

Oracle® Communication Services Gatekeeper

Platform Development Studio - Developer's Guide

Release 4.0

June 2008

ORACLE®

Copyright © 2007, 2008, 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 software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software 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 which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on 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. 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.

Contents

Document Roadmap

Document Scope and Audience	1-1
Guide to This Document	1-1
Terminology	1-2
Related Documentation	1-5

Overview of the Platform Development Studio

Creating New Communication Services	2-1
The Eclipse Wizard	2-2
Example Communication Service	2-2
The Platform Test Environment	2-2
Integration and Customization	2-2
Service Interceptors	2-3
Subscriber-centric Policy	2-3
Integration with External Systems	2-3

Using the Eclipse Wizard

About the Eclipse Wizard	3-1
Configure Eclipse	3-2
Prerequisites	3-2
Basic configuration of Eclipse environment	3-2
Configuring of the Eclipse Wizard	3-2
Using the Eclipse Wizard	3-3

Generating a Communication Service Project	3-3
Adding a Plug-in to a Communication Service Project	3-7
Removing a Plug-in from a Communication Service Project.....	3-8

Description of a Generated Project

Generated project	4-2
Communication Service Project	4-2
Plug-in.....	4-4
SOAP2SOAP Plug-in	4-6
Generated classes for a Plug-in	4-8
Interface: ManagedPluginService.....	4-9
Interface: PluginService.....	4-9
Interface: PluginInstanceFactory	4-9
Interface: PluginServiceLifecycle	4-10
PluginService	4-10
ManagedPlugin Skeleton	4-10
PluginInstance.....	4-11
PluginNorth.....	4-11
PluginNorth skeleton	4-13
RequestFactory Skeleton.....	4-13
Generated classes for a SOAP2SOAP Plug-in.....	4-13
Comparison with a Non-SOAP2SOAP Plug-in	4-14
Client Stubs	4-14
<Web Services Interface>_Stub.....	4-14
<Web Services Interface>	4-14
<Web Services Interface>Service_Impl	4-15
<Web Services Interface>Service	4-15
PluginInstance.....	4-15

PluginNorth	4-15
PluginSouth	4-16
RequestFactory	4-16
Build Files and Targets for a Communication Service Project	4-16
Main Build File	4-16
Communication Service Common Build File	4-17
Plug-in Build File	4-17
Ant Tasks	4-18
cs_gen	4-18
plugin_gen	4-20
cs_package	4-21
javadoc2annotation	4-23

Communication Service Example

Overview	5-1
High-level Flow for sendData (Flow A)	5-3
High-level Flow for startNotification and stopNotification (Flow B)	5-4
High-level flow for notifyDataReception (Flow C)	5-4
Interfaces	5-5
Web Service Interface Definition	5-5
Interface: SendData	5-5
Interface: NotificationManager	5-5
Interface: NotificationListener	5-7
Network Interface Definition	5-8
sendDataToNetwork	5-8
receiveData	5-8
Directory Structure	5-9
Directories for WSDL	5-11

Application-initiated traffic	5-11
Network-triggered traffic	5-11
Directories for Java Source	5-11
Communication Service Common	5-12
Plug-in	5-12
Directories for resources	5-13
Directories for Configuration of Plug-in	5-13
Directories for Build and Configuration of Builds	5-14
Directories for Classes, JAR, and EAR Files	5-15
Classes	5-16
Communication Service Common	5-17
ExceptionType	5-17
NotificationManagerPluginFactory	5-17
Plug-in Layer	5-18
ContextTranslatorImpl	5-18
ExamplePluginService	5-19
ExamplePluginInstance	5-20
ConfigurationStoreHandler	5-22
ExampleMBean	5-24
Management	5-24
NotificationHandlerNorth	5-24
NetworkToNotificationPluginAdapter	5-26
NetworkToNotificationPluginAdapterImpl	5-26
NotificationManagerPluginNorth	5-28
SendDataPluginNorth	5-29
SendDataPluginSouth	5-30
SendDataPluginToNetworkAdapter	5-31
SendDataPluginToNetworkAdapterImpl	5-31

FilterImpl	5-31
NotificationData	5-32
StoreHelper.	5-32
ExamplePluginInstance	5-34
ExamplePluginService	5-35
Store configuration	5-37
SLA Example	5-40

Container Services

Container service APIs	6-2
Class: InstanceFactory	6-3
Class: ClusterHelper	6-4
Service: EventChannel Service.	6-4
Plug-in.	6-5
Management	6-5
EDR	6-5
SLA Enforcement	6-5
Service Correlation	6-6
Interface: ExternalInvocation	6-7
Class: ExternalInvocatorFactory	6-8
Class: ServiceCorrelation	6-8
Implementing the ExternalInvocation Interface	6-8
Parameter Tunneling.	6-9
Storage Services	6-10
ConfigurationStore	6-10
Interfaces	6-11
StorageService	6-15
Store configuration file.	6-17

<store>	6-20
<db_table>	6-20
<query>	6-23
<provider-mapping>	6-26
<providers>	6-27
Shared libraries.	6-27

Communication Service Description

High-level components.	7-1
Communication Service Common	7-2
Plug-in	7-5
Plug-in Service and Plug-in Instance	7-6
States	7-6
PluginPool.	7-10
Interface: Plugin	7-10
Interface: PluginNorth.	7-11
Interface: PluginNorthCallBack	7-11
Interface: PluginSouth.	7-11
Interface: ManagedPluginService.	7-12
Interface: PluginService.	7-12
Interface: PluginInstanceFactory	7-12
Interface: PluginServiceLifecycle	7-12
Interface: ManagedPluginInstance	7-12
Interface: PluginInstance	7-13
Interface: PluginInstanceLifecycle	7-13
Class: RequestFactory.	7-14
Class: CallbackFactory	7-14
Interface: Callback	7-15

Class: RequestInfo	7-15
RequestIdentifierRequestInfo, if the request contains a request identifier.	7-16
Class: ServiceType	7-16
Interface: ContextMapperInfo	7-16
Interface: RequestContext	7-17
Class: ManagedPlugin	7-17
Class: AbstractManagedPlugin	7-17
Management	7-17
SLA Enforcement	7-17
Shared libraries	7-18

Annotations, EDRs, Alarms, and CDRs

About aspects and annotations	8-2
How aspects are applied	8-2
Context Aspect	8-3
EDR Generation	8-6
Exception scenarios	8-8
Adding data to the RequestContext	8-9
Using translators	8-10
Trigger an EDR programmatically	8-12
EDR Content	8-13
Using send lists	8-19
RequestContext and EDR	8-20
Categorizing EDRs	8-22
The EDR descriptor	8-22
Special characters	8-25
Values provided	8-26
Boolean semantic of the filters	8-26

Example filters	8-27
Check-list for EDR generation	8-33
Frequently Asked Questions about EDRs and EDR filters	8-34
Alarm generation	8-37
Trigger an alarm programmatically	8-38
Alarm content	8-39
CDR generation	8-41
Triggering a CDR	8-41
Trigger a CDR programmatically	8-41
CDR content	8-42
Additional_info column	8-46
Out-of-the box (OOTB) CDR support	8-47
Extending Statistics	8-48

Making Communication Services Manageable

Overview	9-1
Create Standard JMX MBeans	9-2
Create an MBean Interface	9-2
Implement the MBean	9-4
Register the MBean with the Runtime MBean Server	9-4
Use the Configuration Store to Persist Values	9-6

Using the Platform Test Environment

Overview	10-1
Installing and Running the Platform Test Environment	10-3
Navigating the Platform Test Environment GUI	10-4
The Tools Panel	10-5
The Tool Selector Panel	10-5

The Tool Action Panel	10-5
The Simulator Panel	10-19
The SLA Editor	10-20
Extending the Platform Test Environment	10-24
The Stateful SPI	10-25
The Stateless SPI	10-27
The Custom Base SPI	10-28
The Custom Results Provider SPI	10-30
The Custom Statistics Provider SPI	10-31
The Context API	10-32
The Module.xml Descriptor File	10-34
Using the Unit Test Framework (UTFW) with the Platform Test Environment	10-43

Service Interceptors

Overview	11-1
Interceptor Decisions and Request Flow	11-2
Decisions	11-2
Flow Control	11-6
Changing the invocation order	11-9
Standard Interceptors	11-9
Retry functionality for plug-ins	11-15
Custom Interceptors	11-16
Developing Custom Interceptors	11-16
Deploying Custom Interceptors	11-18

Subscriber-centric Policy

Service Classes and the Subscriber SLA	12-1
The <reference> tag	12-2

The <restriction> tag	12-3
Managing the Subscriber SLA	12-4
The Profile Provider SPI and Subscriber Contracts	12-4
Deploying the Custom Profile Provider	12-6
Subscriber Policy Enforcement	12-6
Do Relevant Subscriber Contracts Exist	12-6
.	12-7
Is There Adequate Budget for the Contracts	12-11

