to find certificates and CRLs, or as a general purpose certificate and CRL retrieval mechanism.

Unlike the `java.security.KeyStore` class, which provides access to a cache of private keys and trusted certificates, a `CertStore` is designed to provide access to a potentially vast repository of untrusted certificates and CRLs. For example, an LDAP implementation of `CertStore` provides access to certificates and CRLs stored in one or more directories using the LDAP protocol.

All public methods of `CertStore` objects are thread-safe. That is, multiple threads may concurrently invoke these methods on a single `CertStore` object (or more than one) with no ill effects. This allows a `CertPathBuilder` to search for a CRL while simultaneously searching for further certificates, for instance.

Creating a CertStore Object

As with all engine classes, the way to get a `CertStore` object for a particular repository type is to call one of the `getInstance` static factory methods on the `CertStore` class:

```
public static CertStore getInstance(String type,
        CertStoreParameters params)
public static CertStore getInstance(String type,
        CertStoreParameters params, String provider)
public static CertStore getInstance(String type,
        CertStoreParameters params, Provider provider)
```

The `type` parameter is the name of a certificate repository type (for example, "LDAP"). Standard `CertStore` types are listed in Java Security Standard Algorithm Names.

The initialization parameters (`params`) are specific to the repository type. For example, the initialization parameters for a server-based repository may include the hostname and the port of the server. An `InvalidAlgorithmParameterException` is thrown if the

parameters are invalid for this `CertStore` type. The `getCertStoreParameters` method returns the `CertStoreParameters` that were used to initialize a `CertStore`:

```
public final CertStoreParameters getCertStoreParameters()
```

Retrieving Certificates

After you have created a `CertStore` object, you can retrieve certificates from the repository using the `getCertificates` method. This method takes a CertSelector (discussed in more detail later) object as an argument, which specifies a set of selection criteria for determining which certificates should be returned:

```
        public final Collection<? extends Certificate>
getCertificates(CertSelector selector)
                throws CertStoreException
```

This method returns a `Collection` of `java.security.cert.Certificate` objects that satisfy the selection criteria. An empty `Collection` is returned if there are no matches. A `CertStoreException` is usually thrown if an unexpected error condition is encountered, such as a communications failure with a remote repository.

For some `CertStore` implementations, it may not be feasible to search the entire repository for certificates or CRLs that match the specified selection criteria. In these instances, the `CertStore` implementation may use information that is specified in the selectors to locate certificates and CRLs. For instance, an LDAP `CertStore` may not search all entries in the directory. Instead, it may just search entries that are likely to contain the certificates it is looking for. If the `CertSelector` provided does not provide enough information for the LDAP `CertStore` to determine which entries it should look in, the LDAP `CertStore` may throw a `CertStoreException`.

Retrieving CRLs

You can also retrieve CRLs from the repository using the `getCRLs` method. This method takes a CRLSelector (discussed in more detail later) object as an argument, which specifies a set of selection criteria for determining which CRLs should be returned:

```
        public final Collection<? extends CRL> getCRLs(CRLSelector
selector)
                throws CertStoreException
```

This method returns a `Collection` of `java.security.cert.CRL` objects that satisfy the selection criteria. An empty `Collection` is returned if there are no matches.

The CertStoreParameters Interface

The `CertStoreParameters` interface is a transparent representation of the set of parameters used with a particular `CertStore`.

The main purpose of this interface is to group and provide type safety for all certificate storage parameter specifications. The `CertStoreParameters` interface extends the `Cloneable` interface and defines a `clone` method that does not throw an exception. Implementations of this interface should implement and override the `Object.clone()` method, if necessary. This allows applications to clone any `CertStoreParameters` object.

Objects implementing the `CertStoreParameters` interface are passed as arguments to the `getInstance` method of the `CertStore` class. Two classes implementing the `CertStoreParameters` interface are defined in this API: the LDAPCertStoreParameters class and the CollectionCertStoreParameters class.

The LDAPCertStoreParameters Class

The `LDAPCertStoreParameters` class is an implementation of the `CertStoreParameters` interface and holds a set of minimal initialization parameters (host and port number of the directory server) for retrieving certificates and CRLs from a `CertStore` of type **LDAP**.

See `LDAPCertStoreParameters`.

The CollectionCertStoreParameters Class

The `CollectionCertStoreParameters` class is an implementation of the `CertStoreParameters` interface and holds a set of initialization parameters for retrieving certificates and CRLs from a `CertStore` of type **Collection**.

See `CollectionCertStoreParameters`.

The CertSelector and CRLSelector Interfaces

The `CertSelector` and `CRLSelector` interfaces are a specification of the set of criteria for selecting certificates and CRLs from a collection or large group of certificates and CRLs.

The interfaces group and provide type safety for all selector specifications. Each selector interface extends `Cloneable` and defines a `clone()` method that does not throw an exception. This allows applications to clone any `CertSelector` or `CRLSelector` object.

The `CertSelector` and `CRLSelector` interfaces each define a method named `match`. The `match` method takes a `Certificate` or `CRL` object as an argument and returns `true` if the object satisfies the selection criteria. Otherwise, it returns `false`. The `match` method for the `CertSelector` interface is defined as follows:

```
public boolean match(Certificate cert)
```

and for the `CRLSelector` interface:

```
public boolean match(CRL crl)
```

Typically, objects implementing these interfaces are passed as parameters to the `getCertificates` and `getCRLs` methods of the `CertStore` class. These methods return

a `Collection` of `Certificates` or `CRLs` from the `CertStore` repository that match the specified selection criteria. `CertSelectors` may also be used to specify the validation constraints on a target or end-entity certificate in a certification path (see for example, the `PKIXParameters.setTargetCertConstraints` method.)

The X509CertSelector Class

The `X509CertSelector` class is an implementation of the `CertSelector` interface that defines a set of criteria for selecting X.509 certificates.

An `X509Certificate` object must match *all* of the specified criteria to be selected by the `match` method. The selection criteria are designed to be used by a `CertPathBuilder` implementation to discover potential certificates as it builds an X.509 certification path.

For example, the `setSubject` method of `X509CertSelector` allows a PKIX `CertPathBuilder` to filter out `X509Certificates` that do not match the issuer name of the preceding `X509Certificate` in a partially completed chain. By setting this and other criteria in an `X509CertSelector` object, a `CertPathBuilder` is able to discard irrelevant certificates and more easily find an X.509 certification path that meets the requirements specified in the `CertPathParameters` object.

See RFC 5280 for definitions of the X.509 certificate extensions mentioned in this section.

Creating an X509CertSelector Object

An `X509CertSelector` object is created by calling the default constructor:

```
public X509CertSelector()
```

No criteria are initially set (any `X509Certificate` will match).

Setting Selection Criteria

The selection criteria allow a caller to match on different components of an X.509 certificate. A few of the methods for setting selection criteria are described here. See `X509CertSelector`.

The `setIssuer` methods set the issuer criterion:

```
public void setIssuer(X500Principal issuer)
public void setIssuer(String issuerDN)
public void setIssuer(byte[] issuerDN)
```

The specified distinguished name (in `X500Principal`, RFC 2253 String or ASN.1 DER encoded form) must match the issuer distinguished name in the certificate. If null, any issuer distinguished name will do. Note that use of an `X500Principal` to represent a distinguished name is preferred because it is more efficient and suitably typed.

Similarly, the `setSubject` methods set the subject criterion:

```
public void setSubject(X500Principal subject)
public void setSubject(String subjectDN)
public void setSubject(byte[] subjectDN)
```

The specified distinguished name (in `X500Principal`, RFC 2253 String or ASN.1 DER encoded form) must match the subject distinguished name in the certificate. If null, any subject distinguished name will do.

The `setSerialNumber` method sets the serialNumber criterion:

```
public void setSerialNumber(BigInteger serial)
```

The specified serial number must match the certificate serial number in the certificate. If null, any certificate serial number will do.

The `setAuthorityKeyIdentifier` method sets the authorityKeyIdentifier criterion:

```
public void setAuthorityKeyIdentifier(byte[] authorityKeyID)
```

The certificate must contain an Authority Key Identifier extension matching the specified value. If null, no check will be done on the authorityKeyIdentifier criterion.

The `setCertificateValid` method sets the certificateValid criterion:

```
public void setCertificateValid(Date certValid)
```

The specified date must fall within the certificate validity period for the certificate. If null, any date is valid.

The `setKeyUsage` method sets the keyUsage criterion:

```
public void setKeyUsage(boolean[] keyUsage)
```

The certificate's Key Usage Extension must allow the specified key usage values (those which are set to true). If null, no keyUsage check will be done.

Getting Selection Criteria

The current values for each of the selection criteria can be retrieved using an appropriate `get` method. See `X509CertSelector` .

Here is an example of retrieving X.509 certificates from an LDAP `CertStore` with the `X509CertSelector` class.

First, we create the `LDAPCertStoreParameters` object that we will use to initialize the `CertStore` object with the hostname and port of the LDAP server:

```
LDAPCertStoreParameters lcsp = new
        LDAPCertStoreParameters("ldap.sun.com", 389);
```

Next, create the CertStore object, and pass it the `LDAPCertStoreParameters` object, as in the following statement:

```
CertStore cs = CertStore.getInstance("LDAP", lcsp);
```

This call creates a `CertStore` object that retrieves certificates and CRLs from an LDAP repository using the schema defined in RFC 2587.

The following block of code establishes an `X509CertSelector` to retrieve all unexpired (as of the current date and time) end-entity certificates issued to a particular subject with 1) a key usage that allows digital signatures, and 2) a subject alternative name with a specific email address:

```
X509CertSelector xcs = new X509CertSelector();

// select only unexpired certificates
xcs.setCertificateValid(new Date());

// select only certificates issued to
// 'CN=alice, O=xyz, C=us'
xcs.setSubject(new X500Principal("CN=alice, O=xyz, C=us"));

// select only end-entity certificates
xcs.setBasicConstraints(-2);

// select only certificates with a digitalSignature
// keyUsage bit set (set the first entry in the
// boolean array to true)
boolean[] keyUsage = {true};
xcs.setKeyUsage(keyUsage);

// select only certificates with a subjectAltName of
// 'alice@xyz.example.com' (1 is the integer value of
// an RFC822Name)
xcs.addSubjectAlternativeName(1, "alice@xyz.example.com");
```

Then we pass the selector to the `getCertificates` method of our `CertStore` object that we previously created:

```
Collection<Certificate> certs = cs.getCertificates(xcs);
```

A PKIX `CertPathBuilder` may use similar code to help discover and sort through potential certificates by discarding those that do not meet validation constraints or other criteria.

The X509CRLSelector Class

The `X509CRLSelector` class is an implementation of the `CRLSelector` interface that defines a set of criteria for selecting X.509 CRLs.

An `X509CRL` object must match *all* of the specified criteria to be selected by the `match` method. The selection criteria are designed to be useful to a `CertPathValidator` or `CertPathBuilder` implementation that must retrieve CRLs from a repository to check the revocation status of certificates in an X.509 certification path.

For example, the `setDateAndTime` method of `X509CRLSelector` allows a PKIX `CertPathValidator` to filter out `X509CRL`s that have been issued after or expire before the time indicated. By setting this and other criteria in an `X509CRLSelector` object, it allows the `CertPathValidator` to discard irrelevant CRLs and more easily check if a certificate has been revoked.

Please refer to RFC 5280 for definitions of the X.509 CRL fields and extensions mentioned in this section.

Creating an X509CRLSelector Object

An `X509CRLSelector` object is created by calling the default constructor:

```
public X509CRLSelector()
```

No criteria are initially set (any `X509CRL` will match).

Setting Selection Criteria

The selection criteria allow a caller to match on different components of an X.509 CRL. Most of the methods for setting selection criteria are described here. Please refer to the `X509CRLSelector Class` API documentation for details on the remaining methods.

The `setIssuers` and `setIssuerNames` methods set the issuerNames criterion:

```
public void setIssuers(Collection<X500Principal> issuers)
public void setIssuerNames(Collection<?> names)
```

The issuer distinguished name in the CRL must match at least one of the specified distinguished names. The `setIssuers` method is preferred as the use of `X500Principal`s to represent distinguished names is more efficient and suitably typed. For the `setIssuerNames` method, each entry of the `names` argument is either a `String` or a byte array (representing the name, in RFC 2253 or ASN.1 DER encoded form, respectively). If null, any issuer distinguished name will do.

The `setMinCRLNumber` and `setMaxCRLNumber` methods set the minCRLNumber and maxCRLNumber criterion:

```
public void setMinCRLNumber(BigInteger minCRL)
public void setMaxCRLNumber(BigInteger maxCRL)
```

The CRL must have a CRL Number extension whose value is greater than or equal to the specified value if the `setMinCRLNumber` method is called, and less than or equal to the specified value if the `setMaxCRLNumber` method is called. If the value passed to one of these methods is null, the corresponding check is not done.

The `setDateAndTime` method sets the dateAndTime criterion:

```
public void setDateAndTime(Date dateAndTime)
```

The specified date must be equal to or later than the value of the thisUpdate component of the CRL and earlier than the value of the nextUpdate component. If null, no dateAndTime check will be done.

The `setCertificateChecking` method sets the certificate whose revocation status is being checked:

```
public void setCertificateChecking(X509Certificate cert)
```

This is not a criterion. Rather, it is optional information that may help a `CertStore` find CRLs that would be relevant when checking revocation for the specified certificate. If null is specified, then no such optional information is provided. An application should always call this method when checking revocation for a particular certificate, as it may provide the `CertStore` with more information for finding the correct CRLs and filtering out irrelevant ones.

Getting Selection Criteria

The current values for each of the selection criteria can be retrieved using an appropriate `get` method. Please refer to the `X509CRLSelector Class` API documentation for further details on these methods.

Creating an `X509CRLSelector` to retrieve CRLs from an LDAP repository is similar to the `X509CertSelector` example. Suppose we want to retrieve all current (as of the current date and time) CRLs issued by a specific CA and with a minimum CRL number. First, we create an `X509CRLSelector` object and call the appropriate methods to set the selection criteria:

```
X509CRLSelector xcrls = new X509CRLSelector();
// select CRLs satisfying current date and time
xcrls.setDateAndTime(new Date());
// select CRLs issued by 'O=xyz, C=us'
xcrls.addIssuerName("O=xyz, C=us");
```

```
// select only CRLs with a CRL number at least '2'
xcrls.setMinCRLNumber(new BigInteger("2"));
```

Then we pass the selector to the `getCRLs` method of our `CertStore` object (created in the `X509CertSelector` example):

```
Collection<CRL> crls = cs.getCRLs(xcrls);
```

PKIX Classes

The Java Certification Path API includes a set of algorithm-specific classes modeled for use with the PKIX certification path validation algorithm.

The PKIX certification path validation algorithm is defined in RFC 5280: *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*.

Topics

The TrustAnchor Class

The PKIXParameters Class

The CertPathValidatorResult Interface

The PolicyNode Interface and PolicyQualifierInfo Class

The PKIXBuilderParameters Class

The PKIXCertPathBuilderResult Class

The PKIXCertPathChecker Class

Using PKIXCertPathChecker in Certificate Path Validation

The TrustAnchor Class

The `TrustAnchor` class represents a "most-trusted CA", which is used as a trust anchor for validating X.509 certification paths.

A `TrustAnchor` includes the public key of the CA, the CA's name, and any constraints on the set of paths that can be validated using this key. These parameters can be specified in the form of a trusted `X509Certificate` or as individual parameters.

All `TrustAnchor` objects are immutable and thread-safe. That is, multiple threads may concurrently invoke the methods defined in this class on a single `TrustAnchor` object (or more than one) with no ill effects. Requiring `TrustAnchor` objects to be immutable and thread-safe allows them to be passed around to various pieces of code without worrying about coordinating access.

> **Note:**
>
> Although this class is described as a PKIX class it may be used with other X.509 certification path validation algorithms.

Creating a TrustAnchor Object

To instantiate a `TrustAnchor` object, a caller must specify "the most-trusted CA" as a trusted `X509Certificate` or public key and distinguished name pair. The caller may also optionally specify name constraints that are applied to the trust anchor by the validation algorithm during initialization. Note that support for name constraints on trust anchors is not required by the PKIX algorithm, therefore a PKIX `CertPathValidator` or `CertPathBuilder` may choose not to support this parameter and instead throw an exception. Use one of the following constructors to create a `TrustAnchor` object:

```
public TrustAnchor(X509Certificate trustedCert,
        byte[] nameConstraints)
public TrustAnchor(X500Principal caPrincipal, PublicKey pubKey,
        byte[] nameConstraints)
public TrustAnchor(String caName, PublicKey pubKey,
        byte[] nameConstraints)
```

The `nameConstraints` parameter is specified as a byte array containing the ASN.1 DER encoding of a NameConstraints extension. An `IllegalArgumentException` is thrown if the name constraints cannot be decoded (are not formatted correctly).

Getting Parameter Values

Each of the parameters can be retrieved using a corresponding get method:

```
public final X509Certificate getTrustedCert()
public final X500Principal getCA()
public final String getCAName()
public final PublicKey getCAPublicKey()
public final byte[] getNameConstraints()
```

> **Note:**
>
> The `getTrustedCert` method returns `null` if the trust anchor was specified as a public key and name pair. Likewise, the `getCA`, `getCAName` and `getCAPublicKey` methods return `null` if the trust anchor was specified as an `X509Certificate`.

The PKIXParameters Class

The `PKIXParametersClass` class specifies the set of input parameters defined by the PKIX certification path validation algorithm. It also includes a few additional useful parameters.

This class implements the CertPathParameters interface.

An X.509 `CertPath` object and a `PKIXParameters` object are passed as arguments to the `validate` method of a CertPathValidator instance implementing the PKIX algorithm.

The `CertPathValidator` uses the parameters to initialize the PKIX certification path validation algorithm.

Creating a PKIXParameters Object

To instantiate a `PKIXParameters` object, a caller must specify "the most-trusted CA(s)" as defined by the PKIX validation algorithm. The most-trusted CAs can be specified using one of two constructors:

```
public PKIXParameters(Set<TrustAnchor> trustAnchors)
    throws InvalidAlgorithmParameterException
public PKIXParameters(KeyStore keystore)
    throws KeyStoreException, InvalidAlgorithmParameterException
```

The first constructor allows the caller to specify the most-trusted CAs as a `Set` of `TrustAnchor` objects. Alternatively, a caller can use the second constructor and specify a `KeyStore` instance containing trusted certificate entries, each of which will be considered as a most-trusted CA.

Setting Parameter Values

After a `PKIXParameters` object has been created, a caller can set (or replace the current value of) various parameters. A few of the methods for setting parameters are described here. Please refer to the `PKIXParameters` API documentation for details on the other methods.

The `setInitialPolicies` method sets the initial policy identifiers, as specified by the PKIX validation algorithm. The elements of the `Set` are object identifiers (OIDs) represented as a `String`. If the `initialPolicies` parameter is null or not set, any policy is acceptable:

```
public void setInitialPolicies(Set<String> initialPolicies)
```

The `setDate` method sets the time for which the validity of the path should be determined. If the `date` parameter is not set or is null, the current date is used:

```
public void setDate(Date date)
```

The `setPolicyMappingInhibited` method sets the value of the policy mapping inhibited flag. The default value for the flag, if not specified, is false:

```
public void setPolicyMappingInhibited(boolean val)
```

The `setExplicitPolicyRequired` method sets the value of the explicit policy required flag. The default value for the flag, if not specified, is false:

```
public void setExplicitPolicyRequired(boolean val)
```

The `setAnyPolicyInhibited` method sets the value of the any policy inhibited flag. The default value for the flag, if not specified, is false:

```
public void setAnyPolicyInhibited(boolean val)
```

The `setTargetCertConstraints` method allows the caller to set constraints on the target or end-entity certificate. For example, the caller can specify that the target certificate must contain a specific subject name. The constraints are specified as a `CertSelector` object. If the `selector` parameter is null or not set, no constraints are defined on the target certificate:

```
public void setTargetCertConstraints(CertSelector selector)
```

The `setCertStores` method allows a caller to specify a `List` of `CertStore` objects that will be used by a PKIX implementation of `CertPathValidator` to find CRLs for path validation. This provides an extensible mechanism for specifying where to locate CRLs. The `setCertStores` method takes a `List` of `CertStore` objects as a parameter. The first `CertStores` in the list may be preferred to those that appear later.

```
public void setCertStores(List<CertStore> stores)
```

The `setCertPathCheckers` method allows a caller to extend the PKIX validation algorithm by creating implementation-specific certification path checkers. For example, this mechanism can be used to process private certificate extensions. The `setCertPathCheckers` method takes a list of PKIXCertPathChecker (discussed later) objects as a parameter:

```
public void setCertPathCheckers(List<PKIXCertPathChecker> checkers)
```

The `setRevocationEnabled` method allows a caller to disable revocation checking. Revocation checking is enabled by default, since it is a required check of the PKIX validation algorithm. However, PKIX does not define how revocation should be checked. An implementation may use CRLs or OCSP, for example. This method allows the caller to disable the implementation's default revocation checking mechanism if it is not appropriate. A different revocation checking mechanism can then be specified by calling the `setCertPathCheckers` method, and passing it a `PKIXCertPathChecker` that implements the alternate mechanism.

```
public void setRevocationEnabled(boolean val)
```

The `setPolicyQualifiersRejected` method allows a caller to enable or disable policy qualifier processing. When a `PKIXParameters` object is created, this flag is set to `true`. This setting reflects the most common (and simplest) strategy for processing policy

qualifiers. Applications that want to use a more sophisticated policy must set this flag to `false`.

```
public void setPolicyQualifiersRejected(boolean qualifiersRejected)
```

Getting Parameter Values

The current values for each of the parameters can be retrieved using an appropriate `get` method. Please refer to the `Class PKIXParameters` API documentation for further details on these methods.

The PKIXCertPathValidatorResult Class

The `PKIXCertPathValidatorResult` class represents the result of the PKIX certification path validation algorithm.

This class implements the CertPathValidatorResult interface. It holds the valid policy tree and subject public key resulting from the validation algorithm, and includes methods (`getPolicyTree()` and `getPublicKey()`) for returning them. Instances of `PKIXCertPathValidatorResult` are returned by the `validate` method of CertPathValidator objects implementing the PKIX algorithm.

Please refer to the `PKIXCertPathValidatorResult` API documentation for more detailed information on this class.

The PolicyNode Interface and PolicyQualifierInfo Class

The PKIX validation algorithm defines several outputs related to certificate policy processing. Most applications will not need to use these outputs, but all providers that implement the PKIX validation or building algorithm must support them.

The `PolicyNode` interface represents a node of a valid policy tree resulting from a successful execution of the PKIX certification path validation. An application can obtain the root of a valid policy tree using the `getPolicyTree` method of `PKIXCertPathValidatorResult`. Policy Trees are discussed in more detail in the RFC 5280.

The `getPolicyQualifiers` method of `PolicyNode` returns a `Set` of `PolicyQualifierInfo` objects, each of which represents a policy qualifier contained in the Certificate Policies extension of the relevant certificate that this policy applies to.

Most applications will not need to examine the valid policy tree and policy qualifiers. They can achieve their policy processing goals by setting the policy-related parameters in `PKIXParameters`. However, the valid policy tree is available for more sophisticated applications, especially those that process policy qualifiers.

Please refer to the `Interface PolicyNode` and `PolicyQualifierInfo` API documentation for more detailed information on these classes.

Example 10-1 Example of Validating a Certification Path using the PKIX algorithm

This is an example of validating a certification path with the PKIX validation algorithm. The example ignores most of the exception handling and assumes that the certification path and public key of the trust anchor have already been created.

First, create the CertPathValidator, as in the following line:

```
CertPathValidator cpv = CertPathValidator.getInstance("PKIX");
```

The next step is to create a TrustAnchor object. This will be used as an anchor for validating the certification path. In this example, the most-trusted CA is specified as a public key and name (name constraints are not applied and are specified as `null`):

```
TrustAnchor anchor = new TrustAnchor("O=xyz,C=us", pubkey, null);
```

The next step is to create a PKIXParameters object. This will be used to populate the parameters used by the PKIX algorithm. In this example, we pass to the constructor a `Set` containing a single element - the `TrustAnchor` we created in the previous step:

```
    PKIXParameters params = new
PKIXParameters(Collections.singleton(anchor));
```

Next, we populate the parameters object with constraints or other parameters used by the validation algorithm. In this example, we enable the explicitPolicyRequired flag and specify a set of initial policy OIDs (the contents of the set are not shown):

```
// set other PKIX parameters here
params.setExplicitPolicyRequired(true);
params.setInitialPolicies(policyIds);
```

The final step is to validate the certification path using the input parameter set we have created:

```
try {
    PKIXCertPathValidatorResult result =
        (PKIXCertPathValidatorResult) cpv.validate(certPath, params);
    PolicyNode policyTree = result.getPolicyTree();
    PublicKey subjectPublicKey = result.getPublicKey();
} catch (CertPathValidatorException cpve) {
    System.out.println("Validation failure, cert["
        + cpve.getIndex() + "] :" + cpve.getMessage());
}
```

If the validation algorithm is successful, the policy tree and subject public key resulting from the validation algorithm are obtained using the `getPolicyTree` and `getPublicKey` methods of `PKIXCertPathValidatorResult`.

Otherwise, a `CertPathValidatorException` is thrown and the caller can catch the exception and print some details about the failure, such as the error message and the index of the certificate that caused the failure.

The PKIXBuilderParameters Class

The `PKIXBuilderParameters` class specifies the set of parameters to be used with `CertPathBuilder` class.

This class (which extends the PKIXParameters class) specifies the set of parameters to be used with CertPathBuilder class that build certification paths validated against the PKIX certification path validation algorithm.

A `PKIXBuilderParameters` object is passed as an argument to the `build` method of a CertPathBuilder instance implementing the PKIX algorithm. All PKIX `CertPathBuilders` *must* return certification paths which have been validated according to the PKIX certification path validation algorithm.

Please note that the mechanism that a PKIX `CertPathBuilder` uses to validate a constructed path is an implementation detail. For example, an implementation might attempt to first build a path with minimal validation and then fully validate it using an instance of a PKIX `CertPathValidator`, whereas a more efficient implementation may validate more of the path as it is building it, and backtrack to previous stages if it encounters validation failures or dead-ends.

Creating a PKIXBuilderParameters Object

Creating a `PKIXBuilderParameters` object is similar to creating a `PKIXParameters` object. However, a caller *must* specify constraints on the target or end-entity certificate when creating a `PKIXBuilderParameters` object. These constraints should provide the `CertPathBuilder` with enough information to find the target certificate. The constraints are specified as a `CertSelector` object. Use one of the following constructors to create a `PKIXBuilderParameters` object:

```
        public PKIXBuilderParameters(Set<TrustAnchor> trustAnchors,
                CertSelector targetConstraints)
                throws InvalidAlgorithmParameterException
        public PKIXBuilderParameters(KeyStore keystore,
                CertSelector targetConstraints)
                throws KeyStoreException,
InvalidAlgorithmParameterException
```

Getting/Setting Parameter Values

The `PKIXBuilderParameters` class inherits all of the parameters that can be set in the `PKIXParameters` class. In addition, the `setMaxPathLength` method can be called to place a limit on the maximum number of certificates in a certification path:

```
        public void setMaxPathLength(int maxPathLength)
```

The `maxPathLength` parameter specifies the maximum number of non-self-issued intermediate certificates that may exist in a certification path. A `CertPathBuilder` instance implementing the PKIX algorithm must not build paths longer than the length specified. If the value is 0, the path can only contain a single certificate. If the value is -1, the path length is unconstrained (i.e., there is no maximum). The default maximum

path length, if not specified, is 5. This method is useful to prevent the `CertPathBuilder` from spending resources and time constructing long paths that may or may not meet the caller's requirements.

If any of the CA certificates in the path contain a Basic Constraints extension, the value of the pathLenConstraint component of the extension overrides the value of the `maxPathLength` parameter whenever the result is a certification path of smaller length. There is also a corresponding `getMaxPathLength` method for retrieving this parameter:

```
public int getMaxPathLength()
```

Also, the `setCertStores` method (inherited from the `PKIXParameters` class) is typically used by a PKIX implementation of `CertPathBuilder` to find Certificates for path construction as well as finding CRLs for path validation. This provides an extensible mechanism for specifying where to locate Certificates and CRLs.

The PKIXCertPathBuilderResult Class

The `PKIXCertPathBuilderResult` class represents the successful result of the PKIX certification path construction algorithm.

This class extends the `PKIXCertPathValidatorResult` class and implements the CertPathBuilder interface. Instances of `PKIXCertPathBuilderResult` are returned by the `build` method of CertPathBuilder objects implementing the PKIX algorithm.

The `getCertPath` method of a `PKIXCertPathBuilderResult` instance always returns a `CertPath` object validated using the PKIX certification path validation algorithm. The returned `CertPath` object does not include the most-trusted CA certificate that may have been used to anchor the path. Instead, use the `getTrustAnchor` method to get the `Certificate` of the most-trusted CA.

See the `PKIXCertPathBuilderResult` API documentation for more detailed information on this class.

Example 10-2 Example of Building a Certification Path using the PKIX algorithm

This is an example of building a certification path validated against the PKIX algorithm. Some details have been left out, such as exception handling, and the creation of the trust anchors and certificates for populating the `CertStore`.

First, create the CertPathBuilder, as in the following example:

```
CertPathBuilder cpb = CertPathBuilder.getInstance("PKIX");
```

This call creates a `CertPathBuilder` object that returns paths validated against the PKIX algorithm.

The next step is to create a PKIXBuilderParameters object. This will be used to populate the PKIX parameters used by the `CertPathBuilder`:

```
// Create parameters object, passing it a Set of
```

```
// trust anchors for anchoring the path
// and a target subject DN.
X509CertSelector targetConstraints = new X509CertSelector();
targetConstraints.setSubject("CN=alice,O=xyz,C=us");
PKIXBuilderParameters params =
    new PKIXBuilderParameters(trustAnchors, targetConstraints);
```

The next step is to specify the `CertStore` that the `CertPathBuilder` will use to look for certificates and CRLs. For this example, we will populate a Collection `CertStore` with the certificates and CRLs:

```
CollectionCertStoreParameters ccsp =
    new CollectionCertStoreParameters(certsAndCrls);
CertStore store = CertStore.getInstance("Collection", ccsp);
params.addCertStore(store);
```

The next step is to build the certification path using the input parameter set we have created:

```
try {
    PKIXCertPathBuilderResult result =
        (PKIXCertPathBuilderResult) cpb.build(params);
    CertPath cp = result.getCertPath();
} catch (CertPathBuilderException cpbe) {
    System.out.println("build failed: " + cpbe.getMessage());
}
```

If the `CertPathBuilder` cannot build a path that meets the supplied parameters it will throw a `CertPathBuilderException`. Otherwise, the validated certification path can be obtained from the `PKIXCertPathBuilderResult` using the `getCertPath` method.

The PKIXCertPathChecker Class

The PKIXCertPathChecker class allows a user to extend a PKIX `CertPathValidator` or `CertPathBuilder` implementation. This is an advanced feature that most users will not need to understand. However, anyone implementing a PKIX service provider should read this section

The `PKIXCertPathChecker` class is an abstract class that executes one or more checks on an X.509 certificate. Developers should create concrete implementations of the `PKIXCertPathChecker` class when it is necessary to dynamically extend a PKIX `CertPathValidator` or `CertPathBuilder` implementation at runtime. The following are a few examples of when a `PKIXCertPathChecker` implementation is useful:

- If the revocation mechanism supplied by a PKIX `CertPathValidator` or `CertPathBuilder` implementation is not adequate: For example, you can use the `PKIXRevocationChecker` (introduced in JDK 8; see Check Revocation Status of Certificates with PKIXRevocationChecker Class) to have more control over the revocation mechanism, or you can implement your own `PKIXCertPathChecker` to check that certificates have not been revoked.

- If the user wants to recognize certificates containing a critical private extension. Since the extension is private, it will not be recognized by the PKIX `CertPathValidator` or `CertPathBuilder` implementation and a `CertPathValidatorException` will be thrown. In this case, a developer can implement a `PKIXCertPathChecker` that recognizes and processes the critical private extension.
- If the developer wants to record information about each certificate processed for debugging or display purposes.
- If the user wants to reject certificates with certain policy qualifiers.

The `setCertPathCheckers` method of the `PKIXParameters` class allows a user to pass a `List` of `PKIXCertPathChecker` objects to a PKIX `CertPathValidator` or `CertPathBuilder` implementation. Each of the `PKIXCertPathChecker` objects will be called in turn, for each certificate processed by the PKIX `CertPathValidator` or `CertPathBuilder` implementation.

Creating and using a PKIXCertPathChecker Object

The `PKIXCertPathChecker` class does not have a public constructor. This is intentional, since creating an instance of `PKIXCertPathChecker` is an implementation-specific issue. For example, the constructor for a `PKIXCertPathChecker` implementation that uses OCSP to check a certificate's revocation status may require the hostname and port of the OCSP server:

```
        PKIXCertPathChecker checker = new OCSPChecker("ocsp.sun.com",
1321);
```

Once the checker has been instantiated, it can be added as a parameter using the `addCertPathChecker` method of the `PKIXParameters` class:

```
        params.addCertPathChecker(checker);
```

Alternatively, a `List` of checkers can be added using the `setCertPathCheckers` method of the `PKIXParameters` class.

Implementing a PKIXCertPathChecker Object

The `PKIXCertPathChecker` class is abstract. It has four methods (`check`, `getSupportedExtensions`, `init`, and `isForwardCheckingSupported`) that all concrete subclasses must implement.

Implementing a `PKIXCertPathChecker` may be trivial or complex. A `PKIXCertPathChecker` implementation can be stateless or stateful. A stateless implementation does not maintain state between successive calls of the `check` method. For example, a `PKIXCertPathChecker` that checks that each certificate contains a particular policy qualifier is stateless. In contrast, a stateful implementation does maintain state between successive calls of the `check` method. The `check` method of a stateful implementation usually depends on the contents of prior certificates in the certification path. For example, a `PKIXCertPathChecker` that processes the NameConstraints extension is stateful.

Also, the order in which the certificates processed by a service provider implementation are presented (passed) to a `PKIXCertPathChecker` is very important, especially if the implementation is stateful. Depending on the algorithm used by the service provider, the certificates may be presented in *reverse* or *forward* order. A reverse ordering means that the certificates are ordered from the most trusted CA (if present) to the target subject, whereas a forward ordering means that the certificates are ordered from the target subject to the most trusted CA. The order must be made known to the `PKIXCertPathChecker` implementation, so that it knows how to process consecutive certificates.

Initializing a PKIXCertPathChecker Object

The `init` method initializes the internal state of the checker:

```
public abstract void init(boolean forward)
```

All stateful implementations should clear or initialize any internal state in the checker. This prevents a service provider implementation from calling a checker that is in an uninitialized state. It also allows stateful checkers to be reused in subsequent operations without reinstantiating them. The `forward` parameter indicates the order of the certificates presented to the `PKIXCertPathChecker`. If `forward` is `true`, the certificates are presented from target to trust anchor; if `false`, from trust anchor to target.

Forward Checking

The `isForwardCheckingSupported` method returns a `boolean` that indicates if the `PKIXCertPathChecker` supports forward checking:

```
public abstract boolean isForwardCheckingSupported()
```

All `PKIXCertPathChecker` implementations *must*support reverse checking. A `PKIXCertPathChecker` implementation *may*support forward checking.

Supporting forward checking improves the efficiency of `CertPathBuilder`s that build forward, since it allows paths to be checked as they are built. However, some stateful `PKIXCertPathChecker`s may find it difficult or impossible to support forward checking.

Supported Extensions

The `getSupportedExtensions` method returns an immutable `Set` of OID `String`s for the X.509 extensions that the `PKIXCertPathChecker` implementation supports (i.e., recognizes, is able to process):

```
public abstract Set<String> getSupportedExtensions()
```

The method should return `null` if no extensions are processed. All implementations should return the `Set` of OID `String`s that the `check` method may process.

A `CertPathBuilder` can use this information to identify certificates with unrecognized critical extensions, even when performing a forward build with a `PKIXCertPathChecker` that does not support forward checking.

Executing the Check

The following method executes a check on the certificate:

```
        public abstract void
                check(Certificate cert, Collection<String>
unresolvedCritExts)
                throws CertPathValidatorException
```

The `unresolvedCritExts` parameter contains a collection of OIDs as `String`s. These OIDs represent the set of critical extensions in the certificate that have not yet been resolved by the certification path validation algorithm. Concrete implementations of the `check` method should remove any critical extensions that it processes from the `unresolvedCritExts` parameter.