Creating an EDR Listener and Generating SNMP MIBs

Overview of External EDR listeners	13-1
Example using a pure JMS listener	13-3
Example using JMSListener utility with no filter	13-3
Using JMSListener utility with a filter	13-4
Description of EDR listener utility	13-5
Class JMSListener	13-6
Class EdrFilterFactory	13-6
Class EdrData	13-7
Class ConfigDescriptor	13-7
Class EdrConfigDescriptor	13-7
Class AlarmConfigDescriptor	13-8
Class CdrConfigDescriptor	13-8
Updating EDR configuration files	13-8
Generating SNMP MIBs	13-8

Converting Traffic Paths and Plug-ins to Communication Services

Converting Network Protocol Plug-ins	14-1
--	------

Converting Traffic Paths.	14-2
Checklist.	14-2

Policy

Overview.	15-1
Policy Request Data	15-2
Adding a New Rule.	15-4
Mapping PolicyRequest Data	15-5
Creating a New Rule File by Extending an Existing File: an Example	15-7
Using RequestContext Parameters Defined in Service Level Agreements	15-8

Callable Policy Web Service

Introduction.	16-2
Callable Policy Web Service interface definition.	16-3
Endpoints.	16-3
Detailed service description.	16-3
Policy Evaluation	16-3
Policy management	16-4
XML Schema data type definition.	16-4
AdditionalDataValue structure	16-4
AdditionalDataValueType enumeration	16-4
Interface: Policy.	16-5
Operation: evaluate.	16-5
Interface: PolicyManagement	16-6
Operation: viewRuleFile	16-6
Operation: deleteRuleFile.	16-7
Operation: loadRules	16-8
Operation: listRuleFiles	16-9

Rule files.	16-10
---------------------	-------

Checklist

Document Roadmap

The following sections describe the audience for, and organization of, this document:

- [Document Scope and Audience](#)
- [Guide to This Document](#)
- [Terminology](#)
- [Related Documentation](#)

Document Scope and Audience

This document describes the WebLogic Network Gatekeeper Platform Development Studio, a framework for creating and testing new extension Communication Services. The intended audience of this document consists of system integrators and field engineers who need to extend the out-of-the-box functionality of WebLogic Network Gatekeeper.

Guide to This Document

The document contains the following chapters:

[Chapter 1, “Document Roadmap”](#): This chapter

[Chapter 2, “Overview of the Platform Development Studio”](#): A high level description of the capabilities of the PDS

[Chapter 3, “Using the Eclipse Wizard”](#): Setting up the Eclipse Wizard to generate extension projects

[Chapter 4, “Description of a Generated Project”](#): The elements of a generated Communication Service project.

[Chapter 5, “Communication Service Example”](#): A description of the example Communication Service provided with the Platform Development Studio.

[Chapter 6, “Container Services”](#): Accessing Network Gatekeeper’s Core functionality

[Chapter 7, “Communication Service Description”](#): A component by component description of a Communication Service

[Chapter 8, “Annotations, EDRs, Alarms, and CDRs”](#): Creating EDRs, CDRs, and Alarms

[Chapter 9, “Making Communication Services Manageable”](#): Rendering extension Communication Services manageable by the Network Gatekeeper Console extension or other management tools

[Chapter 10, “Using the Platform Test Environment”](#): Using the testing tools framework

[Chapter 11, “Service Interceptors”](#): An overview of service interceptors, a description of the standard out of the box ones, and of developing custom versions

[Chapter 12, “Subscriber-centric Policy”](#): Creating a policy mechanism based on individual subscriber preferences and permissions

[Chapter 13, “Creating an EDR Listener and Generating SNMP MIBs”](#): Using and integrating external EDR listeners and generating SNMP MIBs

[Chapter 14, “Converting Traffic Paths and Plug-ins to Communication Services”](#): Converting extensions built on previous versions of Network Gatekeeper

[Chapter 15, “Policy”](#): Using and extending Policy mechanisms in Communication Services

[Chapter 16, “Callable Policy Web Service”](#): Integrating using the Callable Policy Web Service

[Chapter 17, “Checklist”](#): A checklist for creating extensions.

Terminology

The following terms and acronyms are used in this document:

- **Account**—A registered application or service provider. An account belongs to an account group, which is tied to a common SLA
- **Account group**—Multiple registered service providers or services which share a common SLA

- **Administrative User**—Someone who has privileges on the Network Gatekeeper management tool. This person has an administrative user name and password
- **Alarm**—The result of an unexpected event in the system, often requiring corrective action
- **API**—Application Programming Interface
- **Application**—A TCP/IP based, telecom-enabled program accessed from either a telephony terminal or a computer
- **Application-facing Interface**—The Application Services Provider facing interface
- **Application Service Provider**—An organization offering application services to end users through a telephony network
- **AS**—Application Server
- **Application Instance**—An Application Service Provider from the perspective of internal Network Gatekeeper administration. An Application Instance has a user name and password
- **CBC**—Content Based Charging
- **End User**—The ultimate consumer of the services that an application provides. An end user can be the same as the network subscriber, as in the case of a prepaid service or they can be a non-subscriber, as in the case of an automated mail-ordering application where the subscriber is the mail-order company and the end user is a customer to this company
- **Enterprise Operator** —See Service Provider
- **Event**—A trackable, expected occurrence in the system, of interest to the operator
- **HA** —High Availability
- **HTML**—Hypertext Markup Language
- **IP**—Internet Protocol
- **JDBC**—Java Database Connectivity, the Java API for database access
- **Location Uncertainty Shape**—A geometric shape surrounding a base point specified in terms of latitude and longitude. It is used in terminal location
- **MAP**—Mobile Application Part
- **Mated Pair**—Two physically distributed installations of WebLogic Network Gatekeeper nodes sharing a subset of data allowing for high availability between the nodes

- MM7—A multimedia messaging protocol specified by 3GPP
- MPP—Mobile Positioning Protocol
- Network Plug-in—The WebLogic Network Gatekeeper module that implements the interface to a network node or OSA/Parlay SCS through a specific protocol
- NS—Network Simulator
- OAM —Operation, Administration, and Maintenance
- Operator—The party that manages the Network Gatekeeper. Usually the network operator
- OSA—Open Service Access
- PAP—Push Access Protocol
- Plug-in—See Network Plug-in
- Plug-in Manager—The Network Gatekeeper module charged with routing an application-initiated request to the appropriate network plug-in
- Policy Engine—The Network Gatekeeper module charged with evaluating whether a particular request is acceptable under the rules
- Quotas—Access rule based on an aggregated number of invocations. See also Rates
- Rates—Access rule based on allowable invocations per time period. See also Quotas
- Rules—The customizable set of criteria - based on SLAs and operator-desired additions - according to which requests are evaluated
- SCF—Service Capability Function or Service Control Function, in the OSA/Parlay sense.
- SCS—Service Capability Server, in the OSA/Parlay sense. WebLogic Network Gatekeeper can interact with these on its network-facing interface
- Service Capability—Support for a specific kind of traffic within WebLogic Network Gatekeeper. Defined in terms of Communication Services
- Service Provider—See Application Service Provider
- SIP—Session Initiation Protocol
- SLA—Service Level Agreement
- SMPP—Short Message Peer-to-Peer Protocol

- SMS—Short Message Service
- SMSC—Short Message Service Centre
- SNMP—Simple Network Management Protocol
- SOAP—Simple Object Access Protocol
- SPA—Service Provider APIs
- SS7—Signalling System 7
- Subscriber—A person or organization that signs up for access to an application. The subscriber is charged for the application service usage. See End User
- SQL—Structured Query Language
- TCP—Transmission Control Protocol
- Communication Service—offers a service to an application
- USSD—Unstructured Supplementary Service Data
- VAS—Value Added Service
- VLAN—Virtual Local Area Network
- VPN—Virtual Private Network
- WebLogic Network Gatekeeper Container—The container hosting communication services.
- WSDL —Web Services Definition Language
- XML—Extended Markup Language

Related Documentation

This developer's guide is part of a set of documentation for Network Gatekeeper itself. These documents include:

- [*System Administrator's Guide*](#)
- [*Concepts and Architectural Overview*](#)
- [*Integration Guidelines for Partner Relationship Management*](#)

Document Roadmap

- *SDK User Guide*
- *Managing Accounts and SLAs*
- *Statement of Compliance and Protocol Mapping*
- *Application Development Guide*
- *Communications Service Reference*
- *Handling Alarms*
- *Licensing*
- *Installation Guide*

Overview of the Platform Development Studio

WebLogic Network Gatekeeper provides substantial functionality right out of the box. But because all networks are different, matching the particular requirements and capabilities of some networks sometimes means that Network Gatekeeper must be extended or that certain aspects of it must be closely integrated with existing network functionality. The Platform Development Studio is designed to ease this process. It consists of two main sections:

- [Creating New Communication Services](#)
- [Integration and Customization](#)

Creating New Communication Services

Networks change. Existing functionality is parsed in new ways to support new features. New nodes with new or modified abilities are added. Because of WebLogic Network Gatekeeper's highly modular design, exposing these new features to partners is a straightforward proposition. The extension portion of the Platform Development Studio provides an environment in which much of the mechanics of creating extensions is taken care of, allowing extension developers to focus on only those parts of the system that correspond directly to their specific needs. This aspect consists of three main parts

- [The Eclipse Wizard](#)
- [Example Communication Service](#)
- [The Platform Test Environment](#)

The Eclipse Wizard

At the core of the extension portion of the Platform Development Studio is an Eclipse plug-in that creates projects based on the responses that the developer makes to an Eclipse Wizard. The developer supplies some basic naming information and the location of a WSDL for each application facing interface that the Communication Service is meant to support, and the Wizard generates either a complete Communication Service project, or a network plug-in only project. For more information on setting up the Eclipse Plug-in and running the Wizard, see [Chapter 3, “Using the Eclipse Wizard.”](#) To see an example of a generated project, see [Chapter 4, “Description of a Generated Project.”](#) To get an understanding of the Network Gatekeeper features with which your Communication Service will interact, see [Chapter 15, “Policy,”](#) [Chapter 6, “Container Services,”](#) [Chapter 8, “Annotations, EDRs, Alarms, and CDRs,”](#) and [Chapter 9, “Making Communication Services Manageable.”](#)

Example Communication Service

To give you a concrete sense of the task of generating a new Communication Service, the Platform Development Studio contains an entire example Communication Service, which is buildable and runnable. Based on a very simple Web Service interface and an equally simple model of an underlying network protocol, this Communication Service demonstrates the entire range of tasks that you will encounter in creating your own Communication Service. For more information, see [Chapter 5, “Communication Service Example.”](#)

The Platform Test Environment

To simplify the testing of your Communication Service, the Platform Development Studio includes the Platform Test Environment, which provides an extensible suite of tools for testing Communication Services and the Unit Test Framework, which supplies an abstract base class, `WlngBaseTestCase`, which includes mechanisms for connecting to the Platform Test Environment. As well, there is a complete set of sample tools created to interact with the example Communication Service. For more information, see [Chapter 10, “Using the Platform Test Environment.”](#)

Integration and Customization

New Communication Services are not the only aspect of Network Gatekeeper that can be handled using the Platform Development Studio. To help integrate Network Gatekeeper into the installation environment, three other aspects of customization are supported:

- [Service Interceptors](#)
- [Subscriber-centric Policy](#)
- [Integration with External Systems](#)

Service Interceptors

Service interceptors provide Network Gatekeeper with a mechanism for intercepting and manipulating a request as it flows through any arbitrary Communication Service. They offer an easy way to modify the request flow, simplify routing mechanisms for plug-ins, and centralize policy and SLA enforcement. Out of the box, Network Gatekeeper uses these modules as part of its internal functioning, but operators can also choose to create new interceptors, or to rearrange the order in which the interceptors are used, in order to customize their functionality. [Chapter 11, “Service Interceptors”](#) describes the request flow through interceptors, lists the standard interceptors and explains how to rearrange interceptors or to create new custom versions.

Subscriber-centric Policy

Out of the box the Network Gatekeeper administration model allows operators to manage application service provider access to the network at increasingly granular levels of control. Using the Platform Development Studio, operators can extend that model to encompass their subscribers, giving the operator the ability to offer those subscribers a highly personalized experience while protecting their privacy and keeping their subscriber data safe within the operator’s domain.

Operators create a Subscriber SLA, based on a provided schema, which describes sets of *service classes*. The service classes define access relationships with the services of particular Service Provider and Application Groups, along with default rates and quotas. *Profile providers* created by the operator or integrator using the provided Profile Provider SPI then associate those service classes with subscriber URIs, forming *subscriber contracts*. These contracts are used to evaluate requests and to generate subscriber budgets, which are used by the normal request traffic policy evaluation flow. A single subscriber can be covered by multiple subscriber contracts, based on that individual subscriber’s desires. [Chapter 12, “Subscriber-centric Policy”](#) covers the process for setting this up.

Integration with External Systems

Finally, the Platform Development Studio provides mechanisms to support the integration of Network Gatekeeper with external network systems, including:

- EDR listeners, covered in [Chapter 13, “Creating an EDR Listener and Generating SNMP MIBs”](#)
- Alarm monitoring using SNMP, covered in [Chapter 13, “Creating an EDR Listener and Generating SNMP MIBs”](#)
- Callable policy using JMX, covered in [Chapter 16, “Callable Policy Web Service”](#)

Additional integration points, not covered in the PDS, are provided by:

- The Partner Relationship Management interfaces, for creating Partner Management portals, covered in [Integration Guidelines for Partner Relationship Management](#), a separate document in this set
- JMX for Management, for non-console based management, covered by the WLS documents [Developing Custom Management Utilities with JMX](#) and [Developing Manageable Applications with JMX](#).

Using the Eclipse Wizard

This section describes using the Eclipse Wizard to generate Communication Services:

- [“About the Eclipse Wizard” on page 3-1](#)
- [“Configure Eclipse” on page 3-2](#)
- [“Using the Eclipse Wizard” on page 3-3](#)
 - [“Generating a Communication Service Project” on page 3-3](#)
 - [“Adding a Plug-in to a Communication Service Project” on page 3-7](#)
 - [“Removing a Plug-in from a Communication Service Project” on page 3-8](#)

About the Eclipse Wizard

The Eclipse Wizard is a plug-in that enables an Eclipse user to create Network Gatekeeper Communication Services. The extension projects are created using wizards that customize the project depending on which type of extension is being developed. Two types of extensions can be created:

- Communication Services
- Network protocol plug-ins for existing Service Facades (application-facing interfaces)

The Eclipse Wizard generates classes and Ant build files for both types of extensions. In addition, there is a separate build file with Ant targets for packaging the extension for deployment.

Configure Eclipse

Prerequisites

- Eclipse 3.2 or higher version must be installed
- Ant 1.6.5 must be installed
Use the Ant provided with WebLogic Server. It is located in
`$BEA_HOME/modules/org.apache.ant_1.6.5.`

Basic configuration of Eclipse environment

To do the basic configuration of the Eclipse environment:

1. Start Eclipse
2. Open the Preferences window, **Window**→**Preferences...**
 - a. In **Java**→**Installed JREs**, make sure that the JRE used is the JRE installed with the Network Gatekeeper. This is installed in `$BEA_HOME/<jdk version>/jre`
 - b. In **Ant**→**runtime**, make sure Ant Home is set to the directory in which you have installed Ant.

Configuring of the Eclipse Wizard

To configure the Eclipse Wizard, starting in Eclipse:

1. Open the Preferences window, **Window**→**Preferences...**
2. In **WLNG Platform Development Studio**, configure the following:

WLNG Home Directory The directory of the Network Gatekeeper installation. This provides references to WebLogic Server APIs. In the default installation, this would be `$BEA_HOME/wlng400`.

WLNG PDS Home Directory The directory of the Network Gatekeeper Platform Development Studio installation. This provides references to Network Gatekeeper APIs, extension points and third party APIs. In the default installation, this would be `$BEA_HOME/wlng_pds400`

JDK Installation Directory The JDK installation directory for Network Gatekeeper, for example `$BEA_HOME/jdk150_14`.

Using the Eclipse Wizard

Generating a Communication Service Project

A Communication Service project is based on a WSDL file and a set of attributes given when running the **Communication Service Project** wizard.


The WSDL defining the application-facing interface must adhere to the following:



- Attribute name in `<wsdl:service>` must include the suffix **Service**.
- Attribute name in `<wsdl:port>` must be the same as the name attribute in `<wsdl:service>`, excluding the suffix **Service**.

To generate an Communication Service project:

1. In Eclipse, choose **File**→**New Project**.

This opens the **New Project** window.