If the certificate does not pass the check(s), a `CertPathValidatorException` should be thrown.

Cloning a PKIXCertPathChecker

The `PKIXCertPathChecker` class implements the `Cloneable` interface. All stateful `PKIXCertPathChecker` implementations must override the `clone` method if necessary. The default implementation of the `clone` method calls the `Object.clone` method, which performs a simple clone by copying all fields of the original object to the new object. A stateless implementation should not override the `clone` method. However, all stateful implementations must ensure that the default `clone` method is correct, and override it if necessary. For example, a `PKIXCertPathChecker` that stores state in an array must override the `clone` method to make a copy of the array, rather than just a reference to the array.

The reason that `PKIXCertPathChecker` objects are `Cloneable` is to allow a PKIX `CertPathBuilder` implementation to efficiently backtrack and try another path when a potential certification path reaches a dead end or point of failure. In this case, the implementation is able to restore prior path validation states by restoring the cloned objects.

Example 10-3 Sample Code to Check for a Private Extension

This is an example of a stateless `PKIXCertPathChecker` implementation. It checks if a private extension exists in a certificate and processes it according to some rules.

```
        import java.security.cert.Certificate;
        import java.security.cert.X509Certificate;
        import java.util.Collection;
        import java.util.Collections;
        import java.util.Set;
        import java.security.cert.PKIXCertPathChecker;
        import java.security.cert.CertPathValidatorException;

        public class MyChecker extends PKIXCertPathChecker {
```

```
    private static Set supportedExtensions =
        Collections.singleton("2.16.840.1.113730.1.1");

    /*
     * Initialize checker
     */
    public void init(boolean forward)
        throws CertPathValidatorException {
        // nothing to initialize
    }

    public Set getSupportedExtensions() {
        return supportedExtensions;
    }

    public boolean isForwardCheckingSupported() {
        return true;
    }

    /*
     * Check certificate for presence of Netscape's
     * private extension
     * with OID "2.16.840.1.113730.1.1"
     */
    public void check(Certificate cert,
                      Collection unresolvedCritExts)
        throws CertPathValidatorException
    {
        X509Certificate xcert = (X509Certificate) cert;
        byte[] ext =
            xcert.getExtensionValue("2.16.840.1.113730.1.1");
        if (ext == null)
            return;

        //
        // process private extension according to some
        // rules - if check fails, throw a
        // CertPathValidatorException ...
        // {insert code here}

        // remove extension from collection of unresolved
        // extensions (if it exists)
        if (unresolvedCritExts != null)
            unresolvedCritExts.remove("2.16.840.1.113730.1.1");
    }
}
```

How a PKIX Service Provider implementation should use a PKIXCertPathChecker

Each `PKIXCertPathChecker` object must be initialized by a service provider implementation before commencing the build or validation algorithm, for example:

```
List<PKIXCertPathChecker> checkers = params.getCertPathCheckers();
```

```
for (PKIXCertPathChecker checker : checkers) {
    checker.init(false);
}
```

For each certificate that it validates, the service provider implementation must call the `check` method of each `PKIXCertPathChecker` object in turn, passing it the certificate and any remaining unresolved critical extensions:

```
for (PKIXCertPathChecker checker : checkers) {
    checker.check(cert, unresolvedCritExts);
}
```

If any of the `checks` throw a `CertPathValidatorException`, a `CertPathValidator` implementation should terminate the validation procedure. However, a `CertPathBuilder` implementation may simply log the failure and continue to search for other potential paths. If all of the `checks` are successful, the service provider implementation should check that all critical extensions have been resolved and if not, consider the validation to have failed. For example:

```
if (unresolvedCritExts != null &&
    !unresolvedCritExts.isEmpty())
{
    // note that a CertPathBuilder may have an enclosing
    // try block to catch the exception below and continue on error
    throw new CertPathValidatorException
        ("Unrecognized Critical Extension");
}
```

As discussed in the previous section, a `CertPathBuilder` implementation may need to backtrack when a potential certification path reaches a dead end or point of failure. Backtracking in this context implies returning to the previous certificate in the path and checking for other potential paths. If the `CertPathBuilder` implementation is validating the path as it is building it, it will need to restore the previous state of each `PKIXCertPathChecker`. It can do this by making clones of the `PKIXCertPathChecker` objects *before* each certificate is processed, for example:

```
/* clone checkers */
List newList = new ArrayList(checkers);
ListIterator li = newList.listIterator();
while (li.hasNext()) {
    PKIXCertPathChecker checker = (PKIXCertPathChecker) li.next();
    li.set(checker.clone());
}
```

Using PKIXCertPathChecker in Certificate Path Validation

Using a `PKIXCertPathChecker` to customize certificate path validation is relatively straightforward.

Basic Certification Path Validation

First, consider code that validates a certificate path:

```
Set<TrustAnchor> trustAnchors = getTrustAnchors();
CertPath cp = getCertPath();

PKIXParameters pkixp = new PKIXParameters(trustAnchors);
pkixp.setRevocationEnabled(false);

CertPathValidator cpv = CertPathValidator.getInstance("PKIX");
PKIXCertPathValidatorResult pcpvr =
    (PKIXCertPathValidatorResult)cpv.validate(cp, pkixp);
```

If the validation fails, the `validate()` method throws an exception.

The fundamental steps are as follows:

1. Obtain the CA root certificates and the certification path to be validated.
2. Create a `PKIXParameters` with the trust anchors.
3. Use a `CertPathValidator` to validate the certificate path.

In this example, `getTrustAnchors()` and `getCertPath()` are the methods that obtain CA root certificates and the certification path.

The `getTrustAnchors()` method in the example must return a `Set` of `TrustAnchors` that represent the CA root certificates you wish to use for validation. Here is one simple implementation that loads a single CA root certificate from a file:

```
public Set<TrustAnchor> getTrustAnchors()
    throws IOException, CertificateException {

  CertificateFactory cf = CertificateFactory.getInstance("X.509");

  X509Certificate c;
  try (InputStream in = new FileInputStream("x509_ca-certificate.cer")) {
    c = (X509Certificate)cf.generateCertificate(in);
  }

  TrustAnchor anchor = new TrustAnchor(c, null);
  return Collections.singleton(anchor);
}
```

Similarly, here is a simple implementation of `getCertPath()` that loads a certificate path from a file:

```
public CertPath getCertPath() throws IOException, CertificateException {
  CertificateFactory cf = CertificateFactory.getInstance("X.509");

  CertPath cp;
  try (InputStream in = new FileInputStream("certpath.pkcs7")) {
    cp = cf.generateCertPath(in, "PKCS7");
  }
  return cp;
}
```

Note that PKCS#7 does not require a specific order for the certificates in the file, so this code only works for certification path validation when the certificates are ordered starting from the entity to be validated and progressing back toward the CA root. If the certificates are not in the right order, you need to do some additional processing. `CertificateFactory` has a `generateCertPath()` method that accepts a `Collection`, which is useful for this type of processing.

Adding in a `PKIXCertPathChecker`

To customize certification path validation, add a `PKIXCertPathChecker` as follows. In this example, `SimpleChecker` is a `PKIXCertPathChecker` subclass. The new lines are shown in **bold**.

```
Set<TrustAnchor> trustAnchors = getTrustAnchors();
CertPath cp = getCertPath();

PKIXParameters pkixp = new PKIXParameters(trustAnchors);
pkixp.setRevocationEnabled(false);

SimpleChecker sc = new SimpleChecker();
pkixp.addCertPathChecker(sc);

CertPathValidator cpv = CertPathValidator.getInstance("PKIX");
PKIXCertPathValidatorResult pcpvr =
    (PKIXCertPathValidatorResult)cpv.validate(cp, pkixp);
```

`SimpleChecker` is a rudimentary subclass of `PKIXCertPathChecker`. Its `check()` method is called for every certificate in the certification path that is being validated. `SimpleChecker` uses an `AlgorithmConstraints` implementation to examine the signature algorithm and public key of each certificate.

```
import java.security.AlgorithmConstraints;
import java.security.CryptoPrimitive;
import java.security.Key;
import java.security.cert.*;
import java.util.*;

public class SimpleChecker extends PKIXCertPathChecker {
```

```
  private final static Set<CryptoPrimitive> SIGNATURE_PRIMITIVE_SET =
      EnumSet.of(CryptoPrimitive.SIGNATURE);

  public void init(boolean forward) throws CertPathValidatorException {}

  public boolean isForwardCheckingSupported() { return true; }

  public Set<String> getSupportedExtensions() { return null; }

  public void check(Certificate cert,
      Collection<String> unresolvedCritExts)
      throws CertPathValidatorException {
    X509Certificate c = (X509Certificate)cert;
    String sa = c.getSigAlgName();
    Key key = c.getPublicKey();

    AlgorithmConstraints constraints = new SimpleConstraints();

    if (constraints.permits(SIGNATURE_PRIMITIVE_SET, sa, null) == false)
      throw new CertPathValidatorException("Forbidden algorithm: " + sa);

    if (constraints.permits(SIGNATURE_PRIMITIVE_SET, key) == false)
      throw new CertPathValidatorException("Forbidden key: " + key);
  }
}
```

Finally, `SimpleConstraints` is an `AlgorithmConstraints` implementation that requires RSA 2048.

```
import java.security.AlgorithmConstraints;
import java.security.AlgorithmParameters;
import java.security.CryptoPrimitive;
import java.security.Key;
import java.security.interfaces.RSAKey;
import java.util.Set;

public class SimpleConstraints implements AlgorithmConstraints {
  public boolean permits(Set<CryptoPrimitive> primitives,
      String algorithm, AlgorithmParameters parameters) {
    return permits(primitives, algorithm, null, parameters);
  }

  public boolean permits(Set<CryptoPrimitive> primitives, Key key) {
    return permits(primitives, null, key, null);
  }

  public boolean permits(Set<CryptoPrimitive> primitives,
      String algorithm, Key key, AlgorithmParameters parameters) {
    if (algorithm == null) algorithm = key.getAlgorithm();

    if (algorithm.indexOf("RSA") == -1) return false;

    if (key != null) {
```

```
        RSAKey rsaKey = (RSAKey)key;
        int size = rsaKey.getModulus().bitLength();
        if (size < 2048) return false;
      }

      return true;
    }
  }
```

Check Revocation Status of Certificates with PKIXRevocationChecker Class

An instance of `PKIXRevocationChecker` checks the revocation status of certificates with the Online Certificate Status Protocol (OCSP) or Certificate Revocation Lists (CRLs).

The `PKIXRevocationChecker` (introduced in JDK 8), which is a subclass of `PKIXCertPathChecker`, checks the revocation status of certificates with the PKIX algorithm.

An instance of `PKIXRevocationChecker` checks the revocation status of certificates with the Online Certificate Status Protocol (OCSP) or Certificate Revocation Lists (CRLs). OCSP is described in RFC 2560 and is a network protocol for determining the status of a certificate. A CRL is a time-stamped list identifying revoked certificates, and RFC 5280 describes an algorithm for determining the revocation status of certificates using CRLs.

Each PKIX `CertPathValidator` and `CertPathBuilder` instance provides a default revocation implementation that is enabled by default. If you want more control over the revocation settings used by that implementation, use the `PKIXRevocationChecker` class.

Follow these general steps to check the revocation status of a certificate path with the `PKIXRevocationChecker` class:

1. Obtain a `PKIXRevocationChecker` instance by calling the `getRevocationChecker` method of a PKIX `CertPathValidator` or `CertPathBuilder` instance.
2. Set additional parameters and options specific to certificate revocation with methods contained in the `PKIXRevocationChecker` class. These methods include `setOCSPResponder(URI)`, which sets the URI that identifies the location of the OCSP responder (although normally the URI is included in the certificate and does not have to be set) and `setOptions(Set<PKIXRevocationChecker.Option>)`, which sets revocation options. `PKIXRevocationChecker.Option` is an enumerated type used to specify the following options:
   - `ONLY_END_ENTITY`: Only check the revocation status of end-entity certificates.
   - `PREFER_CRLS`: By default, OCSP is the preferred mechanism for checking revocation status, with CRLs as the fallback mechanism. Switch this preference to CRLs with this option.
   - `SOFT_FAIL`: Ignore network failures.
3. After obtaining an instance of `PKIXRevocationChecker`, add it to a `PKIXParameters` or `PKIXBuilderParameters` object with the `addCertPathChecker` or `setCertPathCheckers` method.

4. Follow one of these steps depending on whether you are using a PKIX `CertPathValidator` or `CertPathBuilder` instance:
   - If you are using a PKIX `CertPathValidator` instance, call the `validate` method using as arguments the certificate path you want to validate and the `PKIXParameters` object that contains a revocation checker.
   - If you are using a PKIX `CertPathBuilder` instance, call the `build` method using as arguments the `PKIXBuilderParameters` object that contains a revocation checker.
5. Call the `validate` method of the PKIX `CertPathValidator` or `CertPathBuilder` instance using as arguments the certificate path you want to validate and the `PKIXParameters` or `PKIXBuilderParameters` object that contains a revocation checker.

The following excerpt checks the revocation status of certificates contained in a certificate path. The `CertPath` object `path` is the certificate path, and `params` is an object of type `PKIXParameters`:

```
    CertPathValidator cpv = CertPathValidator.getInstance("PKIX");
    PKIXRevocationChecker rc =
(PKIXRevocationChecker)cpv.getRevocationChecker();
    rc.setOptions(EnumSet.of(Option.SOFT_FAIL));
    params.addCertPathChecker(rc);
    params.setRevocationEnabled(false);
    CertPathValidatorResult res = cpv.validate(path, params);
```

In this excerpt, the `SOFT_FAIL` option causes the revocation checker to ignore any network failures (such as failing to establish a connection to the OCSP server) when it checks the revocation status.

Implementing a Service Provider

Experienced programmers can create their own provider packages supplying certification path service implementations.

This section assumes that you have read Java Cryptography Architecture (JCA) Reference Guide.

The following engine classes are defined in the Java Certification Path API:

- **`CertPathValidator`** - used to validate certification paths
- **`CertPathBuilder`** - used to build certification paths
- **`CertStore`** - used to retrieve certificates and CRLs from a repository

In addition, the pre-existing `CertificateFactory` engine class also supports the generation of certification paths.

The application interfaces supplied by an engine class are implemented in terms of a "Service Provider Interface" (SPI). The name of each SPI class is the same as that of the corresponding engine class, followed by "Spi". For example, the SPI class corresponding to the `CertPathValidator` engine class is the `CertPathValidatorSpi` class. Each SPI class is abstract. To supply the implementation of a particular type of service, for a specific algorithm or type, a provider must subclass the corresponding

SPI class and provide implementations for all the abstract methods. For example, the `CertStore` class provides access to the functionality of retrieving certificates and CRLs from a repository. The actual implementation supplied in a `CertStoreSpi` subclass would be that for a specific type of certificate repository, such as LDAP.

Steps to Implement and Integrate a Provider

When implementing and integrating a provider for the certification path services, you must ensure that certain information is provided.

Developers should follow the Steps to Implement and Integrate a Provider. Here are some additional rules to follow for certain steps:

Step 3: Write your "Master Class", a subclass of Provider

In Step 3: Write Your Master Class, a Subclass of Provider these are the properties that must be defined for the certification path services, where the algorithm name is substituted for *algName*, and certstore type for *storeType*:

- `CertPathValidator.`*algName*
- `CertPathBuilder.`*algName*
- `CertStore.`*storeType*

See Java Security Standard Algorithm Names for the standard names that are defined for *algName* and *storeType*. The value of each property must be the fully qualified name of the class implementing the specified algorithm, or certstore type. That is, it must be the package name followed by the class name, where the two are separated by a period. For example, a provider sets the `CertPathValidator.PKIX` property to have the value `"sun.security.provider.certpath.PKIXCertPathValidator"` as follows:

```
put("CertPathValidator.PKIX",
"sun.security.provider.certpath.PKIXCertPathValidator")
```

In addition, service attributes can be defined for the certification path services. These attributes can be used as filters for selecting service providers. See Appendix A for the definition of some standard service attributes. For example, a provider may set the `ValidationAlgorithm` service attribute to the name of an RFC or specification that defines the PKIX validation algorithm:

```
put("CertPathValidator.PKIX ValidationAlgorithm", "RFC5280");
```

Step 11: Document your Provider and its Supported Services

In Step 12: Document Your Provider and Its Supported Services, certification path service providers should document the following information for each SPI:

Certificate Factories

A provider should document what types of certification paths (and the version numbers of the certificates in the path, if relevant) can be created by the factory. A provider should describe the ordering of the certificates in the certification path, as well as the contents.

A provider should document the list of encoding formats supported. This is not technically necessary, since the client can request them by calling the `getCertPathEncodings` method. However, the documentation should describe each encoding format in more detail and reference any standards when applicable.

Certification Path Validators

A provider should document any relevant information regarding the `CertPathValidator` implementation, including the types of certification paths that it validates. In particular, a PKIX `CertPathValidator` implementation should document the following information:

- The RFC or specification it is compliant with.
- The mechanism it uses to check that certificates have not been revoked.
- Any optional certificate or CRL extensions that it recognizes and how it processes them.

Certification Path Builders

A provider should document any relevant information regarding the `CertPathBuilder` implementation, including the types of certification paths that it creates and whether or not they are validated. In particular a PKIX `CertPathBuilder` implementation should document the following information:

- The RFC or specification it is compliant with.
- The mechanism it uses to check that certificates have not been revoked.
- Any optional certificate or CRL extensions that it recognizes and how it processes them.
- Details on the algorithm it uses for finding certification paths. Ex: depth-first, breadth-first, forward (i.e., from target to trust anchor(s)), reverse (i.e., from trust anchor(s) to target).
- The algorithm it uses to select and sort potential certificates. For example, given two certificates that are potential candidates for the next certificate in the path, what criteria are used to select one before the other? What criteria are used to reject a certificate?
- If applicable, the algorithm it uses for backtracking or constructing another path (i.e., when potential paths do not meet constraints).
- The types of `CertStore` implementations that have been tested. The implementation should be designed to work with any `CertStore` type, but this information may still be useful.

All `CertPathBuilder` implementations should provide additional debugging support, in order to analyze and correct potential path building problems. Details on how to access this debugging information should be documented.

Certificate/CRL Stores

A provider should document what types of certificates and CRLs (and the version numbers, if relevant) are retrieved by the `CertStore`.

A provider should also document any relevant information regarding the `CertStore` implementation (such as protocols used or formats supported). For example, an LDAP `CertStore` implementation should describe which versions of LDAP are supported and which standard attributes are used for finding certificates and CRLs. It should also

document if the implementation caches results, and for how long (i.e., under what conditions are they refreshed).

If the implementation returns the certificates and CRLs in a particular order, it should describe the sorting algorithm. An implementation should also document any additional or default initialization parameters. Finally, an implementation should document if and how it uses information in the `CertSelector` or `CRLSelector` objects to find certificates and CRLs.

Service Interdependencies

Common types of algorithm interdependencies in certification path service implementations.

The following are some common types of algorithm interdependencies in certification path service implementations:

- Certification Path Validation and Signature Algorithms

  A `CertPathValidator` implementation often requires use of a signature algorithm to verify each certificate's digital signature. The `setSigProvider` method of the `PKIXParameters` class allows a user to specify a specific `Signature` provider.

- Certification Path Builders and Certificate Factories

  A `CertPathBuilder` implementation will often utilize a `CertificateFactory` to generate a certification path from a list of certificates.

- CertStores and Certificate Factories

  A `CertStore` implementation will often utilize a `CertificateFactory` to generate certificates and CRLs from their encodings. For example, an LDAP `CertStore` implementation may use an X.509 `CertificateFactory` to generate X.509 certificates and CRLs from their ASN.1 encoded form.

Certification Path Parameter Specification Interfaces

The Certification Path API contains two interfaces representing *transparent* specifications of parameters, the `CertPathParameters` and `CertStoreParameters` interfaces.

Two implementations of the `CertPathParameters` interface are included, the `PKIXParameters` and `PKIXBuilderParameters` classes. If you are working with PKIX certification path validation and algorithm parameters, you can utilize these classes. If you need parameters for a different algorithm, you will need to supply your own `CertPathParameters` implementation for that algorithm.

Two implementations of the `CertStoreParameters` interface are included, the `LDAPCertStoreParameters` and the `CollectionCertStoreParameters` classes. These classes are to be used with LDAP and Collection `CertStore` implementations, respectively. If you need parameters for a different repository type, you will need to supply your own `CertStoreParameters` implementation for that type.

The `CertPathParameters` and `CertStoreParameters` interfaces each define a `clone` method that implementations should override. A typical implementation will perform a "deep" copy of the object, such that subsequent changes to the copy will not affect the original (and vice versa). However, this is not an absolute requirement for implementations of `CertStoreParameters`. A shallow copy implementation of `clone` is

more appropriate for applications that need to hold a reference to a parameter contained in the `CertStoreParameters`. For example, since `CertStore.getInstance` makes a clone of the specified `CertStoreParameters`, a shallow copy `clone` allows an application to hold a reference to and later release the resources of a particular `CertStore` initialization parameter, rather than waiting for the garbage collection mechanism. This should be done with the utmost care, since the `CertStore` may still be in use by other threads.

Certification Path Result Specification Interfaces

The Certification Path API contains two interfaces representing *transparent* specifications of results, the `CertPathValidatorResult` and `CertPathBuilderResult` interfaces.

One implementation for each of the interfaces is included: the `PKIXCertPathValidatorResult` and `PKIXCertPathBuilderResult` classes. If you are implementing PKIX certification path service providers, you can utilize these classes. If you need certification path results for a different algorithm, you will need to supply your own `CertPathValidatorResult` or `CertPathBuilderResult` implementation for that algorithm.

A PKIX implementation of a `CertPathValidator` or a `CertPathBuilder` may find it useful to store additional information in the `PKIXCertPathValidatorResult` or `PKIXCertPathBuilderResult`, such as debugging traces. In these cases, the implementation should implement a subclass of the appropriate result class with methods to retrieve the relevant information. These classes must be shipped with the provider classes, for example, as part of the provider JAR file.

Certification Path Exception Classes

The Certification Path API contains a set of exception classes for handling errors. `CertPathValidatorException, CertPathBuilderException`, and `CertStoreException` are subclasses of `GeneralSecurityException`.

You may need to extend these classes in your service provider implementation.

For example, a `CertPathBuilder` implementation may provide additional information such as debugging traces when a `CertPathBuilderException` is thrown. The implementation may throw a subclass of `CertPathBuilderException` that holds this information. Likewise, a `CertStore` implementation can provide additional information when a failure occurs by throwing a subclass of `CertStoreException` . Also, you may want to implement a subclass of `CertPathValidatorException` to describe a particular failure mode of your `CertPathValidator` implementation.

In each case, the new exception classes must be shipped with the provider classes, for example, as part of the provider JAR file. Each provider should document the exception subclasses.

Appendix A: Standard Names

The Java Certification Path API requires and utilizes a set of standard names for certification path validation algorithms, encodings and certificate storage types.

The standard names previously found here in Appendix A and in the other security specifications (JCA/JSSE/etc.) have been combined in the Java Security Standard Algorithm Names. Specific provider information can be found in the JDK Providers.

Please note that a service provider may choose to define a new name for a proprietary or non-standard algorithm that is not mentioned in the Standard Names document. However, to prevent name collisions, it is recommended that the name be prefixed with the reverse Internet domain name of the provider's organization (for example: `com.sun.MyCertPathValidator`).

Appendix B: CertPath Implementation in SUN Provider

The "SUN" provider supports the following standard algorithms, types and encodings:

- `CertificateFactory`: **X.509** `CertPath` type with **PKCS7** and **PkiPath** encodings
- `CertPathValidator`: **PKIX** algorithm
- `CertPathBuilder`: **PKIX** algorithm
- `CertStore`: **Collection** `CertStore` type

Each of these service provider interface implementations is discussed in more detail below.

CertificateFactory

The "SUN" provider for the `CertificateFactory` engine class supports generation of X.509 `CertPath` objects. The PKCS7 and PkiPath encodings are supported. The PKCS#7 implementation supports a subset of RFC 2315 (only the SignedData ContentInfo type is supported). The certificates in the `CertPath` are ordered in the forward direction (from target to trust anchor). Each certificate in the `CertPath` is of type `java.security.cert.X509Certificate` , and versions 1, 2 and 3 are supported.

CertPathValidator

The "SUN" provider supplies a PKIX implementation of the `CertPathValidator` engine class. The implementation validates `CertPaths` of type X.509 and implements the certification path validation algorithm defined in RFC 5280: *PKIX Certificate and CRL Profile*. This implementation sets the `ValidationAlgorithm` service attribute to "RFC5280".

Weak cryptographic algorithms can be disabled in the "SUN" provider using the `jdk.certpath.disabledAlgorithms` Security Property. See Appendix E: Disabling Cryptographic Algorithms for a description and examples of this property.

The PKIX Certificate and CRL Profile has many optional features. The "SUN" provider implements support for the policy mapping, authority information access and CRL distribution point certificate extensions, the issuing distribution point CRL extension, and the reason code and certificate issuer CRL entry extensions. It does not implement support for the freshest CRL or subject information access certificate extensions. It also does not include support for the freshest CRL and delta CRL Indicator CRL extensions and the invalidity date and hold instruction code CRL entry extensions.

The implementation supports a CRL revocation checking mechanism that conforms to section 6.3 of the PKIX Certificate and CRL Profile. OCSP (RFC 2560) is also currently supported as a built in revocation checking mechanism. See Appendix C:

OCSP Support for more details on the implementation and configuration and how it works in conjunction with CRLs.

The implementation does not support the `nameConstraints` parameter of the `TrustAnchor` class and the `validate` method throws an `InvalidAlgorithmParameterException` if it is specified.

CertPathBuilder

The "SUN" provider supplies a PKIX implementation of the `CertPathBuilder` engine class. The implementation builds `CertPaths` of type X.509. Each `CertPath` is validated according to the PKIX algorithm defined in RFC 5280: *PKIX Certificate and CRL Profile*. This implementation sets the `ValidationAlgorithm` service attribute to "RFC5280".

The implementation requires that the `targetConstraints` parameter of a `PKIXBuilderParameters` object is an instance of `X509CertSelector` and the subject criterion is set to a non-null value. Otherwise the `build` method throws an `InvalidAlgorithmParameterException`.

The implementation builds `CertPath` objects in a forward direction using a depth-first algorithm. It backtracks to previous states and tries alternate paths when a potential path is determined to be invalid or exceeds the `PKIXBuilderParameters` `maxPathLength` parameter.

Validation of the path is performed in the same manner as the `CertPathValidator` implementation. The implementation validates most of the path as it is being built, in order to eliminate invalid paths earlier in the process. Validation checks that cannot be executed on certificates ordered in a forward direction are delayed and executed on the path after it has been constructed (but before it is returned to the application).

As with `CertPathValidator`, the `jdk.certpath.disabledAlgorithms` Security Property can be used to exclude cryptographic algorithms that are not considered safe.

When two or more potential certificates are discovered that may lead to finding a path that meets the specified constraints, the implementation uses the following criteria to prioritize the certificates (in the examples below, assume a `TrustAnchor` distinguished name of "ou=D,ou=C,o=B,c=A" is specified):

1. The issuer DN of the certificate matches the DN of one of the specified `TrustAnchors` (ex: issuerDN = "ou=D,ou=C,o=B,c=A").
2. The issuer DN of the certificate is a descendant of the DN of one of the `TrustAnchors`, ordered by proximity to the anchor (ex: issuerDN = "ou=E,ou=D,ou=C,o=B,c=A").
3. The issuer DN of the certificate is an ancestor of the DN of one of the `TrustAnchors`, ordered by proximity to the anchor (ex: issuerDN = "ou=C,o=B,c=A".
4. The issuer DN of the certificate is in the same namespace of one of the `TrustAnchors`, ordered by proximity to the anchor (ex: issuerDN = "ou=G,ou=C,o=B,c=A").
5. The issuer DN of the certificate is an ancestor of the subject DN of the certificate, ordered by proximity to the subject.

These are followed by certificates which don't meet any of the above criteria.

This implementation has been tested with the LDAP and Collection `CertStore` implementations included in this release of the "SUN" provider.

Debugging support can be enabled by setting the `java.security.debug` property to `certpath`. For example:

```
java -Djava.security.debug=certpath BuildCertPath
```

This will print additional debugging information to standard error.

Collection CertStore

The SUN provider supports the Collection implementation of the `CertStore` engine class.

The Collection `CertStore` implementation can hold any objects that are an instance of `java.security.cert.Certificate` or `java.security.cert.CRL`.

The certificates and CRLs are not returned in any particular order and will not contain duplicates.

Support for the CRL Distribution Points Extension

Support for the CRL Distribution Points extension is available. It is disabled by default for compatibility and can be enabled by setting the system property `com.sun.security.enableCRLDP` to the value true.

If set to true, Sun's PKIX implementation uses the information in a certificate's CRL Distribution Points extension (in addition to `CertStores` that are specified) to find the CRL, provided the distribution point is an X.500 distinguished name or a URI of type ldap, http, or ftp.

> **Note:**
>
> Depending on your network and firewall setup, it may be necessary to also configure your networking proxy servers.

Support for the Authority Information Access (AIA) Extension

Support for the `caIssuers` access method of the Authority Information Access extension is available. It is disabled by default for compatibility and can be enabled by setting the system property `com.sun.security.enableAIAcaIssuers` to the value true.

If set to true, Sun's PKIX implementation of `CertPathBuilder` uses the information in a certificate's AIA extension (in addition to CertStores that are specified) to find the issuing CA certificate, provided it is a URI of type ldap, http, or ftp.

> **Note:**
>
> Depending on your network and firewall setup, it may be necessary to also configure your networking proxy servers.

Maximum Network Connection Timeout for CRL Retrievals

Set the maximum connection timeout for CRL retrievals, in seconds, with the system property `com.sun.security.crl.timeout`. If the property has not been set, or if its value is negative, it's set to the default value of 15 seconds. A value of 0 means an infinite timeout.

Appendix C: OCSP Support

Client-side support for the On-Line Certificate Status Protocol (OCSP) as defined in RFC 2560 is supported.

OCSP checking is controlled by the following five Security Properties:

Property Name	Description
`ocsp.enable`	This property's value is either true or false. If true, OCSP checking is enabled when doing certificate revocation checking; if false or not set, OCSP checking is disabled.
`ocsp.responderURL`	This property's value is a URL that identifies the location of the OCSP responder. Here is an example `ocsp.responderURL=http://ocsp.example.net:80` By default, the location of the OCSP responder is determined implicitly from the certificate being validated. The property is used when the Authority Information Access extension (defined in RFC 5280) is absent from the certificate or when it requires overriding.

Property Name	Description
`ocsp.responderCertSubjectName`	This property's value is the subject name of the OCSP responder's certificate. Here is an example `ocsp.responderCertSubjectName="CN=OCSP Responder, O=XYZ Corp"` By default, the certificate of the OCSP responder is that of the issuer of the certificate being validated. This property identifies the certificate of the OCSP responder when the default does not apply. Its value is a string distinguished name (defined in RFC 2253) which identifies a certificate in the set of certificates supplied during cert path validation. In cases where the subject name alone is not sufficient to uniquely identify the certificate, then both the `ocsp.responderCertIssuerName` and `ocsp.responderCertSerialNumber` properties must be used instead. When this property is set, then those two properties are ignored.
`ocsp.responderCertIssuerName`	This property's value is the issuer name of the OCSP responder's certificate . Here is an example `ocsp.responderCertIssuerName="CN=Enterprise CA, O=XYZ Corp"` By default, the certificate of the OCSP responder is that of the issuer of the certificate being validated. This property identifies the certificate of the OCSP responder when the default does not apply. Its value is a string distinguished name (defined in RFC 2253) which identifies a certificate in the set of certificates supplied during cert path validation. When this property is set then the `ocsp.responderCertSerialNumber` property must also be set. Note that this property is ignored when the `ocsp.responderCertSubjectName` property has been set.

Property Name	Description
`ocsp.responderCertSerialNumber`	This property's value is the serial number of the OCSP responder's certificate Here is an example `ocsp.responderCertSerialNumber=2A:FF:00` By default, the certificate of the OCSP responder is that of the issuer of the certificate being validated. This property identifies the certificate of the OCSP responder when the default does not apply. Its value is a string of hexadecimal digits (colon or space separators may be present) which identifies a certificate in the set of certificates supplied during cert path validation. When this property is set then the `ocsp.responderCertIssuerName` property must also be set. Note that this property is ignored when the `ocsp.responderCertSubjectName` property has been set.

These properties may be set either statically in the Java runtime's `<java_home>/conf/security/java.security` file, or dynamically using the `java.security.Security.setProperty()` method.

By default, OCSP checking is not enabled. It is enabled by setting the `ocsp.enable` property to `"true"`. Use of the remaining properties is optional. Note that enabling OCSP checking only has an effect if revocation checking has also been enabled. Revocation checking is enabled via the `PKIXParameters.setRevocationEnabled()` method.

OCSP checking works in conjunction with Certificate Revocation Lists (CRLs) during revocation checking. Below is a summary of the interaction of OCSP and CRLs. Failover to CRLs occurs only if an OCSP problem is encountered. Failover does not occur if the OCSP responder confirms either that the certificate has been revoked or that it has not been revoked.

PKIXParameters RevocationEnabled (default=true)	`ocsp.enable` (default=false)	Behavior
true	true	Revocation checking using OCSP, failover to using CRLs
true	false	Revocation checking using CRLs only
false	true	No revocation checking
false	false	No revocation checking

Maximum Allowable Clock Skew

You might encounter connection failures during revocation checking because the network is slow or the system clock is off by some amount. Set the maximum allowable clock skew (the time difference between response time and local time), in seconds, used for revocation checks with the system property `com.sun.security.ocsp.clockskew`. If the property has not been set, or if its value is negative, it's set to the default value of 900 seconds (15 minutes).

Appendix D: CertPath Implementation in JdkLDAP Provider

The JdkLDAP provider supports the LDAP implementation of the `CertStore` engine class.

LDAP CertStore

The LDAP `CertStore` implementation retrieves certificates and CRLs from an LDAP directory using the LDAP schema defined in RFC 2587.

The LDAPSchema service attribute is set to "RFC2587".

The implementation fetches certificates from different locations, depending on the values of the subject, issuer, and basicConstraints selection criteria specified in the `X509CertSelector`. It performs as many of the following operations as possible:

1. Subject non-null, basicConstraints <= -1
   Looks for certificates in the subject DN's "userCertificate" attribute.
2. Subject non-null, basicConstraints >= -1
   Looks for certificates in the forward element of the subject DN's "crossCertificatePair" attribute AND in the subject's "caCertificate" attribute.
3. Issuer non-null, basicConstraints >= -1
   Looks for certificates in the reverse element of the issuer DN's "crossCertificatePair" attribute AND in the issuer DN's "caCertificate" attribute.

In each case, certificates are checked using `X509CertSelector.match()` before adding them to the resulting collection.

If none of the conditions specified above applies, then an exception is thrown to indicate that it was impossible to fetch certificates using the criteria supplied. Note that even if one or more of the conditions apply, the Collection returned may still be empty if there are no certificates in the directory.

The implementation fetches CRLs from the issuer DNs specified in the `setCertificateChecking`, `addIssuerName` or `setIssuerNames` methods of the `X509CRLSelector` class. If no issuer DNs have been specified using one of these methods, the implementation throws an exception indicating it was impossible to fetch CRLs using the criteria supplied. Otherwise, the CRLs are searched as follows:

1. The implementation first creates a list of issuer names. If a certificate was specified in the `setCertificateChecking` method, it uses the issuer of that certificate. Otherwise, it uses the issuer names specified using the `addIssuerName` or `setIssuerNames` methods.
2. Next, the implementation iterates through the list of issuer names. For each issuer name, it searches first in the issuer's `"authorityRevocationList"` attribute and then, if no matching CRL was found there, in the issuer's

"certificateRevocationList" attribute. One exception to the above is that if the issuer name was obtained from the certificate specified in the setCertificateChecking method, it only checks the issuer's "authorityRevocationList" attribute if the specified certificate is a CA certificate.

3. All CRLs are checked using X509CRLSelector.match() before adding them to the resulting collection.
4. If no CRLs satisfying the selection criteria can be found, an empty Collection is returned.