In this window...	Perform the following action...
Select Wizard	<p>Make sure WLNG Platform Development Studio→Communication Service Project is selected.</p> <p>Click Next to proceed. You may cancel the wizard at any time by clicking Exit. You may go back to a previous window by clicking Previous.</p>
Create a Communication Service	<p>Enter a Project Name and choose a location for your project.</p> <p>You can choose:</p> <ol style="list-style-type: none">1. To create an entirely new Communication Service2. To create a new Service Facade (application-facing interface) and the common parts of the Service Enabler layer for an existing plug-in3. To create a new network plug-in that uses the Service Facade and common parts of the Service Enabler of a currently existing Communication Service. <p>If you wish to do 3, check the check-box Use predefined communication service and from the drop-down list select the Service Facade for which you want to create a plug-in.</p> <p>If you wish to do 1 or 2, leave the box unchecked.</p> <p>Click Next to continue.</p> <p>If you checked the Use predefined check box, the Define the Plug-in Information window is displayed. Go to Define the Plug-in Information instructions below.</p> <p>If you did not check it, the Define the Communication Service is displayed.</p>
<div>Define the Communication Service</div> <div><div>Configure Service WSDL Files</div></div>	<p>For each WSDL file that includes the service definition <i>to be implemented</i> by the new Communication Service:</p> <div><div>Click  , browse to the WSDL file, select it, and click OK.</div></div>

In this window...	Perform the following action...
Define the Communication Service <ul style="list-style-type: none"> Configure Callback WSDL Files 	<p>For each WSDL file that includes the callback service definition <i>to be used</i> by the new Communication Service in sending information to the service provider's application:</p> <div style="text-align: center;">  </div> <p>Click , browse to the WSDL file, select it, and click OK.</p>
Define the Communication Service <ul style="list-style-type: none"> Communication Service Properties 	<p>Company: Set your company name, to be used in META-INF/MANIFEST.MF.</p> <p>Version: Set the version, to be used in META-INF/MANIFEST.MF.</p> <p>Identifier: Create an identifier to tie together a collection of Web Services. Will be a part of the names of the generated war and jar files and the service type for the Communication Service:</p> <pre><Communication Service identifier>.war and <Communication Service identifier>_callback.jar</pre> <p>Service Type: Set the service type. Used in EDRs, statistics, etc. For example: SmsServiceType, MultimediaMessagingServiceType.</p> <p>Java Class Package Name: Set the package names to be used. For example: com.mycompany.service</p> <p>Web Services Context path: Set the context path for the Web Service. For example: myService</p>

In this window...	Perform the following action...
Define the Communication Service <ul style="list-style-type: none"> WSDL Properties 	<p>Support SOAP Attachment Check this box if SOAP with attachments must be supported.</p> <p>Enable SOAP Plug-in Generation Check this box if generating a SOAP-SOAP Communication Service.</p> <p>Include WLNG Exceptions in WSDL Check this box if Network Gatekeeper specific exception should be added to the Web Service for SOAP-SOAP Communication Services. If this option is selected, WLNG custom exceptions are appended to the WSDL. This is done to provide a better error handling capability. The client application should handle these exceptions. In SOAP-SOAP communication case, it is recommended, though not required, to select this option</p> <p>If you are creating an entirely new Communication Service, including a new plug-in, click Next. The Define the Plug-in Information window opens.</p> <p>If you are not creating a new plug-in, you have completed the Wizard. Click Finish to start the code generation process.</p>


In this window...	Perform the following action...
Define the Plug-in information	<p>For each plug-in to be created in the Communication Service project:</p> <div data-bbox="534 414 610 487" data-label="Image"> </div> <p>Click</p> <p>This opens a pop-up window with the following fields:</p> <p>Protocol: An identifier for the network protocol the plug-in implements. Used as a part of the names of the generated jar file: <code><Communication Service identifier>_<protocol>.jar</code> and the service name <code>Plugin_<Communication Service identifier>_<protocol></code></p> <p>Schemes: Address schemes the plug-in can handle. Use a comma separated list if multiple schemes are supported. For example: tel: or sip:</p> <p>Package Name: Package names to be used.</p> <p>Company: Used in <code>META-INF/MANIFEST.MF</code>.</p> <p>Version: Used in <code>META-INF/MANIFEST.MF</code>.</p> <p>SOAP Plug-in: Check this box if the plug-in is a SOAP2SOAP plug-in. Enable SOAP Plug-in Generation must have been selected in the previous step.</p> <p>Click OK.</p> <p>The plug-in definitions are added to the list of plug-ins.</p> <div data-bbox="534 1150 610 1223" data-label="Image"> </div> <p>Click to remove the selected plug-in.</p> <p>Click Finish to start the code generation for the plug-in(s).</p>

Adding a Plug-in to a Communication Service Project

To add a plug-in to an existing Communication Service project:

1. In the Eclipse package explorer, right-click the project for the Communication Service project, and choose **Properties**.

This opens the **Properties** window for the Communication Service project.

In this window...	Perform the following action...
Plugin Configuration	<p>A list of plug-ins defined for the Communication Service project is displayed.</p> <p>For each plug-in to be created in the Communication Service project:</p> <div data-bbox="471 536 547 607"></div> <p>Click</p> <p>This opens a pop-up window with the following fields:</p> <p>Protocol: An identifier for the network protocol the plug-in implements. Used as a part of the names of the generated jar file: <code><Communication Service identifier>_<protocol>.jar</code> and the service name <code>Plugin_<Communication Service identifier>_<protocol></code></p> <p>Schemes: Address schemes the plug-in can handle. Use a comma separated list if multiple schemes are supported. For example: tel: or sip:</p> <p>Package Name: Package names to be used.</p> <p>Company: Used in <code>META-INF/MANIFEST.MF</code>.</p> <p>Version: Used in <code>META-INF/MANIFEST.MF</code>.</p> <p>SOAP Plug-in: Check this box if the plug-in is a SOAP2SOAP plug-in.</p> <p>Click OK.</p> <p>The plug-in definitions are added to the list of plug-ins.</p> <p>Click Finish to start the code generation for the plug-in(s).</p>

Removing a Plug-in from a Communication Service Project

To remove a a plug-in from an existing Communication Service project:

1. In the Eclipse package explorer, right-click the project for the Communication Service project, and choose **Properties**.

This opens the Properties Window for the Communication Service project.

In this window...	Perform the following action...
Plugin Configuration	<p>A list of plug-ins defined for the Communication Service project is displayed.</p> <p>For each plug-in to be removed from the Communication Service project:</p> <div data-bbox="920 534 997 607" data-label="Image"> </div> <ul style="list-style-type: none"> Select the plug-in to be removed and click <p>The plug-in definitions are removed from the list.</p> <ul style="list-style-type: none"> Click Apply to remove the plug-in part(s) from the Communication Service project. <p>Warning: This removes <i>all</i> parts of the project, including any manually edited or added files.</p> <ul style="list-style-type: none"> Click Restore Defaults to restore the plug-in definition list.

- Click **OK** or **Cancel** to close the **Properties** window.

Using the Eclipse Wizard

Description of a Generated Project

The section describes a project generated from the Eclipse Wizard:

- [Generated project](#)
 - [Communication Service Project](#)
 - [Plug-in](#)
 - [SOAP2SOAP Plug-in](#)
- [Generated classes for a Plug-in](#)
 - [Interface: ManagedPluginService](#)
 - [PluginService](#)
 - [PluginInstance](#)
 - [PluginNorth](#)
 - [RequestFactory Skeleton](#)
- [Generated classes for a SOAP2SOAP Plug-in](#)
 - [Comparison with a Non-SOAP2SOAP Plug-in](#)
 - [Client Stubs](#)
 - [PluginInstance](#)
 - [PluginNorth](#)
 - [PluginSouth](#)

- [RequestFactory](#)
- [Build Files and Targets for a Communication Service Project](#)
 - [Main Build File](#)
 - [Communication Service Common Build File](#)
 - [Plug-in Build File](#)
 - [Ant Tasks](#)

Generated project

Communication Service Project

Generating a Communication Service project creates the directory structure illustrated in [Listing 4-1](#).

The base directory is the directory given in the Eclipse Wizard input field **Identifier**. It contains the following files:

- `build.properties`: properties file for the build files:
 - `wlng.home` is set to `$WLNG_HOME`, the base directory for the Network Gatekeeper installation.
 - `pds.home` is set to `$PDS_HOME`, the base directory for the Platform Development Studio.
- `build.xml`: the main file for the project, that is the build file for the Communication Service and references to any other plug-in specific build files in the project. See [Main Build File](#).
- `common.xml`: properties, ant task and targets used by all build files in the project.

The directories and files in bold in [Listing 4-1](#) are generated when building the common parts of the Communication Service; the others are generated by the Eclipse Wizard.

Listing 4-1 Generated project for Communications Services Common

```
<Eclipse Project Name>  
+- build.properties
```

```

+- common.xml
+- build.xml
+- <Identifier given in Eclipse Wizard>
| +- dist //Generated by target dist in <Eclipse Project Name>/build.xml
| | +- <Package name>.store_<version.jar // Example store configuration
| | +- wlng_at_<Identifier>.ear //Deployable in access tier
| | +- wlng_nt_<Identifier>.ear //Deployable in network tier
| +- common
| | +- build.xml //Build file for the common parts of the communication service
| | +- dist //Generated by target dist on
| | | //<Eclipse Project Name>/common/build.xml
| | | +- request_factory_skel //Skeletons for the RequestFactory,
| | | | //one class for each service WSDL
| | | +- tmp //Used during build. Contains classes, source,
| | | | //definitions, WSDLs, templates, and more.
| | | +- <Identifier>.war // Web Service implementation
| | | +- <Identifier>_callback.jar // Service callback EJB for
| | | | //the communication service
| | | +- <Identifier>_callback_client.jar //Service call-back EJB used by
| | | | // the plug-in.
| | | +- <Identifier>_service.jar // Service EJB
| | | | // for the communication service
| | +- resources // Contains application.xml and weblogic-application.xml
| | | // for the access and network tier EAR files respectively.
| | | // The files are packaged in the EAR files META-INF directory
| | +- src // Source directory for communication service common
| | | +- <Package name>/plugin
| | | | +- <Web Services interface>PluginFactory // One per interface

```

```
// defined in the
// Service WSDL files.
```

If the check-box **Include WLNG Exceptions** was checked when generating the Communication Service, the following exception definitions are added to the Web Service:

- PolicyException - Any policy based exceptions.
- RoutingException - Any exceptions during the routing of the request.
- ServiceException - Any other internal exceptions.

The exceptions are added only to the service facade, not to the plug-in to network interface.

If the exceptions listed above are present in the original WSDL they are reused; if not they are added.

Plug-in

When creating a plug-in for a given Communication Service, the directory structure illustrated in [Listing 4-2](#) is created under the top-level directory. The base directory depends on the type of Communication Service the plug-in belongs to, such as, for example, `px21_multimedia_messaging`, or `px21_sms`. It also depends on whether the plug-in is for an existing Communication Service or for a new one.

If the plug-in is for an existing Communication Service, it is generated under one of the following directories:

- `px30_audio_call` for plug-ins for Parlay X 30 Audio Call
- `px21_call_notification` for Parlay X 2.1 Call Notification
- `px30_call_notification` for Parlay X 3.0 Call Notification
- `px21_multimedia_messaging` for Parlay X 2.1 Multimedia Messaging
- `px21_presence` for Parlay X 2.1 Presence
- `ews_push_message` for Extended Web Services WAP Push
- `px21_sms` for Parlay X 2.1 Short Messaging
- `ews_susbcriber_profile` for Extended Web Services Subscriber Profile

- `px21_terminal_location` for Parlay X 2.1 Terminal Location
- `px21_third_party_call` for Parlay X 2.1 Third Party Call
- `px30_third_party_call` for Parlay X 3.0 Thrid Party Call

If it is for a new Communication Service, the base directory is given in the **Identifier** entry field in the Eclipse Wizard.

The base directory contains the directory `plugins`, which contains subdirectories for each protocol that is being added. The names of the directories are the same as the name chosen for the **Protocol** field in the Eclipse Wizard.

Each of the sub-directories for a plug-in contains the following files:

- `build.xml`: The build file for the plug-in, see [Plug-in Build File](#).

Each plug-in sub-directory also contains the directories:

- `confi`: The directory that includes an instancemap that will be used by the InstanceFactory to create instances for the plug-in interface implementations.
- `dist`: The directory where the final deployable plug-in jar will end up. If a new plug-in skeleton is generated from the build file it will be generated here.
- `resources`: The directory that contains deployment descriptors for the plug-in.
- `src`: The directory that contains the generated plug-in code.
- `storage`: The directory that contains the configuration file for the Storage service.

The directories and files in bold in [Listing 4-2](#) are generated when building the plug-in, the others are generated by the Eclipse Wizard.

Listing 4-2 Generated project for a plug-in

```
| +- plugins // Container directory for all plug-ins for
|           // the communication service
| | +- <Protocol> // One specific plug-in
| | | +- build.xml // Build file for the plug-in
| | | +- build // Used during the build process
| | | +- config //
```

Description of a Generated Project

```
| | | | +- instance_factory
| | | | +- instancemap //Instance map
| | | +- dist // Generated by target dist in build.xml for the plug-in
| | | | +- <Identifier>_<Protocol>_plugin.jar
| | | | +- <Package name>.store_<version>.jar
| | | +- resources // Contains parts of weblogic-extension.xml
| | | | // for the network tier EAR file.
| | | | // the file is packaged in the EAR file's META-INF directory
| | | +- src
| | | | +- <Package name> // Directory structure reflecting
| | | | | // plug-in package name
| | | | | +- management // Example MBean
| | | | | | +- MyTypeMBean.java
| | | | | | +- MyTypeMBeanImpl.java
| | | | | +- <Web Services interface> // One per Service WSDL
| | | | | | +- north
| | | | | | | +- <Web Services interface>PluginImpl.java
| | | | | | | | // Implmentation of the interface
| | | | | | +- <Type>PluginInstance.java
| | | | | | +- <Type>PluginService.java
| | | | | | | // PluginService implementation
| | | +- storage //Example of a store configuration.
| | | | +- wlng-cachestore-config-extensions.xml
```

SOAP2SOAP Plug-in

When creating a SOAP2SOAP plug-in, the directory structure described in [Plug-in](#) is created under the top-level directory. In addition, the directories and files in [Listing 4-3](#) are generated.

The directories and files in bold are created when building the plug-in; the others are generated by the Eclipse Wizard.

Note: Only the deployable artifacts are relevant. The generated code for SOAP2SOAP type of plug-ins should not be modified.

Listing 4-3 Generated project for a SOAP2SOAP plug-in

```
| +- plugins // Container directory for all plug-ins for
|           // the communication service
| | +- <Protocol> // One specific plug-in
| | | +- build.xml // Build file for the plug-in
| | | +- build // Used during the build process
| | | +- config //
| | | | +- instance_factory
| | | | | +- instancemap //Instance map
| | | +- dist // Generated by target dist in build.xml for the plug-in
| | | | +- <Identifier>_<Protocol>_plugin.jar
| | | | +- <Package name>.store_<version>.jar //unused, empty
| | | +- resources // Contains parts of weblogic-extension.xml
| |           // for the network tier EAR file.
|           // the file is packaged in the EAR file's META-INF directory
| | | +- src
| | | | +- <Package name> // Directory structure reflecting
| | | |           // plug-in package name
| | | | | +- client // Implementation of Web Service client
| | | | | | +- <Web Services interface>_Stub.java
| | | | | | +- <Web Services interface>.java
| | | | | | +- <Web Services interface>Service_Impl.java
```

Description of a Generated Project

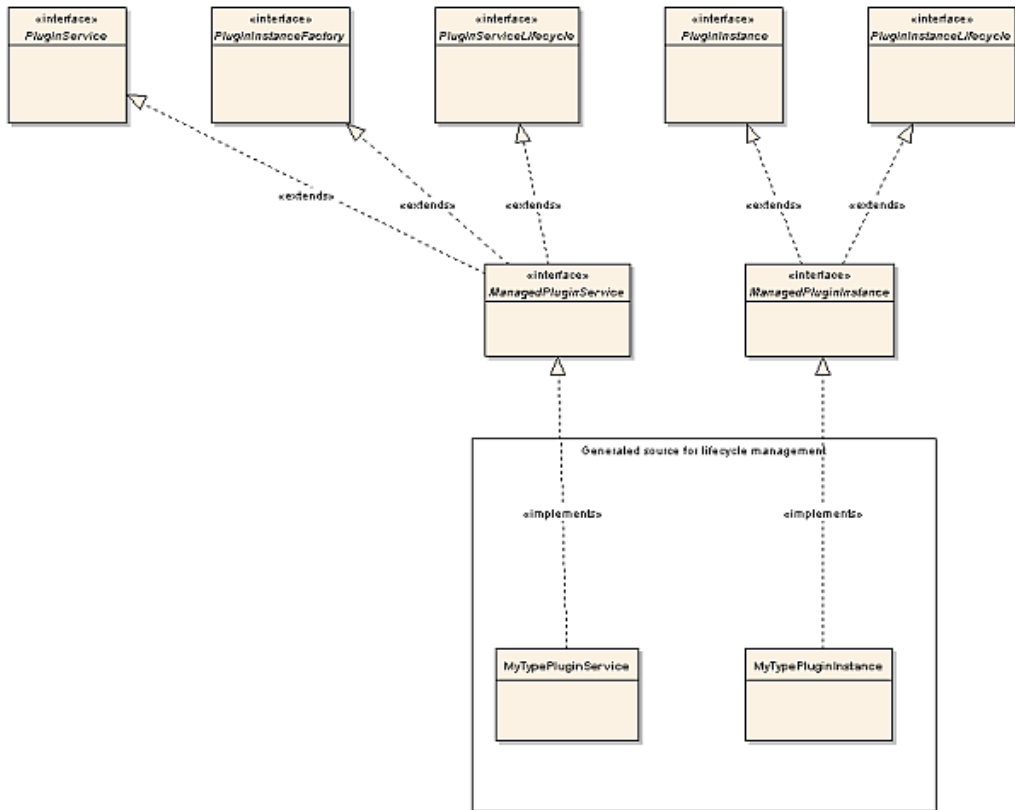
```
| | | | | +- <Web Services interface>Service.java
| | | | | +- <Web Services call-back interface> // One per Call-back WSDL
| | | | | +- south
| | | | | | +- <Web Services interface>PluginSouth.java
| | | | | | | | // Interface for network-triggered requests
| | | | | | +- <Web Services interface>PluginSouthImpl.java
| | | | | | | | // Implementation of the interface
| | | | | +- <Web Services interface> // One per Service WSDL
| | | | | +- north
| | | | | | +- <Web Services interface>PluginImpl.java
| | | | | | | | // Implementation of the interface
| | | | | +- <Type>PluginInstance.java
| | | | | +- <Type>PluginService.java
| | | | | | | | // PluginService implementation
| | | | | +- schema // Java Representation of the schemas in the WSDLs
| | | | | | +- <Package name> // Directory structure reflecting
| | | | | | | | // namespace in WSDL
| | | +- storage //Example of a store configuration. Empty.
| | | +- wsdl // WSDLs and XML-to-Java mappings.
| +- <Protocol Identifier_callback.war // Web Service implementation
| | | | | // for the SOAP2SOAP plug-in
```

As illustrated in [Listing 4-3](#), a WAR file for the plug-in is generated. This WAR file contains the Web Service for network-triggered requests. It is only generated if there is a notification WSDL defined at generation-time. It will be packaged in the EAR for the Service Enabler.