Caching

By default each LDAP CertStore instance caches lookups for a maximum of 30 seconds. The cache lifetime can be changed by setting the system property sun.security.certpath.ldap.cache.lifetime to a value in seconds. A value of 0 disables the cache completely. A value of -1 means unlimited lifetime.

Appendix E: Disabling Cryptographic Algorithms

The jdk.certpath.disabledAlgorithms Security Property contains a list of cryptographic algorithms and key size constraints that are considered weak or broken. Certificates and other data (CRLs, OCSPResponses) containing any of these algorithms or key sizes will be blocked during certification path building and validation. This property is used by Oracle's PKIX implementation, other implementations might not examine and use it.

The exact syntax of the jdk.certpath.disabledAlgorithms property is described in the java.security file. In Java SE 9, the default value of the property is:

```
jdk.certpath.disabledAlgorithms=MD2, MD5, SHA1 jdkCA & usage TLSServer, \
    RSA keySize < 1024, DSA keySize < 1024, EC keySize < 224
```

In this syntax:

MD2
Any MD2-based algorithm will be blocked.
For example, a certificate, CRL, or OCSPResponse signed with an MD2withRSA signature algorithm.

MD5
Any MD5-based algorithm will be blocked.
For example, a certificate, CRL, or OCSPResponse signed with an MD5withRSA signature algorithm.

SHA1 jdkCA & usage TLSServer
All SHA1 certificates that chain to trust anchors pre-installed in the cacerts keystore and that are used for authentication of TLS Servers. See JEP 288.

RSA keySize < 1024
Any RSA key less than 1024 bits will be blocked.
For example, a certificate with a 768-bit RSA public key.

DSA keySize < 1024
Any DSA key less than 1024 bits will be blocked.

For example, a certificate with a 512-bit DSA public key.

EC keySize < 224
Any EC key less than 224 bits will be blocked.
For example, a certificate with a 160-bit EC public key.

Administrators or users can modify the value of the `jdk.certpath.disabledAlgorithms` property to address additional security requirements. However, removing any of the current algorithms or key sizes is not recommended.

11

Java SASL API Programming and Deployment Guide

Simple Authentication and Security Layer, or SASL, is an Internet standard (RFC 2222) that specifies a protocol for authentication and optional establishment of a security layer between client and server applications. SASL defines how authentication data is to be exchanged but does not itself specify the contents of that data. It is a framework into which specific authentication mechanisms that specify the contents and semantics of the authentication data can fit.

SASL is used by protocols, such as the Lightweight Directory Access Protocol, version 3 (LDAP v3), and the Internet Message Access Protocol, version 4 (IMAP v4) to enable pluggable authentication. Instead of hardwiring an authentication method into the protocol, LDAP v3 and IMAP v4 use SASL to perform authentication, thus enabling authentication via various SASL mechanisms.

There are a number of standard SASL mechanisms defined by the Internet community for various levels of security and deployment scenarios. These range from no security (for example, anonymous authentication) to high security (for example, Kerberos authentication) and levels in between.

The Java SASL API

The Java SASL API defines classes and interfaces for applications that use SASL mechanisms. It is defined to be mechanism-neutral: the application that uses the API need not be hardwired into using any particular SASL mechanism. The API supports both client and server applications. It allows applications to select the mechanism to use based on desired security features, such as whether they are susceptible to passive dictionary attacks or whether they accept anonymous authentication.

The Java SASL API also allows developers to use their own, custom SASL mechanisms. SASL mechanisms are installed by using the Java Cryptography Architecture (JCA); see Java Cryptography Architecture (JCA) Reference Guide.

When to Use SASL

SASL provides a pluggable authentication and security layer for network applications. There are other features in Java SE that provide similar functionality, including Java Secure Socket Extension (JSSE) (see Java Secure Socket Extension (JSSE) Reference Guide) and the Java Generic Security Service. JSSE provides a framework and an implementation for a Java language version of the SSL, TLS, and DTLS protocols. Java GSS is the Java language bindings for the Generic Security Service Application Programming Interface (GSS-API). The only mechanism currently supported underneath this API on Java SE is Kerberos v5.

With the exception of defining and building protocols from scratch, protocol definition is often the biggest factor that goes into determining which API to use. When compared with JSSE and Java GSS, SASL is relatively lightweight and is popular among some protocols. It also has the advantage that several popular, lightweight (in terms of infrastructure support) SASL mechanisms have been defined. Primary JSSE and Java

GSS mechanisms, on the other hand, have relatively heavyweight mechanisms that require more elaborate infrastructures (Public Key Infrastructure and Kerberos, respectively).

SASL, JSSE, and Java GSS are often used together. For example, a common pattern is for an application to use JSSE for establishing a secure channel, and to use SASL for client, username/password-based authentication. There are also SASL mechanisms layered on top of GSS-API mechanisms; one popular example is a SASL GSS-API/Kerberos v5 mechanism that is used with LDAP.

With the exception of defining and building protocols from scratch, protocol definition is often the biggest factor in determining which API to use. For example, LDAP and IMAP are defined to use SASL, so software related to these protocols should use the Java SASL API. When building Kerberos applications and services, the API to use is Java GSS. When building applications and services that use SSL/TLS as their protocol, the API to use is JSSE.

Java SASL API Overview

SASL is a challenge-response protocol. The server issues a challenge to the client, and the client sends a response based on the challenge. This exchange continues until the server is satisfied and issues no further challenge. These challenges and responses are binary tokens of arbitrary length. The encapsulating protocol (such as LDAP or IMAP) specifies how these tokens are encoded and exchanged. For example, LDAP specifies how SASL tokens are encapsulated within LDAP bind requests and responses.

The Java SASL API is modeled according to this style of interaction and usage. It has interfaces, `SaslClient` and `SaslServer`, that represent client-side and server-side mechanisms, respectively. The application interacts with the mechanisms via byte arrays that represent the challenges and responses. The server-side mechanism iterates, issuing challenges and processing responses, until it is satisfied, while the client-side mechanism iterates, evaluating challenges and issuing responses, until the server is satisfied. The application that is using the mechanism drives each iteration. That is, it extracts the challenge or response from a protocol packet and supplies it to the mechanism, and then puts the response or challenge returned by the mechanism into a protocol packet and sends it to the peer.

Creating the Mechanisms

The client and server code that use the SASL mechanisms are not hardwired to use specific mechanism(s). In many protocols that use SASL, the server advertises (either statically or dynamically) a list of SASL mechanisms that it supports. The client then selects one of these based on its security requirements.

The `Sasl` class is used for creating instances of `SaslClient` and `SaslServer`. Here is an example of how an application creates a SASL client mechanism using a list of possible SASL mechanisms.

```
String[] mechanisms = new String[]{"DIGEST-MD5", "PLAIN"};
SaslClient sc = Sasl.createSaslClient(
    mechanisms, authzid, protocol, serverName, props, callbackHandler);
```

Based on the availability of the mechanisms supported by the platform and other configuration information provided via the parameters, the Java SASL framework selects one of the listed mechanisms and return an instance of `SaslClient`.

The name of the selected mechanism is usually transmitted to the server via the application protocol. Upon receiving the mechanism name, the server creates a corresponding `SaslServer` object to process client-sent responses. Here is an example of how the server would create an instance of `SaslServer`.

```
SaslServer ss = Sasl.createSaslServer(
    mechanism, protocol, myName, props, callbackHandler);
```

Passing Input to the Mechanisms

Because the Java SASL API is a general framework, it must be able to accommodate many different types of mechanisms. Each mechanism needs to be initialized with input and may need input to make progress. The API provides three means by which an application gives input to a mechanism:

1. **Common input parameters**: The application uses predefined parameters to supply information that are defined by the SASL specification and commonly required by mechanisms. For`SaslClient` mechanisms, the input parameters are authorization id, protocol id, and server name. For`SaslServer` mechanisms, the common input parameters are protocol id and (its own fully qualified) server name.
2. **Properties parameter**: The application uses the properties parameter, a mapping of property names to (possibly non-string) property values, to supply configuration information. The Java SASL API defines some standard properties, such as `Sasl.QOP` (quality-of-protection), `Sasl.STRENGTH` (cipher strength), and `Sasl.MAX_BUFFER` (maximum buffer size). The parameter can also be used to pass in non-standard properties that are specific to particular mechanisms.
3. **Callbacks**: The application uses the `CallbackHandler` parameter to supply input that cannot be predetermined or might not be common across mechanisms. When a mechanism requires input data, it uses the callback handler supplied by the application to collect the data, possibly from the end-user of the application. For example, a mechanism might require the end-user of the application to supply a name and password.

   Mechanisms can use the callbacks defined in the `javax.security.auth.callback` package; these are generic callbacks useful for building applications that perform authentication. Mechanisms might also need SASL-specific callbacks, such as those for collecting realm and authorization information, or even (non-standardized) mechanism-specific callbacks. The application should be able to accommodate a variety of mechanisms. Consequently, its callback handler must be able to service all of the callbacks that the mechanisms might request. This is not possible in general for arbitrary mechanisms, but is usually feasible due to the limited number of mechanisms that are typically deployed and used.

Using the Mechanisms

Once the application has created a mechanism, it uses the mechanism to obtain SASL tokens to exchange with the peer. The client typically indicates to the server via the application protocol which mechanism to use. Some protocols allow the client to

accompany the request with an optional initial response for mechanisms that have an initial response. This feature can be used to lower the number of message exchanges required for authentication. Here is an example of how a client might use `SaslClient` for authentication.

```
    // Get optional initial response
    byte[] response =
        (sc.hasInitialResponse() ? sc.evaluateChallenge(new byte[]) :
null);

    String mechanism = sc.getMechanismName();

    // Send selected mechanism name and optional initial response to server
    send(mechanism, response);

    // Read response
    msg = receive();
    while (!sc.isComplete() && (msg.status == CONTINUE || msg.status ==
SUCCESS)) {
        // Evaluate server challenge
        response = sc.evaluateChallenge(msg.contents);

        if (msg.status == SUCCESS) {
            // done; server doesn't expect any more SASL data
             if (response != null) {
                throw new IOException(
                    "Protocol error: attempting to send response after
completion");
            }
            break;
        } else {
            send(mechanism, response);
            msg = receive();
        }
    }
```

The client application iterates through each step of the authentication by using the mechanism (`sc`) to evaluate the challenge gotten from the server and to get a response to send back to the server. It continues this cycle until either the mechanism or application-level protocol indicates that the authentication has completed, or if the mechanism cannot evaluate a challenge. If the mechanism cannot evaluate the challenge, it throws an exception to indicate the error and terminates the authentication. Disagreement between the mechanism and protocol about the completion state must be treated as an error because it might indicate a compromise of the authentication exchange.

Here is an example of how a server might use `SaslServer`.

```
    // Read request that contains mechanism name and optional initial
response
    msg.receive();

    // Obtain a SaslServer to perform authentication
    SaslServer ss = Sasl.createSaslServer(msg.mechanism,
```

```
        protocol, myName, props, callbackHandler);

    // Perform authentication steps until done
    while (!ss.isComplete()) {
        try {
            // Process response
            byte[] challenge = sc.evaluateResponse(msg.contents);

            if (ss.isComplete()) {
                send(mechanism, challenge, SUCCESS);
            } else {
                send(mechanism, challenge, CONTINUE);
                msg.receive();
            }
        } catch (SaslException e) {
            send(ERROR);
            sc.dispose();
            break;
        }
    }
```

The server application iterates through each step of the authentication by giving the client's response to the mechanism (`ss`) to process. If the response is incorrect, the mechanism indicates the error by throwing a `SaslException` so that the server can report the error and terminate the authentication. If the response is correct, the mechanism returns challenge data to be sent to the client and indicates whether the authentication is complete. Note that challenge data can accompany a "success" indication. This might be used, for example, to tell the client to finalize some negotiated state.

Using the Negotiated Security Layer

Some SASL mechanisms support only authentication while others support use of a negotiated security layer after authentication. The security layer feature is often not used when the application uses some other means, such as SSL/TLS, to communicate securely with the peer.

When a security layer has been negotiated, all subsequent communication with the peer must take place using the security layer. To determine whether a security layer has been negotiated, get the negotiated `Sasl.QOP` from the mechanism. Here is an example of how to determine whether a security layer has been negotiated.

```
String qop = (String) sc.getNegotiatedProperty(Sasl.QOP);
boolean hasSecurityLayer = (qop != null &&
    (qop.equals("auth-int") || qop.equals("auth-conf")));
```

A security layer has been negotiated if the `Sasl.QOP` property indicates that either integrity and/or confidentiality has been negotiated.

To communicate with the peer using the negotiated layer, the application first uses the `wrap` method to encode the data to be sent to the peer to produce a "wrapped" buffer. It then transfers a length field representing the number of octets in the wrapped buffer followed by the contents of the wrapped buffer to the peer. The peer receiving the

stream of octets passes the buffer (without the length field) to `unwrap` to obtain the decoded bytes sent by the peer. Details of this protocol are described in RFC 2222. Example 11-1 illustrates how a client application sends and receives application data using a security layer.

Example 11-1 Sample Code for SASL Client Send and Receive Data

```
// Send outgoing application data to peer
byte[] outgoing = ...;
byte[] netOut = sc.wrap(outgoing, 0, outgoing.length);

send(netOut.length, netOut);   // send to peer

// Receive incoming application data from peer
byte[] netIn = receive();      // read length and ensuing bytes from peer

byte[] incoming = sc.unwrap(netIn, 0, netIn.length);
```

How SASL Mechanisms are Installed and Selected

SASL mechanism implementations are provided by SASL security providers. Each provider may support one or more SASL mechanisms and is registered with the JCA.

By default, the SunSASL provider is automatically registered as a JCA provider. To remove it or reorder its priority as a JCA provider, change the line

```
security.provider.7=SunSASL
```

in the Java security properties file (*java-home*/`conf/security/java.security`).

To add or remove a SASL provider, you add or remove the corresponding line in the security properties file. For example, if you want to add a SASL provider and have its mechanisms be chosen over the same ones implemented by the SunSASL provider, then you would add a line to the security properties file with a lower number.

```
security.provider.7=com.example.MyProvider
security.provider.8=SunSASL
```

Alternatively, you can programmatically add your own provider using the `java.security.Security` class. For example, the following sample code registers the `com.example.MyProvider` to the list of available SASL security providers.

```
Security.addProvider(new com.example.MyProvider());
```

See Step 8: Prepare for Testing in Steps to Implement and Integrate a Provider for more information about adding providers to the security properties file and programmatically adding your own providers.

When an application requests a SASL mechanism by supplying one or more mechanism names, the SASL framework looks for registered SASL providers that support that mechanism by going through, in order, the list of registered providers. The providers must then determine whether the requested mechanism matches the

selection policy properties in the `Sasl` and if so, return an implementation for the mechanism.

The selection policy properties specify the security aspects of a mechanism, such as its susceptibility to certain attacks. These are characteristics of the mechanism (definition), rather than its implementation so all providers should come to the same conclusion about a particular mechanism. For example, the PLAIN mechanism is susceptible to plaintext attacks regardless of how it is implemented. If no selection policy properties are supplied, there are no restrictions on the selected mechanism. Using these properties, an application can ensure that it does not use unsuitable mechanisms that might be deployed in the execution environment. For example, an application might use the following sample code if it does not want to allow the use of mechanisms susceptible to plaintext attacks.

```
Map<String, String> props = new HashMap<>();
props.put(Sasl.POLICY_NOPLAINTEXT, "true");
SaslClient sc = Sasl.createSaslClient(
    mechanisms, authzid, protocol, serverName, props, callbackHandler);
```

The SunSASL Provider

The SunSASL provider supports the following client and server mechanisms:

> **Note:**
>
> Some of these mechanisms are weak; you're responsible for determining whether the algorithm meets the security requirements of your applications.

- Client Mechanisms
  - PLAIN (RFC 2595). This mechanism supports cleartext user name/password authentication.
  - CRAM-MD5 (RFC 2195). This mechanism supports a hashed user name/ password authentication scheme.
  - DIGEST-MD5 (RFC 2831). This mechanism defines how HTTP Digest Authentication can be used as a SASL mechanism.
  - EXTERNAL (RFC 2222). This mechanism obtains authentication information from an external channel (such as TLS or IPsec).
  - NTLM. This mechanism supports NTLM authentication.
- Server Mechanisms
  - CRAM-MD5
  - DIGEST-MD5
  - NTLM

The SunSASL Provider Client Mechanisms

The SunSASL provider supports several SASL client mechanisms used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the client mechanisms and their required input.

Table 11-1 SunSASL Provider Client Mechanisms

Client Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy
CRAM-MD5	authorization id (as default user name)	`PasswordCallback` `NameCallback`	None	`Sasl.POLICY_NOANONYMOUS` `Sasl.POLICY_NOPLAINTEXT`

DIGEST-MD5	authorization id
	protocol id
	server name
		`NameCallback`
		`PasswordCallback`
		`RealmCallback`
		`RealmChoiceCallback`
			`Sasl.QOP`
			`Sasl.STRENGTH`
			`Sasl.MAX_BUFFER`
			`Sasl.SERVER_AUTH`
			`javax.security.sasl.sendmaxbuffer`
			`com.sun.security.sasl.digest.cipher`
				`Sasl.POLICY_NOANONYMOUS`
				`Sasl.POLICY_NOPLAINTEXT`

EXTERNAL	authorization id
	external channel
		None	None	`Sasl.POLICY_NOPLAINTEXT`
				`Sasl.POLICY_NOACTIVE`
				`Sasl.POLICY_NODICTIONARY`

NTLM	authzId (as default user name)
	serverName (as default domain)
		`RealmCallback`
		`NameCallback`
		`PasswordCallback`
			`Sasl.QOP`
			`com.sun.security.sasl.ntlm.version`
			`com.sun.security.sasl.ntlm.random`
			`com.sun.security.sasl.ntlm.hostname`
				`Sasl.POLICY_NOANONYMOUS`
				`Sasl.POLICY_NOPLAINTEXT`

PLAIN	authorization id	`NameCallback` `PasswordCallback`	None	`Sasl.POLICY_NOANONYMOUS`

An application that uses these mechanisms from the SunSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults. Most of the parameters, callbacks, and properties are described in the API documentation. The following sections describe mechanism-specific behaviors and parameters not already covered by the API documentation.

Cram-MD5

The Cram-MD5 client mechanism uses the authorization id parameter, if supplied, as the default user name in the `NameCallback` to solicit the application/end-user for the authentication id. The authorization id is otherwise not used by the Cram-MD5 mechanism; only the authentication id is exchanged with the server.