Generated classes for a Plug-in

The generated classes are listed below.

Figure 4-1 Example class diagram of the generated plug-in classes for life-cycle management and relationship with other interfaces



Interface: ManagedPluginService

The interface a plug-in service needs to implement.

It extends the interfaces `PluginService`, `PluginInstanceFactory` and `PluginServiceLifecycle`.

Interface: PluginService

The interface that defines the plug-in service when it is registered in the Plug-in Manager.

Interface: PluginInstanceFactory

The factory that allows a plug-in service to create plug-in instances.

Interface: PluginServiceLifecycle

The interface that defines the lifecycle for a plug-in service. See [States](#).

PluginService

Class.

Implements `com.bea.wlcp.wlng.api.plugin.ManagedPluginService`.

Defines the life-cycle for a plug-in service, see [States](#).

Also holds the data that is specific for the plug-in instance.

The actual class name is `<Communication Service Type>PluginService`. This class manages the life-cycle for the plug-in service, including implementing the necessary interfaces that make the plug-in deployable in Network Gatekeeper. It is also responsible for registering the north interfaces with the Plug-in Manager. At startup time it uses the InstanceFactory to create one instance of each plug-in service and at activation time it registers these with the Plug-in Manager. The InstanceFactory uses an instancemap to find out which class it should instantiate for each plug-in interface implementation. The instance map is found under the `resource` directory.

ManagedPlugin Skeleton

The `ManagedPlugin` skeleton implements the following methods related to life-cycle management and should be adjusted for the plug-in:

- `doStarted()` - plug-in specific functionality for being started.
- `doActivated()` - plug-in specific functionality for being activated.
- `doDeactivated()` - plug-in specific functionality for being deactivated.
- `doStopped()` - plug-in specific functionality for being stopped.
- `handleForceSuspending()` - Called when a FORCE STOP/SHUTDOWN has been issued.
- `handleResuming()` - Transitions the plug-in from ADMIN to ACTIVE state in which it begins to accept traffic.
- `handleSuspending(CompletionBarrier barrier)` - Called in a normal re-deployment when the plug-in is taken from ACTIVE to ADMIN state.

- `isActive()` - reports back true or false. If false, no application-initiated requests are routed to the plug-in.

In addition, this class defines which address schemes the plug-in can handle, in `private static final String[] SUPPORTED_SCHEMES`.

PluginInstance

Class.

Implements `com.bea.wlcp.wlng.api.plugin.ManagedPluginInstance`.

Defines the life-cycle for a plug-in instance, see [States](#).

The actual class name is `<Communication service Type>PluginInstance`. This class manages the life-cycle for the plug-in instance including implementing the necessary interfaces that make the plug-in an instance in Network Gatekeeper.

It is also responsible for instantiating the classes that implement the traffic interfaces, and initializing stores to use and relevant MBeans.

See [Interface: ManagedPluginInstance](#).

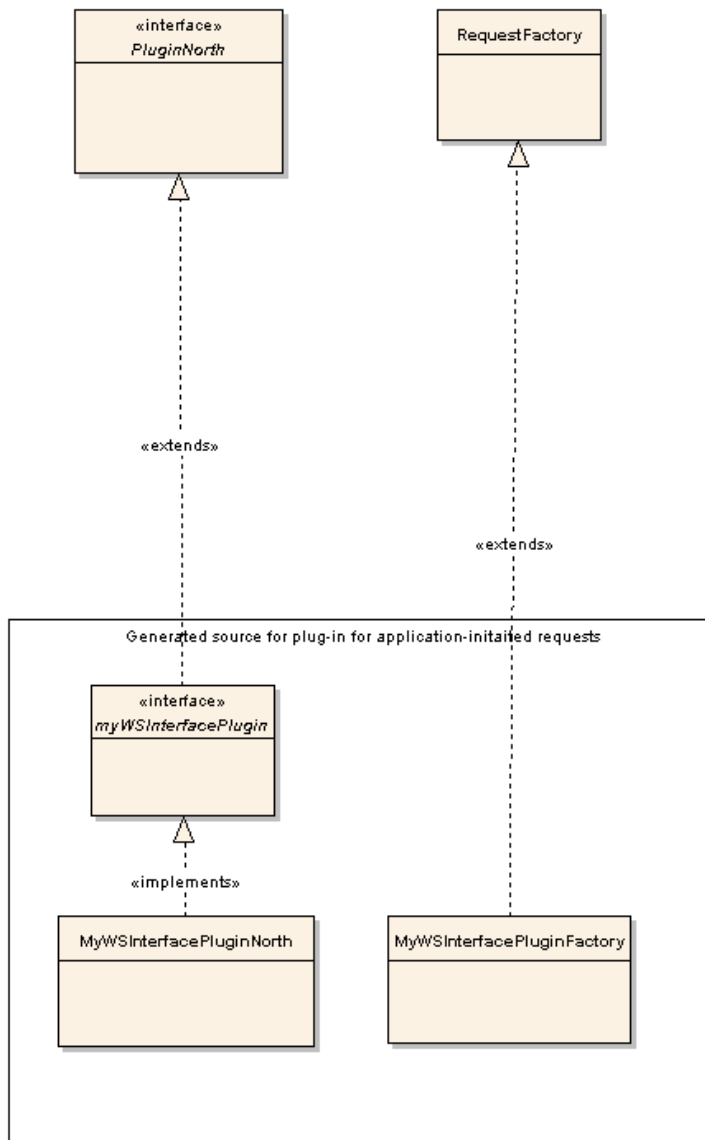
PluginNorth

This is an empty implementation of the Plug-in North interface. This interface is generated based on the WSDL files that define the application-facing interface. This is the starting point for the plug-in implementation.

The following files will be generated in the directory under `src/.../<service name>/north`:

- `<web service interface name>PluginNorth`: This class implements the plug-in interface. One file is generated for each plug-in interface. There is one plug-in interface for each service WSDL.

Figure 4-2 Class diagram of the generated PluginNorth and RequestFactory.



PluginNorth skeleton

Below is an outline on what needs to be implemented in the plug-in skeleton.

The class contains a Java mapping of the methods defined in the Web Service. The methods are mapped one-to-one. The name of each method is the same as the name of the operation defined in the WSDL. The parameter is a class that mirrors the parameters in the input message in the Web Service request. The return type is a class that represents the output message in the Web Service Request.

RequestFactory Skeleton

The actual class name is `<Communication service identifier>PluginFactory`, such as, `NotificationManagerPluginFactory`. This is a helper class used by the Service EJB. It serves two purposes:

- It creates the routing information requested by the Plug-in Manager when routing the method call to a plug-in.
- It converts exceptions thrown either by the Plug-in Manager or by the plug-in to exception types that are supported by the application-facing interface. This is the place to convert exceptions specific to an extension plug-in to exceptions specific to the application-facing interface. It is a *best practice* to have one single place for performing these conversions in order to document and locate exception mappings.

The following files will be generated in the `dist` directory under `request_factory_skel/src`:

- `<webservice_interface_name>PluginFactory`: This class extends the `RequestFactory` class. There will be one file generated for each plug-in interface.

Generated classes for a SOAP2SOAP Plug-in

In addition to the generated classes for a regular plug-in, a SOAP2SOAP plug-in adds a few extra classes, because the network protocol is known.

Note: Only the deployable artifacts are relevant. The generated code for SOAP2SOAP type of plug-ins should not be modified.