Digest-MD5

The Digest-MD5 mechanism is used for digest authentication and optional establishment of a security layer. It specifies the following ciphers for use with the security layer: Triple DES, DES and RC4 (128, 56, and 40 bits). The Digest-MD5 mechanism can support only ciphers that are available on the platform. For example, if the platform does not support the RC4 ciphers, then the Digest-MD5 mechanism will not use those ciphers.

The `Sasl.STRENGTH` property supports `high`, `medium`, and `low` settings; its default is `high,medium,low`. The ciphers are mapped to the strength settings as follows:

Table 11-2 Cipher Strength

Strength	Cipher	Cipher Id
high	Triple DES RC4 128 bits	3des rc4
medium	DES RC4 56 bits	des rc4-56
low	RC4 40 bits	rc4-40

When there is more than one choice for a particular strength, the cipher selected depends on the availability of the ciphers in the underlying platform. To explicitly name the cipher to use, set the `com.sun.security.sasl.digest.cipher` property to the corresponding cipher id. Note that this property setting must be compatible with `Sasl.STRENGTH` and the ciphers available in the underlying platform. For example, `Sasl.STRENGTH` being set to `low` and `com.sun.security.sasl.digest.cipher` being set to `3des` are incompatible. The `com.sun.security.sasl.digest.cipher` property has no default.

The `javax.security.sasl.sendmaxbuffer`property specifies (the string representation of) the maximum send buffer size in bytes. The default is 65536. The actual maximum number of bytes will be the minimum of this number and the peer's maximum receive buffer size.

NTLM

> **Note:**
>
> This section applies both to the NTLM client mechanism and the NTLM server mechanism.

NT LAN Manager (NTLM) is an security protocol from Microsoft used to access their various services such as IIS Web Server and Exchange Mail Server. As a SASL mechanism, it can be used to access Microsoft Exchange Server. It is also useful for HTTP authentication with the NTLM scheme.

The NTLM mechanism is used for NTLM authentication. It does not provide a security layer. This means that you can only set the `javax.security.sasl.qop` environment property to `auth`.

If the LMCompatibilityLevel registry value is set to a high value on the server, certain low value requests are not supported. However, there's no protocol for the server to inform the client to use a higher version, so the user must manually choose the correct version on the client side.

Set the system property `ntlm.debug` to any value to turn on debugging

Provide the following information either at mechanism creation or through callbacks:

Table 11-3 NTLM Required Information

Information	Type	Required or Optional	Description
Name	`String`	Required	Provided through `NameCallback` with the `authzid` input argument as the default value
Password	`char[]`	Required	Provided through `PasswordCallback` If the password contains non-ASCII characters, the original LM version might fail. In this case, do not choose `LM` as the version.

Table 11-3 (Cont.) NTLM Required Information

Information	Type	Required or Optional	Description
Domain	`String`	Optional	Provided through `RealmCallback` with the `serverName` input argument as the default value. The domain provided on the client side is used to create the Type 1 message. The negotiated property `com.sun.security.sasl.ntlm.domain` is determined by the server's Type 2 message.

Table 11-3 (Cont.) NTLM Required Information

Information	Type	Required or Optional	Description
NTLM version	`String`	Optional	Specifies a specific version to use. Provided through the `com.sun.security.sasl.ntlm.version` property. It can have one of the following values: • `LM/NTLM`: Original NTLM v1 • `LM`: Original NTLM v1, LM only • `NTLM`: Original NTLM v1, NTLM only • `NTLM2`: NTLM v1 with Client Challenge • `LMv2/NTLMv2`: NTLM v2 • `LMv2`: NTLM v2, LM only • `NTLMv2`: NTLM v2, NTLM only If not provided, then the system property `ntlm.version` is used. If still not provided, then the value `LMv2/NTLMv2` is used, and on the server side, all values are accepted. Note: these types are only different on the client side. On the server side, because authentication succeeds if only one of LM (or LMv2) or NTLM (or NTLMv2) is verified, the first three types are effectively the same; this is also true for the last three types.
Host name	`String`	Optional	Provided through the `com.sun.security.sasl.ntlm.hostname` property, which will be sent to the server. If not provided, then the system will automatically derive a host name. This property is only used on the client side.

Table 11-3 (Cont.) NTLM Required Information

Information	Type	Required or Optional	Description
Random source	`java.util.Random`	Optional	Used as random source to derive nonce bytes. Provided through the `com.sun.security.sasl.ntlm.random` property. If not provided, then an internal `java.util.Random` object is used.

After authentication, the client will receive a negotiated property named `com.sun.security.sasl.html.domain`, which is provided by the server, and the server will receive a negotiated property named `com.sun.security.sasl.ntlm.hostname`, which is he host name the client used to access this server.

The SunSASL Provider Server Mechanisms

The SunSASL provider supports several SASL server mechanisms used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the server mechanisms and the required input:

Table 11-4 Server Mechanisms

Server Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy

CRAM-MD5	server name	`AuthorizeCallback`
		`NameCallback`
		`PasswordCallback`
			None	`Sasl.POLICY_NOANONYMOUS`
				`Sasl.POLICY_NOPLAINTEXT`

DIGEST-MD5	protocol id
	server name
		`AuthorizeCallback`
		`NameCallback`
		`PasswordCallback`
		`RealmCallback`
			`Sasl.QOP`
			`Sasl.STRENGTH`
			`Sasl.MAX_BUFFER`
			`javax.security.sasl.sendmaxbuffer`
			`com.sun.security.sasl.digest.realm`
			`com.sun.security.sasl.digest.utf8`
				`Sasl.POLICY_NOANONYMOUS`
				`Sasl.POLICY_NOPLAINTEXT`

Table 11-4 (Cont.) Server Mechanisms

Server Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy

NTLM	serverName (as domain, can be overridden by properties)	`RealmCallback`, providing request user's domain
		`NameCallback`, providing request user's name
		`PasswordCallback`
			`Sasl.QOP`
			`com.sun.security.sasl.ntlm.random`
			`com.sun.security.sasl.ntlm.version`
			`com.sun.security.sasl.ntlm.domain`
				`Sasl.POLICY_NOANONYMOUS`
				`Sasl.POLICY_NOPLAINTEXT`

An application that uses these mechanisms from the SunSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults.

All users of server mechanisms must have a callback handler that deals with the `AuthorizeCallback`. This is used by the mechanisms to determine whether the authenticated user is allowed to act on behalf of the requested authorization id, and also to obtain the canonicalized name of the authorized user (if canonicalization is applicable).

Most of the parameters, callbacks, and properties are described in the API documentation. The following sections describe mechanism-specific behaviors and parameters not already covered by the API documentation.

Cram-MD5

The Cram-MD5 server mechanism uses the `NameCallback` and `PasswordCallback` to obtain the password required to verify the SASL client's response. The callback handler should use the `NameCallback.getDefaultName()` as the key to fetch the password.

Digest-MD5

The Digest-MD5 server mechanism uses the `RealmCallback`, `NameCallback`, and `PasswordCallback` to obtain the password required to verify the SASL client's response. The callback handler should use `RealmCallback.getDefaultText()` and `NameCallback.getDefaultName()` as keys to fetch the password.

The `javax.security.sasl.sendmaxbuffer` property specifies (the string representation of) the maximum send buffer size in bytes. The default is 65536. The actual maximum number of bytes will be the minimum of this number and the peer's maximum receive buffer size.

The `com.sun.security.sasl.digest.realm` property is used to specify a list of space-separated realm names that the server supports. The list is sent to the client as part of the challenge. If this property has not been set, the default realm is the server's name (supplied as a parameter).

The `com.sun.security.sasl.digest.utf8` property is used to specify the character encoding to use. The value `true` means to use UTF-8 encoding; the value `false` means to use ISO Latin 1 (ISO-8859-1). The default value is `true`.

The JdkSASL Provider

The JdkSASL provider supports the following client and server mechanisms:

- Client Mechanisms
  - GSSAPI (RFC 2222). This mechanism uses the GSSAPI for obtaining authentication information. It supports Kerberos v5 authentication.
- Server Mechanisms
  - GSSAPI (Kerberos v5)

The JdkSASL Provider Client Mechanism

The JdkSASL provider supports the GSSAPI client mechanism used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the GSSAPI client mechanism and its required input.

Table 11-5 JdkSASL Provider Client Mechanism

Client Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy

GSSAPI	JAAS Subject
	authorization id
	protocol id
	server name
		None	`Sasl.QOP`
			`Sasl.MAX_BUFFER`
			`Sasl.SERVER_AUTH`
			`javax.security.sasl.sendmaxbuffer`
			`com.sun.security.jgss.inquiretype.`*`type_name`*
				`Sasl.POLICY_NOACTIVE`
				`Sasl.POLICY_NOANONYMOUS`
				`Sasl.POLICY_NOPLAINTEXT`

An application that uses the GSSAPI mechanism from the JdkSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults. Most of the parameters, callbacks, and properties are described in the API documentation. The following section describes further GSSAPI behaviors and parameters not already covered by the API documentation.

GSSAPI

> **Note:**
>
> The GSSAPI server mechanism has the same requirements as the GSSAPI client mechanism in terms of Kerberos credentials and the `javax.security.sasl.sendmaxbuffer` property.

The GSSAPI mechanism is used for Kerberos v5 authentication and optional establishment of a security layer. The mechanism expects the calling thread's `Subject` to contain the client's Kerberos credentials or that the credentials could be obtained by implicitly logging in to Kerberos. To obtain the client's Kerberos credentials, use the Java Authentication and Authorization Service (JAAS) to log in using the Kerberos login module. See Introduction to JAAS and Java GSS-API Tutorials for details and examples. After using JAAS authentication to obtain the Kerberos credentials, you put the code that uses the SASL GSSAPI mechanism within `doAs` or `doAsPrivileged`.

```
LoginContext lc = new LoginContext("JaasSample", new
TextCallbackHandler());
lc.login();
lc.getSubject().doAs(new SaslAction());

class SaslAction implements java.security.PrivilegedAction<Void> {
   public Void run() {
       // ...
       String[] mechanisms = new String[]{"GSSAPI"};
       SaslClient sc = Sasl.createSaslClient(
           mechanisms, authzid, protocol, serverName, props,
callbackHandler);
       // ...
   }
}
```

To obtain Kerberos credentials without doing explicit JAAS programming, see Use of Java GSS-API for Secure Message Exchanges Without JAAS Programming. When using this approach, there is no need to wrap the code within `doAs` or `doAsPrivileged`

The `javax.security.sasl.sendmaxbuffer` property specifies (the string representation of) the maximum send buffer size in bytes. The default is 65536. The actual maximum number of bytes will be the minimum of this number and the peer's maximum receive buffer size.

The `com.sun.security.jgss.inquiretype.`*`type_name`* negotiated property contains the value returned by the `ExtendedGSSContext.inquireSecContext(InquireType)` method, where *`type_name`* is the string form of the `InquireType` enum parameter in lower case.

The JdkSASL Provider Server Mechanism

The JdkSASL provider supports the GSSAPI mechanism used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the GSSAPI server mechanism and the required input:

Table 11-6 Server mechanism

Server Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy

GSSAPI	`Subject`
	protocol id
	server name
		`AuthorizeCallback`	`Sasl.QOP`
			`Sasl.MAX_BUFFER`
			`javax.security.sasl.sendmaxbuffer`
				`Sasl.POLICY_NOACTIVE`
				`Sasl.POLICY_NOANONYMOUS`
				`Sasl.POLICY_NOPLAINTEXT`

An application that uses the GSSAPI mechanism from the JdkSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults.

All users of server mechanism must have a callback handler that deals with the `AuthorizeCallback`. This is used by the mechanism to determine whether the authenticated user is allowed to act on behalf of the requested authorization id, and also to obtain the canonicalized name of the authorized user (if canonicalization is applicable).

Most of the parameters, callbacks, and properties are described in the API documentation.

Debugging and Monitoring

The SunSASL and JdkSASL providers uses the Logging APIs to provide implementation logging output. This output can be controlled by using the logging configuration file and programmatic API (`java.util.logging`). The logger name used by the SunSASL provider is `javax.security.sasl`. Here is a sample logging configuration file that enables the `FINEST` logging level for the SunSASL provider:

```
javax.security.sasl.level=FINEST
handlers=java.util.logging.ConsoleHandler
java.util.logging.ConsoleHandler.level=FINEST
```

The table below shows the mechanisms and the logging output that they generate:

Table 11-7 Logging Output

Mechanism	Logging Level	Information Logged
CRAM-MD5	`FINE`	Configuration properties; challenge/response messages
DIGEST-MD5	`INFO`	Message discarded due to encoding problem (for example, unmatched MACs, incorrect padding)

Table 11-7 (Cont.) Logging Output

Mechanism	Logging Level	Information Logged
DIGEST-MD5	`FINE`	Configuration properties; challenge/response messages
DIGEST-MD5	`FINER`	More detailed information about challenge/response messages
DIGEST-MD5	`FINEST`	Buffers exchanged at the security layer
GSSAPI	`FINE`	Configuration properties; challenge/response messages
GSSAPI	`FINER`	More detailed information about challenge/response messages
GSSAPI	`FINEST`	Buffers exchanged at the security layer

Implementing a SASL Security Provider

There are three basic steps in implementing a SASL security provider:

1. Write a class that implements the `SaslClient` or `SaslServer` interface.

   This involves providing an implementation for the SASL mechanism. To implement a client mechanism, you need to implement the methods declared in the `SaslClient` interface. Similarly, for a server mechanism, you need to implement the methods declared in the `SaslServer` interface. For the purposes of this discussion, suppose you are developing an implementation for the client mechanism "SAMPLE-MECH", implemented by the class, `com.example.SampleMechClient`. You must decide what input are needed by the mechanism and how the implementation is going to collect them. For example, if the mechanism is username/password-based, then the implementation would likely need to collect that information via the callback handler parameter.

2. Write a factory class (that implements `SaslClientFactory` or `SaslServerFactory`) that creates instances of the class.

   This involves providing a factory class that will create instances of `com.example.SampleMechClient`. The factory needs to determine the characteristics of the mechanism that it supports (as described by the `Sasl.POLICY_*` properties) so that it can return an instance of the mechanism when the API user requests it using compatible policy properties. The factory may also check for validity of the parameters before creating the mechanism. For the purposes of this discussion, suppose the factory class is named `com.example.MySampleClientFactory`. Although our sample factory is responsible for only one mechanism, a single factory can be responsible for any number of mechanisms.

3. Write a JCA provider that registers the factory.

   This involves creating a JCA provider. The steps for creating a JCA provider is described in detail in Steps to Implement and Integrate a Provider. SASL client factories are registered using property names of the form `SaslClientFactory.`*mechName* while SASL server factories are registered using property names of the form `SaslServerFactory.`*mechName*

mechName is the SASL mechanism's name. This is what's returned by `SaslClient.getMechanismName()` and `SaslServer.getMechanismName()`. Continuing with our example, here is how the provider would register the "SAMPLE-MECH" mechanism.

```
    put("SaslClientFactory.SAMPLE-MECH",
"com.example.MySampleClientFactory");
```

A single SASL provider might be responsible for many mechanisms. Therefore, it might have many invocations of `put` to register the relevant factories. The completed SASL provider can then be made available to applications using the instructions described in How SASL Mechanisms are Installed and Selected.

12

XML Digital Signature API Overview and Tutorial

The Java XML Digital Signature API is a standard Java API for generating and validating XML Signatures. This API was defined under the Java Community Process as JSR 105.

XML Signatures can be applied to data of any type, XML or binary (see XML Signature Syntax and Processing). The resulting signature is represented in XML. An XML Signature can be used to secure your data and provide data integrity, message authentication, and signer authentication.

After providing a brief overview of XML Signatures and the XML Digital Signature API, this document presents two examples that demonstrate how to use the API to validate and generate an XML Signature. This document assumes that you have a basic knowledge of cryptography and digital signatures.

The API is designed to support all of the required or recommended features of the W3C Recommendation for XML-Signature Syntax and Processing. The API is extensible and pluggable and is based on the Java Cryptography Service Provider Architecture; see Java Cryptography Architecture (JCA) Reference Guide. The API is designed for two types of developers:

- Developers who want to use the XML Digital Signature API to generate and validate XML signatures
- Developers who want to create a concrete implementation of the XML Digital Signature API and register it as a cryptographic service of a JCA provider (see The Provider Class).

Package Hierarchy

The following six packages, which are contained in the `java.xml.crypto` module, comprise the XML Digital Signature API:

- `javax.xml.crypto`
- `javax.xml.crypto.dsig`
- `javax.xml.crypto.dsig.keyinfo`
- `javax.xml.crypto.dsig.spec`
- `javax.xml.crypto.dom`
- `javax.xml.crypto.dsig.dom`

The `javax.xml.crypto` package contains common classes that are used to perform XML cryptographic operations, such as generating an XML signature or encrypting XML data. Two notable classes in this package are the `KeySelector` class, which allows developers to supply implementations that locate and optionally validate keys

using the information contained in a `KeyInfo` object, and the `URIDereferencer` class, which allows developers to create and specify their own URI dereferencing implementations.

The `javax.xml.crypto.dsig` package includes interfaces that represent the core elements defined in the W3C XML digital signature specification. Of primary significance is the `XMLSignature` class, which allows you to sign and validate an XML digital signature. Most of the XML signature structures or elements are represented by a corresponding interface (except for the KeyInfo structures, which are included in their own package and are discussed in the next paragraph). These interfaces include: `SignedInfo`, `CanonicalizationMethod`, `SignatureMethod`, `Reference`, `Transform`, `DigestMethod`, `XMLObject`, `Manifest`, `SignatureProperty`, and `SignatureProperties`. The `XMLSignatureFactory` class is an abstract factory that is used to create objects that implement these interfaces.

The `javax.xml.crypto.dsig.keyinfo` package contains interfaces that represent most of the `KeyInfo` structures defined in the W3C XML digital signature recommendation, including `KeyInfo`, `KeyName`, `KeyValue`, `X509Data`, `X509IssuerSerial`, `RetrievalMethod`, and `PGPData`. The `KeyInfoFactory` class is an abstract factory that is used to create objects that implement these interfaces.

The `javax.xml.crypto.dsig.spec` package contains interfaces and classes representing input parameters for the digest, signature, transform, or canonicalization algorithms used in the processing of XML signatures.

Finally, the `javax.xml.crypto.dom` and `javax.xml.crypto.dsig.dom` packages contains DOM-specific classes for the `javax.xml.crypto` and `javax.xml.crypto.dsig` packages, respectively. Only developers and users who are creating or using a DOM-based `XMLSignatureFactory` or `KeyInfoFactory` implementation will need to make direct use of these packages.

Service Providers

A Java XML Signature is a concrete implementation of the abstract `XMLSignatureFactory` and `KeyInfoFactory` classes and is responsible for creating objects and algorithms that parse, generate and validate XML Signatures and `KeyInfo` structures. A concrete implementation of `XMLSignatureFactory` must provide support for each of the required algorithms as specified by the W3C recommendation for XML Signatures. It can optionally support other algorithms as defined by the W3C recommendation or other specifications.

The Java XML Digital Signature API leverages the JCA provider model for registering and loading `XMLSignatureFactory` and `KeyInfoFactory` implementations.

Each concrete `XMLSignatureFactory` or `KeyInfoFactory` implementation supports a specific XML mechanism type that identifies the XML processing mechanism that an implementation uses internally to parse and generate XML signature and `KeyInfo` structures. This JSR supports one standard type, DOM. The XML Digital Signature provider implementation that is bundled with the JDK supports the DOM mechanism.

An XML Digital Signature API implementation should use underlying JCA engine classes, such as `java.security.Signature` and `java.security.MessageDigest`, to perform cryptographic operations.

In addition to the `XMLSignatureFactory` and `KeyInfoFactory` classes, JSR 105 supports a service provider interface for transform and canonicalization algorithms. The `TransformService` class allows you to develop and plug in an implementation of a specific transform or canonicalization algorithm for a particular XML mechanism type. The `TransformService` class uses the standard JCA provider model for registering and loading implementations. Each JSR 105 implementation should use the `TransformService` class to find a provider that supports transform and canonicalization algorithms in XML Signatures that it is generating or validating.

Introduction to XML Signatures

You can use an XML Signature to sign any arbitrary data, whether it is XML or binary. The data is identified via URIs in one or more Reference elements. XML Signatures are described in one or more of three forms: detached, enveloping, or enveloped. A detached signature is over data that is external, or outside of the signature element itself. Enveloping signatures are signatures over data that is inside the signature element, and an enveloped signature is a signature that is contained inside the data that it is signing.

Example of an XML Signature

The easiest way to describe the contents of an XML Signature is to show an actual sample and describe each component in more detail. Example 12-3 is an enveloped XML Signature generated over the contents of an XML document. The root element, `Envelop`, contains a `Signature` element:

```
<Envelope xmlns="urn:envelope">
  <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
    <!-- ... -->
  </Signature>
</Envelope>
```

This `Signature` element has been inserted inside the content that it is signing, thereby making it an enveloped signature. The required `SignedInfo` element contains the information that is actually signed:

```
<Envelope xmlns="urn:envelope">
  <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
    <SignedInfo>
      <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xml-
c14n-20010315#WithComments"/>
      <SignatureMethod Algorithm="http://www.w3.org/2001/04/xmldsig-
more#rsa-sha256"/>
      <Reference URI="">
        <Transforms>
          <Transform Algorithm="http://www.w3.org/2000/09/
xmldsig#enveloped-signature"/>
        </Transforms>
```

```
        <DigestMethod Algorithm="http://www.w3.org/2001/04/xmlenc#sha256"/>
        <DigestValue>/juoQ4bDxElf1M+KJauO20euW+QAvvPP0nDCruCQooM=</
DigestValue>
      </Reference>
    </SignedInfo>
    <!-- ... -->
  </Signature>
</Envelope>
```

The required `CanonicalizationMethod` element defines the algorithm used to canonicalize the `SignedInfo` element before it is signed or validated. Canonicalization is the process of converting XML content to a canonical form, to take into account changes that can invalidate a signature over that data. Canonicalization is necessary due to the nature of XML and the way it is parsed by different processors and intermediaries, which can change the data such that the signature is no longer valid but the signed data is still logically equivalent.

The required `SignatureMethod` element defines the digital signature algorithm used to generate the signature, in this case RSA with SHA-256.

One or more `Reference` elements identify the data that is digested. Each `Reference` element identifies the data via a URI. In this example, the value of the URI is the empty String ("""), which indicates the root of the document. The optional `Transforms` element contains a list of one or more `Transform` elements, each of which describes a transformation algorithm used to transform the data before it is digested. In this example, there is one Transform element for the enveloped transform algorithm. The enveloped transform is required for enveloped signatures so that the signature element itself is removed before calculating the signature value. The required `DigestMethod` element defines the algorithm used to digest the data, in this case SHA-256. Finally the required `DigestValue` element contains the actual base64-encoded digested value.

The required `SignatureValue` element contains the base64-encoded signature value of the signature over the `SignedInfo` element.

The optional `KeyInfo` element contains information about the key that is needed to validate the signature:

```
    <KeyInfo>
      <KeyValue>
        <RSAKeyValue>
          <Modulus>
9hSmAKw/4TTw/1l1u1pYzdFm6lOjRB/5NfdGWl/fB8iAa/tiK0f1u/
VWoK6SMtogYgSDKqQThbAu
9dy9rRnOWRGY2He1JtpOvGh0WCmIFUEs2P22HvEf+JGKVEpkoP4hv53ucT69T+7nKGK3/bjxgp
+T
C7fbnVj651+jAHuDFlC8Txt1R8ZymfN5cUeHIH96dvNFrtai/uwZDbVMfhV9chL//
+Vyhx4O5nHv
jfS
+0So9Qi52YAbEyLu6+BLdu8wnMWapC88CfXsRwrpx8b6aCU0e6QSZyOvdgXWz3+9ifVTBDIxE
kjhL5OASx0qjvc+dPUOMvq7fJE05RRZLyb0YJw==
          </Modulus>
          <Exponent>AQAB</Exponent>
        </RSAKeyValue>
```

```
    </KeyValue>
  </KeyInfo>
```

This `KeyInfo` element contains a `KeyValue` element, which in turn contains a `RSAKeyValue` element consisting of the public key needed to validate the signature. `KeyInfo` can contain various content such as X.509 certificates and PGP key identifiers. See The `KeyInfo` Element in XML Signature Syntax and Processing for more information on the different `KeyInfo` types.

XML Signature Secure Validation Mode

XML Signature secure validation mode can protect you from XML Signatures that may contain potentially hostile constructs that can cause denial-of-service or other types of security issues.

XML Signature secure validation mode is enabled by default when you run your application with a security manager.

You can also enable XML Signature secure validation mode by setting the `org.jcp.xml.dsig.secureValidation` property to `TRUE`. You must set this property to `TRUE` before validating an XML Signature.

To set this property in an application, call the `javax.xml.crypto.dsig.dom.DOMValidateContext.setProperty` method:

```
DOMValidateContext context = new DOMValidateContext(key, element);
context.setProperty("org.jcp.xml.dsig.secureValidation", Boolean.TRUE);
```

When XML Signature secure validation mode is enabled, XML Signatures are processed more securely. Limits are set on various XML Signature constructs to avoid conditions such as denial-of-service attacks. By default, it enforces the following restrictions:

- Forbids the use of XSLT transforms
- Restricts the number of `SignedInfo` or `Manifest` `Reference` elements to 30 or less
- Restricts the number of `Reference` transforms to 5 or less
- Forbids the use of MD5-related signatures or MAC algorithms
- Ensures that `Reference` IDs are unique to help prevent signature wrapping attacks
- Forbids `Reference` URIs of type `http`, `https`, or `file`
- Does not allow a `RetrievalMethod` element to reference another `RetrievalMethod` element
- Forbids RSA or DSA keys less than 1024 bits

In addition, you can use the `jdk.xml.dsig.secureValidationPolicy` Security Property to control and fine-tune the restrictions listed previously or add additional restrictions. See the definition of this Security Property in the `java.security` file for more information.

XML Digital Signature API Examples

The following sections describe two examples that show how to use the XML Digital Signature API:

Validate Example

To compile and run the example, execute the following commands:

```
$ javac Validate.java
$ java Validate signature.xml
```

The sample program will validate the signature in the file `signature.xml` in the current working directory.

Example 12-1 Validate.java

```
import javax.xml.crypto.*;
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dom.*;
import javax.xml.crypto.dsig.dom.DOMValidateContext;
import javax.xml.crypto.dsig.keyinfo.*;
import java.io.FileInputStream;
import java.security.*;
import java.util.Collections;
import java.util.Iterator;
import java.util.List;
import javax.xml.parsers.DocumentBuilderFactory;
import org.w3c.dom.Document;
import org.w3c.dom.NodeList;

/**
 * This is a simple example of validating an XML Signature using
 * the XML Signature API. It assumes the key needed to validate
 * the signature is contained in a KeyValue KeyInfo.
 */
public class Validate {

    //
    // Synopsis: java Validate [document]
    //
    //    where "document" is the name of a file containing the XML
document
    //    to be validated.
    //
    public static void main(String[] args) throws Exception {

        // Instantiate the document to be validated
        DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
        dbf.setNamespaceAware(true);
        Document doc = null;
        try (FileInputStream fis = new FileInputStream(args[0])) {
            doc = dbf.newDocumentBuilder().parse(fis);
        }
```

```
        // Find Signature element
        NodeList nl =
            doc.getElementsByTagNameNS(XMLSignature.XMLNS, "Signature");
        if (nl.getLength() == 0) {
            throw new Exception("Cannot find Signature element");
        }

        // Create a DOM XMLSignatureFactory that will be used to unmarshal
the
        // document containing the XMLSignature
        XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");

        // Create a DOMValidateContext and specify a KeyValue KeySelector
        // and document context
        DOMValidateContext valContext = new DOMValidateContext
            (new KeyValueKeySelector(), nl.item(0));

        // unmarshal the XMLSignature
        XMLSignature signature = fac.unmarshalXMLSignature(valContext);

        // Validate the XMLSignature (generated above)
        boolean coreValidity = signature.validate(valContext);

        // Check core validation status
        if (coreValidity == false) {
            System.err.println("Signature failed core validation");
            boolean sv =
signature.getSignatureValue().validate(valContext);
            System.out.println("signature validation status: " + sv);
            // check the validation status of each Reference
            Iterator<Reference> i =
                signature.getSignedInfo().getReferences().iterator();
            for (int j=0; i.hasNext(); j++) {
                boolean refValid = i.next().validate(valContext);
                System.out.println("ref["+j+"] validity status: " +
refValid);
            }
        } else {
            System.out.println("Signature passed core validation");
        }
    }

    /**
     * KeySelector which retrieves the public key out of the
     * KeyValue element and returns it.
     * NOTE: If the key algorithm doesn't match signature algorithm,
     * then the public key will be ignored.
     */
    private static class KeyValueKeySelector extends KeySelector {
        public KeySelectorResult select(KeyInfo keyInfo,
                                        KeySelector.Purpose purpose,
                                        AlgorithmMethod method,
                                        XMLCryptoContext context)
            throws KeySelectorException {
            if (keyInfo == null) {
```

```
                throw new KeySelectorException("Null KeyInfo object!");
            }
            SignatureMethod sm = (SignatureMethod) method;
            List<XMLStructure> list = keyInfo.getContent();

            for (int i = 0; i < list.size(); i++) {
                XMLStructure xmlStructure = list.get(i);
                if (xmlStructure instanceof KeyValue) {
                    PublicKey pk = null;
                    try {
                        pk = ((KeyValue)xmlStructure).getPublicKey();
                    } catch (KeyException ke) {
                        throw new KeySelectorException(ke);
                    }
                    // make sure algorithm is compatible with method
                    if (algEquals(sm.getAlgorithm(), pk.getAlgorithm())) {
                        return new SimpleKeySelectorResult(pk);
                    }
                }
            }
            throw new KeySelectorException("No KeyValue element found!");
        }

        static boolean algEquals(String algURI, String algName) {
            if (algName.equalsIgnoreCase("DSA") &&
                algURI.equalsIgnoreCase("http://www.w3.org/2009/
xmldsig11#dsa-sha256")) {
                return true;
            } else if (algName.equalsIgnoreCase("RSA") &&
                       algURI.equalsIgnoreCase("http://www.w3.org/2001/04/
xmldsig-more#rsa-sha256")) {
                return true;
            } else {
                return false;
            }
        }
    }

    private static class SimpleKeySelectorResult implements
KeySelectorResult {
        private PublicKey pk;
        SimpleKeySelectorResult(PublicKey pk) {
            this.pk = pk;
        }

        public Key getKey() { return pk; }
    }
}
```

Example 12-2 envelope.xml

```
<Envelope xmlns="urn:envelope">
</Envelope>
```

Example 12-3 signature.xml

This file has been indented and formatted for readability.


```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<Envelope xmlns="urn:envelope">
  <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
    <SignedInfo>
      <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xml-
c14n-20010315#WithComments"/>
      <SignatureMethod Algorithm="http://www.w3.org/2001/04/xmldsig-
more#rsa-sha256"/>
      <Reference URI="">
        <Transforms>
          <Transform Algorithm="http://www.w3.org/2000/09/
xmldsig#enveloped-signature"/>
        </Transforms>
        <DigestMethod Algorithm="http://www.w3.org/2001/04/xmlenc#sha256"/>
        <DigestValue>/juoQ4bDxElf1M+KJauO20euW+QAvvPP0nDCruCQooM=</
DigestValue>
      </Reference>
    </SignedInfo>
    <SignatureValue>
Vorr4nABCD7eWOjh4jn8pdM5iseGJPt4BmlgjEbxr05TsR9ObHq7WLVOBOtJfb3M6pXv6NnTucp
H
e/97zHbuUMaNeGxCs/gN7YDUGOkQE1Gs4HAbGwXuTcif3pw
+066ZW4uxyzapwS6lZHmqIm7PRl8I
NIQXVL4dezLe+rx77Kh+rZRheVe4UlTTP
+TmIOaBZo93GQ5FudreMhSiuIC0Nx2SP7mAkt6+8kVH
luZouFbqriSvyhzIxDgyOXpm/PHCuuPU2scCokwjEZBtlZXDOl6lIWGllnyrptWntQ6F/
ngQObI5
c2+npgCshq1svGuS/xx18MAFHGWi98Vj+07QCg==
    </SignatureValue>
    <KeyInfo>
      <KeyValue>
        <RSAKeyValue>
          <Modulus>
9hSmAKw/4TTw/1l1u1pYzdFm6lOjRB/5NfdGWl/fB8iAa/tiK0f1u/
VWoK6SMtogYgSDKqQThbAu
9dy9rRnOWRGY2He1JtpOvGh0WCmIFUEs2P22HvEf+JGKVEpkoP4hv53ucT69T+7nKGK3/bjxgp
+T
C7fbnVj651+jAHuDFlC8Txt1R8ZymfN5cUeHIH96dvNFrtai/uwZDbVMfhV9chL//
+Vyhx4O5nHv
jfS
+0So9Qi52YAbEyLu6+BLdu8wnMWapC88CfXsRwrpx8b6aCU0e6QSZyOvdgXWz3+9ifVTBDIxE
kjhL5OASx0qjvc+dPUOMvq7fJE05RRZLyb0YJw==
          </Modulus>
          <Exponent>AQAB</Exponent>
        </RSAKeyValue>
      </KeyValue>
    </KeyInfo>
  </Signature>
</Envelope>
```

Validating an XML Signature

This example shows you how to validate an XML Signature using the Java XML Digital Signature API. The example uses DOM (the Document Object Model) to parse an XML document containing a Signature element and a DOM implementation to validate the signature.

Instantiating the Document that Contains the Signature

First we use a JAXP `DocumentBuilderFactory` to parse the XML document containing the `Signature`. An application obtains the default implementation for `DocumentBuilderFactory` by calling the following line of code:

```
DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
```

We must also make the factory namespace-aware:

```
dbf.setNamespaceAware(true);
```

Next, we use the factory to get an instance of a `DocumentBuilder`, which is used to parse the document:

```
Document doc = null;
try (FileInputStream fis = new FileInputStream(args[0])) {
    doc = dbf.newDocumentBuilder().parse(fis);
}
```

Specifying the Signature Element to be Validated

We need to specify the `Signature` element that we want to validate, since there could be more than one in the document. We use the DOM method `Document.getElementsByTagNameNS`, passing it the XML Signature namespace URI and the tag name of the `Signature` element, as shown:

```
NodeList nl =
    doc.getElementsByTagNameNS(XMLSignature.XMLNS, "Signature");
if (nl.getLength() == 0) {
    throw new Exception("Cannot find Signature element");
}
```

This returns a list of all `Signature` elements in the document. In this example, there is only one `Signature` element.

Creating a Validation Context

We create an `XMLValidateContext` instance containing input parameters for validating the signature. Since we are using DOM, we instantiate a `DOMValidateContext` instance (a subclass of `XMLValidateContext`), and pass it two parameters, a `KeyValueKeySelector` object and a reference to the Signature element to be validated (which is the first entry of the `NodeList` we generated earlier):

```
DOMValidateContext valContext = new DOMValidateContext
    (new KeyValueKeySelector(), nl.item(0));
```

The `KeyValueKeySelector` is explained in greater detail in Using KeySelectors.

Unmarshalling the XML Signature

We extract the contents of the `Signature` element into an `XMLSignature` object. This process is called unmarshalling. The `Signature` element is unmarshalled using an `XMLSignatureFactory` object. An application can obtain a DOM implementation of `XMLSignatureFactory` by calling the following line of code:

```
XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
```

We then invoke the `unmarshalXMLSignature` method of the factory to unmarshal an `XMLSignature` object, and pass it the validation context we created earlier:

```
XMLSignature signature = fac.unmarshalXMLSignature(valContext);
```

Validating the XML Signature

Now we are ready to validate the signature. We do this by invoking the validate method on the `XMLSignature` object, and pass it the validation context as follows:

```
boolean coreValidity = signature.validate(valContext);
```

The validate method returns "true" if the signature validates successfully according to the core validation rules in the W3C XML Signature Recommendation, and false otherwise.

What If the XML Signature Fails to Validate?