See [Managing and Configuring SOAP2SOAP Communication Services](#) in the *System Administrator's Guide* for information on how to configure and manage a SOAP2SOAP plug-in

Comparison with a Non-SOAP2SOAP Plug-in

The following generated code is similar to the code generated for the non-SOAP2SOAP plug-ins:

- [Interface: ManagedPluginService](#)
- [Interface: PluginService](#)
- [Interface: PluginInstanceFactory](#)
- [Interface: PluginServiceLifecycle](#)
- [ManagedPlugin Skeleton](#)
- [RequestFactory Skeleton](#)

Client Stubs

These classes and interfaces are generated for each interface, based on the Service WSDLs:

- [<Web Services Interface>_Stub](#)
- [<Web Services Interface>](#)
- [<Web Services Interface>Service_Impl](#)
- [<Web Services Interface>Service](#)

<Web Services Interface>_Stub

Class.

Extends `weblogic.wsee.jaxrpc.StubImp`

Implements [<Web Services Interface>](#)

Used by the corresponding `PluginNorth` class.

<Web Services Interface>

Interface.

Extends `java.rmi.Remote`.

Implemented by [<Web Services Interface>_Stub](#).

Defines the traffic methods.

<Web Services Interface>Service_Impl

Class.

Extends `weblogic.wsee.jaxrpc.ServiceImpl`.

Implements the Web Service.

<Web Services Interface>Service

Interface.

Extends `javax.xml.rpc.Service`.

Defines the traffic interfaces.

PluginInstance

In addition to the functionality in described in [PluginInstance](#), in the PluginInstance generated for SOAP2SOAP plug-ins, the following occurs:

- In the implementation of `activate()` it:
 - instantiates and registers a class implementing `com.bea.wlcp.wlng.httpproxy.management.HTTPProxyManagement`
 - instantiates and registers a a class implementing `com.bea.wlcp.wlng.heartbeat.management.HeartbeatManagement`
- It unregisters the above in the implementation of `deactivate()`.
- In the implementation of `isConnected()`, `HeartbeatManagement` is used to check the connection towards the network node.
- `getHttpProxyManagement()` is added for use by [PluginSouth](#).

`HTTPProxyManagement` is described in section [Managing and Configuring SOAP2SOAP Communication Services](#) in *Network Gatekeeper System Administrator's Guide*.

`HeartbeatManagement` is described in section [Configuring Heartbeats](#) in *Network Gatekeeper System Administrator's Guide*.

PluginNorth

In addition to the functionality described in [PluginNorth](#), this class:

- Checks whether there is an endpoint to the network node registered in the `HttpProxyManagement` MBean.
- Instantiates the client stubs used to make Web Services call to the network node: see [Client Stubs](#).
- Invokes the corresponding method on the stubs.

PluginSouth

This class implements a Java representation of the Web Service implementation. It implements `PluginSouth`: see [Interface: PluginSouth](#). When a network-triggered method is invoked, it:

- gets the handle to the callback EJB, see [Class: CallbackFactory](#).
- Resolves the endpoint used for the application instance by querying the [PluginInstance](#) for the endpoint by calling `getApplicationEndPoint(getApplicationInstanceId)`.
- Passes on the request to the callback EJB.

RequestFactory

The `RequestFactory` for a SOAP2SOAP plug-in has the same functionality as described in [RequestFactory Skeleton](#), but instead of serving as a skeleton, it is an implementation. It contains an implementation of `createRequestInfo(...)` which means that the Plug-in Manager does no routing based on destination address.

Build Files and Targets for a Communication Service Project

Main Build File

The main build file for the Communication Service contains the following targets:

- `build_csc`, builds the common parts of the Communication Service .
- `build_plugins`, builds the plug-ins for the Communicaiton Service .
- `stage`, copies the JARs for the plug-ins to the directory `stage`.
- `make-facade`, creates a deployable EAR for the access tier in the directory `dist`.

- `make-enabler`, creates a deployable EAR for the network tier in the directory `dist`.
- `deploy-facade`, deploys the service facade EAR to the access tier.
- `undeploy-facade`, undeploys the service facade EAR from the access tier.
- `deploy-enabler`, deploys the service enabler EAR from the network tier.
- `undeploy-enabler`, undeploys the service enabler EAR from the network tier
- `clean`, clears the directory `dist`.
- `dist`, calls the `prepare`, `build_csc`, `build_plugins`, `stage`, `make-facade`, `make-enabler` targets.

Note: When using the `deploy` and `undeploy` targets, make sure to adapt the settings for `user`, `password`, `adminurl`, `targets`, and `appversion` in the parameters to `wldeploy`. By default Web Services Security is not enabled for new Communication Services. See section [Setting up WS-Policy and JMX Policy](#) in *System Administrator's Guide* for instructions on how to configure this.

Communication Service Common Build File

The build file for the common parts of the Communication Service contains the following targets:

- `dist`, Calls the `csc_gen` ant task that generates the Java source for each `PluginFactory`. The source is generated under the directory `dist/request_factory_skel/src`
- `clean`: Deletes the `dist` directory.

Plug-in Build File

The build file for the plug-in contains the following targets:

- `compile`, compiles the source code under the `src` directory and puts the class files under the `build` directory.
- `jar`, calls the `compile` target and then creates a plug-in jar file under the `dist` directory.
- `instrument`, weaves the aspects that should apply into the plug-in.
- `build.schema`, builds the schema file and the classes used by the storage service.
- `dist`, calls the `clean`, `compile`, `jar` and `instrument`, and `build.schema` targets.
- `clean`, deletes the `build` and `dist` directories.

Ant Tasks

The build files use a set of ant tasks and macros, described below:

- [cs_gen](#)
- [plugin_gen](#)
- [cs_package](#)
- [javadoc2annotation](#)

The ant tasks are defined in `$PDS_HOME/lib/wlng/ant-tasks.jar`

cs_gen

This ant task builds the common parts of the Communication Service. Below is a description of the attributes.

Table 4-1 cs_gen ant task

Attribute	Description
destDir	Defines the destination directory for the generated files.
packageName	Define the package name to be used. Example: com.mycompany.service
serviceType	Defines the service type. Used in EDRs, statistics, etc. Example: SmsServiceType, MultimediaMessagingServiceType.
company	Defines the company name, to be used in META-INF/MANIFEST.MF.
version	Defines the version, to be used in META-INF/MANIFEST.MF.
contextPath	Defines the context path for the Web Service. Example: myService
soapAttachmentSupport	Use <code>true</code> if SOAP with attachments shall be supported. Use <code>false</code> if not.
wlngHome	Path to \$WLNG_HOME, this depends on the installation. Example: <code>c:/bea/wlng400</code>

Table 4-1 cs_gen ant task

Attribute	Description
pdsHome	Path to \$PDS_Home, this depends on the installation. Example c:/bea/wlmg_pds400
classpath	Defines the necessary classpaths. Must include: \$WLNG_HOME/server/lib/weblogic.jar \$WLNG_HOME/server/lib/webservices.jar \$WLNG_HOME/server/lib/api.jar \$PDS_HOME/lib/wlmg/wlmg.jar \$PDS_HOME/lib/log4j/log4j.jar
servicewSDL	URL to the WSDL that defines the service.

Example:

```
<cs_gen destDir="${dist.dir}"
    packageName="com.bea.wlcp.wlmg.example"
    name="say_hello"
    serviceType="example"
    company="BEA"
    version="4.0"
    contextPath="sayHello"
    soapAttachmentSupport="false"
    wlmgHome="${wlmg.home}"
    pdsHome="${pds.home}">
    <classpath refid="wls.classpath"/>
    <classpath refid="wlmg.classpath"/>
    <servicewSDL file="${wsdl}/example_hello_say_service.wsdl"/>
</cs_gen>
```

plugin_gen

This ant task builds a plug-in for a Communication Service. Below is a description of the attributes.

Table 4-2 plugin_gen ant task

Attribute	Description
destDir	Defines the destination directory for the generated files.
packageName	Define the package name to be used. Example: com.mycompany.service
name	Name and directory of the plug-in JAR.
serviceType	Defines the service type. Used in EDRs, statistics, etc. Example: SmsServiceType, MultimediaMessagingServiceType.
esPackageName	Communication Service package name used to import relevant classes.
protocol	An identifier for the network protocol the plug-in implements. Used as a part of the names of the generated jar file: <Communication Service identifier>_<protocol>.jar and the service name Plugin_<Communication Service identifier>_<protocol>.
schemes	Address schemes the plug-in can handle. Use a comma separated list if multiple schemes are supported. For example: tel: or sip:.
company	Defines the company name, to be used in META-INF/MANIFEST.MF.
version	Defines the version, to be used in META-INF/MANIFEST.MF.
pluginifjar	The name of the JAR file for the plug-in.

Table 4-2 plugin_gen ant task

Attribute	Description
classpath	Defines the necessary classpaths. Must include: \$WLNG_HOME/server/lib/weblogic.jar \$WLNG_HOME/server/lib/webservices.jar \$WLNG_HOME/server/lib/api.jar \$PDS_HOME/lib/wlng/wlng.jar \$PDS_HOME/lib/log4j/log4j.jar
servicewSDL	URL to the WSDL that defines the service.