If the `XMLSignature.validate` method returns false, we can try to narrow down the cause of the failure. There are two phases in core XML Signature validation:

- Signature validation (the cryptographic verification of the signature)
- Reference validation (the verification of the digest of each reference in the signature)

Each phase must be successful for the signature to be valid. To check if the signature failed to cryptographically validate, we can check the status, as follows:

```
boolean sv = signature.getSignatureValue().validate(valContext);
System.out.println("signature validation status: " + sv);
```

We can also iterate over the references and check the validation status of each one, as follows:

```
Iterator<Reference> i =
    signature.getSignedInfo().getReferences().iterator();
for (int j=0; i.hasNext(); j++) {
    boolean refValid = i.next().validate(valContext);
    System.out.println("ref["+j+"] validity status: " + refValid);
}
```

Using KeySelectors

`KeySelectors` are used to find and select keys that are needed to validate an `XMLSignature`. Earlier, when we created a `DOMValidateContext` object, we passed a `KeyValueKeySelector` object as the first argument:

```
DOMValidateContext valContext = new DOMValidateContext
    (new KeyValueKeySelector(), nl.item(0));
```

Alternatively, we could have passed a `PublicKey` as the first argument if we already knew what key is needed to validate the signature. However, we often don't know.

The `KeyValueKeySelector` class is a concrete implementation of the abstract `KeySelector` class. The `KeyValueKeySelector` implementation tries to find an appropriate validation key using the data contained in `KeyValue` elements of the `KeyInfo` element of an `XMLSignature`. It does not determine if the key is trusted. This is a very simple `KeySelector` implementation, designed for illustration rather than real-world usage. A more practical example of a `KeySelector` is one that searches a `KeyStore` for trusted keys that match `X509Data` information (for example, `X509SubjectName`, `X509IssuerSerial`, `X509SKI`, or `X509Certificate` elements) contained in a `KeyInfo`.

The implementation of the `KeyValueKeySelector` class is as follows:

```
private static class KeyValueKeySelector extends KeySelector {
    public KeySelectorResult select(KeyInfo keyInfo,
                                    KeySelector.Purpose purpose,
                                    AlgorithmMethod method,
                                    XMLCryptoContext context)
        throws KeySelectorException {
        if (keyInfo == null) {
            throw new KeySelectorException("Null KeyInfo object!");
        }
        SignatureMethod sm = (SignatureMethod) method;
        List<XMLStructure> list = keyInfo.getContent();
```

```
            for (int i = 0; i < list.size(); i++) {
                XMLStructure xmlStructure = list.get(i);
                if (xmlStructure instanceof KeyValue) {
                    PublicKey pk = null;
                    try {
                        pk = ((KeyValue)xmlStructure).getPublicKey();
                    } catch (KeyException ke) {
                        throw new KeySelectorException(ke);
                    }
                    // make sure algorithm is compatible with method
                    if (algEquals(sm.getAlgorithm(), pk.getAlgorithm())) {
                        return new SimpleKeySelectorResult(pk);
                    }
                }
            }
            throw new KeySelectorException("No KeyValue element found!");
        }

        static boolean algEquals(String algURI, String algName) {
            if (algName.equalsIgnoreCase("DSA") &&
                algURI.equalsIgnoreCase("http://www.w3.org/2009/
xmldsig11#dsa-sha256")) {
                return true;
            } else if (algName.equalsIgnoreCase("RSA") &&
                       algURI.equalsIgnoreCase("http://www.w3.org/2001/04/
xmldsig-more#rsa-sha256")) {
                return true;
            } else {
                return false;
            }
        }
    }

    private static class SimpleKeySelectorResult implements
KeySelectorResult {
        private PublicKey pk;
        SimpleKeySelectorResult(PublicKey pk) {
            this.pk = pk;
        }

        public Key getKey() { return pk; }
    }
```

GenEnveloped Example

To compile and run this sample, execute the following command:

```
$ javac GenEnveloped.java
$ java GenEnveloped envelope.xml envelopedSignature.xml
```

The sample program will generate an enveloped signature of the document in the file `envelope.xml` and store it in the file `envelopedSignature.xml` in the current working directory.

Example 12-4 GenEnveloped.java

```
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dsig.dom.DOMSignContext;
import javax.xml.crypto.dsig.keyinfo.*;
import javax.xml.crypto.dsig.spec.*;
import java.io.FileInputStream;
import java.io.FileOutputStream;
import java.io.OutputStream;
import java.security.*;
import java.util.List;
import javax.xml.parsers.DocumentBuilderFactory;
import javax.xml.transform.*;
import javax.xml.transform.dom.DOMSource;
import javax.xml.transform.stream.StreamResult;
import org.w3c.dom.Document;

/**
 * This is a simple example of generating an Enveloped XML
 * Signature using the Java XML Digital Signature API. The
 * resulting signature will look like (key and signature
 * values will be different):
 *
 * <pre><code>
 *<Envelope xmlns="urn:envelope">
 * <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
 *   <SignedInfo>
 *     <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-
xml-c14n-20010315"/>
 *     <SignatureMethod Algorithm="http://www.w3.org/2001/04/xmldsig-
more#rsa-sha256"/>
 *     <Reference URI="">
 *       <Transforms>
 *         <Transform Algorithm="http://www.w3.org/2000/09/
xmldsig#enveloped-signature"/>
 *       </Transforms>
 *       <DigestMethod Algorithm="http://www.w3.org/2001/04/
xmlenc#sha256"/>
 *       <DigestValue>/juoQ4bDxElf1M+KJauO20euW
+QAvvPP0nDCruCQooM=<DigestValue>
 *     </Reference>
 *   </SignedInfo>
 *   <SignatureValue>
 *     YeS+F0uiYv0h946M69Q9pKFNnD6dxUwLA8QT3GX/0H3cSPKRnNFyZiR4RPgaA1ir/
ztb4rt6Lqb8
 *     hgwPERIa5qhoGUJyHDfUTcQ0Xqn1jYCVoC3ho
+oUgJPXNVgtMAtpvOgxcWXUPATYdyimO6RrHF8+
 *     JXDkeICI9BPA4NKN1i77CAy6JJbaA87aNIpMJPImwJf8CM7mYsXremZz
+RsafNE2cXXRzAoNOynC
 *     pi4oPYpE7CBLzhd23gf7zYRoyT06/
bVIj4j3qOlVY1TQofsQ20NtAz6PbqAs7QkNoDzkX1CYlDSJ
```

```
 *     U8cGHuwXpul/UIpOiL6MZF8I/YI4ZlJn+O8Mvg==
 *   </SignatureValue>
 *   <KeyInfo>
 *     <KeyValue>
 *       <RSAKeyValue>
 *         <Modulus>
 *           mH0S/iw2K2tFTFHI75BtB67pzjR52HvQ8K7Xi5UX3NJm0oA
+KX2mm0IrVcUuv609vbAAyQoW7CWm
 *           4kswVgStCm68dlw36309cxrEmPhG
+PKBmUaGuBmRzwityjXRyRZJ6yaLenE8SJO/DC5ntQvmHqQQ
 *
qeOJYvz2Cbi2bi6x9XwmpqOfZCE5iTvYwioEsrglhP1uLG9fiXyNR2PXUTyLqD91HLhZFj1CEiU
7
 *           aE++WfkKaowIx5p8e3F6hQ+VFRNXjtemK5aajuL0gwU
+Oujg9ijgbyMh19vBoI8LruJoMOBrYFNN
 *           2boQJ3wP0Ek7CPIqAzQB5MnmvKc9jICKiiZVZw==
 *         </Modulus>
 *         <Exponent>AQAB</Exponent>
 *       </RSAKeyValue>
 *     </KeyValue>
 *   </KeyInfo>
 * </Signature>
 *</Envelope>
 * </code></pre>
 */
public class GenEnveloped {

    //
    // Synopsis: java GenEnveloped [document] [output]
    //
    //    where "document" is the name of a file containing the XML
document
    //    to be signed, and "output" is the name of the file to store the
    //    signed document. The 2nd argument is optional - if not specified,
    //    standard output will be used.
    //
    public static void main(String[] args) throws Exception {

        // Instantiate the document to be signed
        DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
        dbf.setNamespaceAware(true);
        Document doc = null;
        try (FileInputStream fis = new FileInputStream(args[0])) {
            doc = dbf.newDocumentBuilder().parse(fis);
        }

        // Create a RSA KeyPair
        KeyPairGenerator kpg = KeyPairGenerator.getInstance("RSA");
        kpg.initialize(2048);
        KeyPair kp = kpg.generateKeyPair();

        // Create a DOMSignContext and specify the RSA PrivateKey and
        // location of the resulting XMLSignature's parent element
        DOMSignContext dsc = new DOMSignContext
            (kp.getPrivate(), doc.getDocumentElement());
```

```
        // Create a DOM XMLSignatureFactory that will be used to generate
the
        // enveloped signature
        XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");

        // Create a Reference to the enveloped document (in this case we
are
        // signing the whole document, so a URI of "" signifies that) and
        // also specify the SHA256 digest algorithm and the ENVELOPED
Transform.
        Reference ref = fac.newReference
            ("", fac.newDigestMethod(DigestMethod.SHA256, null),
             List.of
              (fac.newTransform
                (Transform.ENVELOPED, (TransformParameterSpec) null)),
             null, null);

        // Create the SignedInfo
        SignedInfo si = fac.newSignedInfo
            (fac.newCanonicalizationMethod
             (CanonicalizationMethod.INCLUSIVE_WITH_COMMENTS,
              (C14NMethodParameterSpec) null),
             fac.newSignatureMethod("http://www.w3.org/2001/04/xmldsig-
more#rsa-sha256", null),
             List.of(ref));

        // Create a KeyValue containing the RSA PublicKey that was
generated
        KeyInfoFactory kif = fac.getKeyInfoFactory();
        KeyValue kv = kif.newKeyValue(kp.getPublic());

        // Create a KeyInfo and add the KeyValue to it
        KeyInfo ki = kif.newKeyInfo(List.of(kv));

        // Create the XMLSignature (but don't sign it yet)
        XMLSignature signature = fac.newXMLSignature(si, ki);

        // Marshal, generate (and sign) the enveloped signature
        signature.sign(dsc);

        // output the resulting document
        OutputStream os;
        if (args.length > 1) {
           os = new FileOutputStream(args[1]);
        } else {
           os = System.out;
        }

        TransformerFactory tf = TransformerFactory.newInstance();
        Transformer trans = tf.newTransformer();
        trans.transform(new DOMSource(doc), new StreamResult(os));
    }
}
```

Example 12-5 envelope.xml

```
<Envelope xmlns="urn:envelope">
</Envelope>
```

Generating an XML Signature

This example shows you how to generate an XML Signature using the XML Digital Signature API. More specifically, the example generates an enveloped XML Signature of an XML document. An enveloped signature is a signature that is contained inside the content that it is signing. The example uses DOM (the Document Object Model) to parse the XML document to be signed and a DOM implementation to generate the resulting signature.

A basic knowledge of XML Signatures and their different components is helpful for understanding this section. See XML Signature Syntax and Processing Version 1.1 for more information.

Instantiating the Document to be Signed

First, we use a JAXP `DocumentBuilderFactory` to parse the XML document that we want to sign. An application obtains the default implementation for `DocumentBuilderFactory` by calling the following line of code:

```
DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
```

We must also make the factory namespace-aware:

```
dbf.setNamespaceAware(true);
```

Next, we use the factory to get an instance of a `DocumentBuilder`, which is used to parse the document:

```
Document doc = null;
try (FileInputStream fis = new FileInputStream(args[0])) {
    doc = dbf.newDocumentBuilder().parse(fis);
}
```

Creating a Public Key Pair

We generate a public key pair. Later in the example, we will use the private key to generate the signature. We create the key pair with a `KeyPairGenerator`. In this example, we will create a RSA `KeyPair` with a length of 2048 bytes:

```
KeyPairGenerator kpg = KeyPairGenerator.getInstance("RSA");
kpg.initialize(2048);
KeyPair kp = kpg.generateKeyPair();
```

In practice, the private key is usually previously generated and stored in a `KeyStore` file with an associated public key certificate.

Creating a Signing Context

We create an `XMLSignContext` containing input parameters for generating the signature. Since we are using DOM, we instantiate a `DOMSignContext` (a subclass of `XMLSignContext`), and pass it two parameters, the private key that will be used to sign the document and the root of the document to be signed:

```
DOMSignContext dsc = new DOMSignContext
    (kp.getPrivate(), doc.getDocumentElement());
```

Assembling the XML Signature

We assemble the different parts of the Signature element into an `XMLSignature` object. These objects are all created and assembled using an `XMLSignatureFactory` object. An application obtains a DOM implementation of `XMLSignatureFactory` by calling the following line of code:

```
XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
```

We then invoke various factory methods to create the different parts of the `XMLSignature` object as shown below. We create a `Reference` object, passing to it the following:

- The URI of the object to be signed (We specify a URI of "", which implies the root of the document.)
- The `DigestMethod` (we use SHA256)
- A single `Transform`, the enveloped `Transform`, which is required for enveloped signatures so that the signature itself is removed before calculating the signature value

```
Reference ref = fac.newReference
    ("", fac.newDigestMethod(DigestMethod.SHA256, null),
     List.of
      (fac.newTransform
        (Transform.ENVELOPED, (TransformParameterSpec) null)),
     null, null);
```

Next, we create the `SignedInfo` object, which is the object that is actually signed, as shown below. When creating the `SignedInfo`, we pass as parameters:

- The `CanonicalizationMethod` (we use inclusive and preserve comments)
- The `SignatureMethod` (we use RSA)
- A list of `References` (in this case, only one)

```
SignedInfo si = fac.newSignedInfo
    (fac.newCanonicalizationMethod
```

```
        (CanonicalizationMethod.INCLUSIVE_WITH_COMMENTS,
         (C14NMethodParameterSpec) null),
        fac.newSignatureMethod("http://www.w3.org/2001/04/xmldsig-
more#rsa-sha256", null),
        List.of(ref));
```

Next, we create the optional `KeyInfo` object, which contains information that enables the recipient to find the key needed to validate the signature. In this example, we add a `KeyValue` object containing the public key. To create `KeyInfo` and its various subtypes, we use a `KeyInfoFactory` object, which can be obtained by invoking the `getKeyInfoFactory` method of the *XMLSignatureFactory*, as follows:

```
KeyInfoFactory kif = fac.getKeyInfoFactory();
```

We then use the `KeyInfoFactory` to create the `KeyValue` object and add it to a `KeyInfo` object:

```
KeyValue kv = kif.newKeyValue(kp.getPublic());
KeyInfo ki = kif.newKeyInfo(List.of(kv));
```

Finally, we create the `XMLSignature` object, passing as parameters the `SignedInfo` and `KeyInfo` objects that we created earlier:

```
XMLSignature signature = fac.newXMLSignature(si, ki);
```

Notice that we haven't actually generated the signature yet; we'll do that in the next step.

Generating the XML Signature

Now we are ready to generate the signature, which we do by invoking the sign method on the `XMLSignature` object, and pass it the signing context as follows:

```
signature.sign(dsc);
```

The resulting document now contains a signature, which has been inserted as the last child element of the root element.

Printing or Displaying the Resulting Document

You can use the following code to print the resulting signed document to a file or standard output:

```
OutputStream os;
if (args.length > 1) {
   os = new FileOutputStream(args[1]);
} else {
   os = System.out;
}
```

```
TransformerFactory tf = TransformerFactory.newInstance();
Transformer trans = tf.newTransformer();
trans.transform(new DOMSource(doc), new StreamResult(os));
```

13
Security API Specification

General Security

- `java.security`
- `javax.crypto`
- `javax.security.cert`
- `javax.crypto.spec`
- `java.security.spec`
- `java.security.interfaces`
- `javax.crypto.interfaces`
- `javax.rmi.ssl`
- `jdk.security.jarsigner`

Certification Path

- `java.security.cert`

JAAS

- `javax.security.auth`
- `javax.security.auth.callback`
- `javax.security.auth.kerberos`
- `javax.security.auth.login`
- `javax.security.auth.spi`
- `javax.security.auth.x500`

Java GSS-API

- `org.ietf.jgss`
- `com.sun.security.jgss`

JSSE

- `javax.net`
- `javax.net.ssl`
- `java.security.cert`

Java SASL

- `javax.security.sasl`

SSL/TLS-Based RMI Socket Factories

- `javax.rmi.ssl`

XML Digital Signature

- `javax.xml.crypto`
- `javax.xml.crypto.dom`
- `javax.xml.crypto.dsig`
- `javax.xml.crypto.dsig.dom`
- `javax.xml.crypto.dsig.keyinfo`
- `javax.xml.crypto.dsig.spec`

Smart Card I/O

- `javax.smartcardio`

14

Deprecated Security APIs Marked for Removal

The following APIs are deprecated and eligible to be removed in a future release.

You can check the API dependencies using the `jdeprscan` tool. See jdeprscan in *Java Platform, Standard Edition Tools Reference*.

The following classes are deprecated and marked for removal:

- `com.sun.security.auth.PolicyFile`
- `com.sun.security.auth.SolarisNumericGroupPrincipal`
- `com.sun.security.auth.SolarisNumericUserPrincipal`
- `com.sun.security.auth.SolarisPrincipal`
- `com.sun.security.auth.X500Principal`
- `com.sun.security.auth.module.SolarisLoginModule`
- `com.sun.security.auth.module.SolarisSystem`

The following methods are deprecated and marked for removal:

- `java.lang.SecurityManager.getInCheck`
- `java.lang.SecurityManager.checkMemberAccess`
- `java.lang.SecurityManager.classDepth`
- `java.lang.SecurityManager.currentClassLoader`
- `java.lang.SecurityManager.currentLoadedClass`
- `java.lang.SecurityManager.inClass`
- `java.lang.SecurityManager.inClassLoader`
- `java.lang.SecurityManager.checkAwtEventQueueAccess`
- `java.lang.SecurityManager.checkTopLevelWindow`
- `java.lang.SecurityManager.checkSystemClipboardAccess`

The following field is deprecated and marked for removal:

- `java.lang.SecurityManager.incheck`

15
Security Tools

See Security Tools and Commands in *Java Platform, Standard Edition Tools Reference*.

16
Related Security Topics

- The Security Features in Java SE trail of the Java Tutorial
- Serialization Filtering in *Java Platform, Standard Edition Java Core Libraries Developer's Guide*
- RMI:
  - RMI Security Recommendations in *Java Platform, Standard Edition Java Remote Method Invocation User's Guide*
  - Using Custom Socket Factories with Java RMI in the JDK 8 documentation
- JAXP:
  - JAXP Processing Limits in the Java Tutorials
  - External Access Restriction Properties in the Java Tutorials