Example:

```
<plugin_gen destDir="${dist.dir}"
    packageName="com.bea.wlcp.wlng.example.bla"
    name="say_hello"
    serviceType="example"
    esPackageName="com.bea.wlcp.wlng.example"
    protocol="bla"
    schemes=" "
    company="BEA"
    version="4.0"
    pluginifjar="${dist.dir}/say_hello/common/dist/say_hello_service.jar">
    <classpath refid="wls.classpath"/>
    <classpath refid="wlng.classpath"/>
    <servicewSDL file="${wSDL}/example_hello_say_service.wSDL"/>
</plugin_gen>
```

cs_package

This ant task packages a Communication Service. Below is a description of the attributes.

Table 4-3 cs_package ant task

Attribute	Description
destfile	Defines the destination directory for the generated files.
duplicate	Defines the package name to be used. Example: com.mycompany.service
displayname	Used in applicaiton.xml for the display name of the application.
descriptorfileset	Defines the service type. Used in EDRs, statistics, etc. Example: SmsServiceType, MultimediaMessagingServiceType.
manifest	Description of the manifest file use. Enter values for the following attributes: name="Bundle-Name" value should be the name of the EAR for the service enabler. name="Bundle-Version" value should be the version to use. name="Bundle-Vendor" value should be vendor name name="Weblogic-Application-Version" value should be the version of the EAR
fileset	Should point to the Communication Service JAR.
zipfileset	Should point to the plug-in JAR(s).

Example:

```
<cs_package destfile="${cs.dist}/${enabler.ear.name}.ear"
  duplicate = "fail"
  displayname="${enabler.ear.name}">
  <descriptorfileset dir="${csc.dir}/resources/enabler/META-INF"
    includes="*.xml"/>
  <descriptorfileset dir="${cs.name}/plugins"
    includes="*/resources/META-INF/*.xml"/>
  <manifest>
```



```

<attribute name="Bundle-Name"
    value="${enabler.ear.name}" />
<attribute name="Bundle-Version"
    value="${manifest.bundle.version}" />
<attribute name="Bundle-Vendor"
    value="${manifest.bundle.vendor}" />
<attribute name="Weblogic-Application-Version"
    value="${manifest.bundle.version}" />
</manifest>
<fileset dir="${csc.dist}">
    <include name="*_service.jar" />
</fileset>
<zipfileset dir="${cs.stage}">
    <include name="*plugin.jar" />
</zipfileset>
</cs_package>

```

javadoc2annotation

This ant macro annotates an MBean interface based on the JavaDoc. The macro is defined in the common.xml build file for the

The annotations are rendered as descriptive information by the Gatekeeper Administration console. Below is a description of the attributes.

Table 4-4 javadoc2annotation ant macro

Attribute	Description
tempDir	Temporary directory for the generated files.
destDir	Destination directory for the generated MBean interface.

Table 4-4 javadoc2annotation ant macro

Attribute	Description
sourceDir	Source directory for the MBean interface with JavaDoc annotations.
classpath	Defines the necessary classpaths. Depending on which interfaces that are used from the MBean, include: \$WLNG_HOME/server/lib/weblogic.jar \$WLNG_HOME/server/lib/webservices.jar \$WLNG_HOME/server/lib/api.jar \$PDS_HOME/lib/wlng/wlng.jar \$PDS_HOME/lib/log4j/log4j.jar

Example:

```
<javadoc2annotation  
    tempDir="${plugin.generated.dir}/mbean_gen_tmpdir"  
    destDir="${plugin.classes.dir}"  
    sourceDir="${plugin.src.dir}"  
    classpath="javadoc.classpath">  
</javadoc2annotation>
```

Communication Service Example

This section describes the example Communication Service in the Platform Development Studio:

- [Overview](#)
 - [High-level Flow for sendData \(Flow A\)](#)
 - [High-level Flow for startNotification and stopNotification \(Flow B\)](#)
 - [High-level flow for notifyDataReception \(Flow C\)](#)
- [Interfaces](#)
 - [Web Service Interface Definition](#)
 - [Network Interface Definition](#)
- [Directory Structure](#)
- [Classes](#)
- [Store configuration](#)
- [SLA Example](#)

Overview

The Communication service example demonstrates the following:

- Structure and execution workflow in a Communication Service.
- Parameter validation

Communication Service Example

- Hitless upgrade
- Retry
- Simple TCP/IP protocol-based simulator
- Testability with the PTE

The example is based on an end-to-end Communication Service, with a set of simple interfaces

- `SendData`, which defines the operation `sendData` used to send data to a given address.
- `NotificationManager`, which defines these operations:
 - `startEventNotification`, that starts a subscription for network-triggered events.
 - `stopEventNotification`, that ends the subscription for network-triggered events.
- `Notification`, which defines the operation:
 - `notifyDataReception`, used to notify the application on a network-triggered event.

The `SendData` and `NotificationManager` interfaces are used by an application and implemented by the Communication Service.

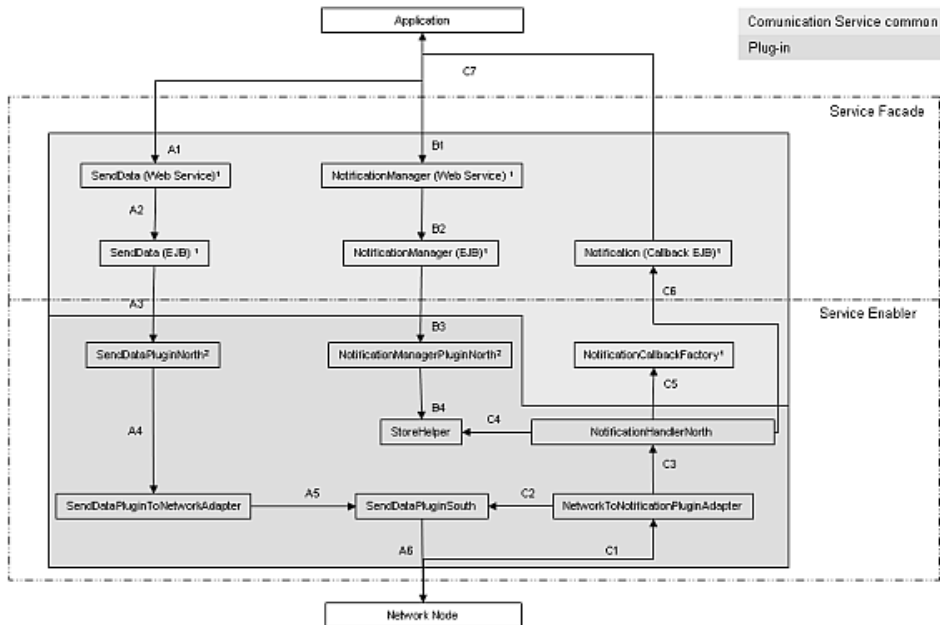
The `Notification` interface is used by the Communication Service and implemented by an application.

The Communication Service to network node interface is a simple TCP/IP based interface that defines the two commands:

- `sendDataToNetwork`, that sends data to the network node.
- `receiveData`, that is used by the network node to send data to a receiver - in this case the network protocol plug-in.

[Figure 5-1](#) illustrates the flow for these operations.

Figure 5-1 Overview of example Communication Service



The flow marked A* is for `sendData`, the flow marked B* is for `startNotification` and `stopNotification`, and the flow marked C* is for `notifyDataReception`.

The modules marked with 1 are automatically generated based on the WSDL files that defines the application-facing interface and code generation templates provided by the Platform Development Studio. The modules marked with 2 are skeletons generated at build time.

High-level Flow for `sendData` (Flow A)

1. A1: An application invokes the Web Service `SendData`, with the operation `sendData`.
2. A2: The request is passed on the EJB for the interface, which passes it on to the network protocol plug-in. The diagram is simplified, but at this stage the Plug-in Manager is invoked and makes a routing decision to route to the appropriate plug-in.

3. A3: The Plug-in Manager invokes the `sendData` method in the class `SendDataPluginNorth`. It will always invoke a class named `PluginNorth`, that has a prefix that is the same as the Java representation of the Web Service interface.
4. A4: The request is passed on to class `SendDataPluginToNetworkAdapter` that performs the protocol translation according to the network-interface.
5. A5: The request is passed to `SendDataPluginSouth`.
6. A6: The request is handed off to the network node.

High-level Flow for `startNotification` and `stopNotification` (Flow B)

The initial steps (B1-B3) are similar to flow A*. Instead of translating the request to a command on the network node, `NotificationManagerNorth` uses the `StoreHelper` to either store a new or remove a previously registered subscription for notifications. The data stored, the `NotificationData`, is used in network-triggered requests to resolve which application started the notification and the destination to which to send it. In the example the notification is started on an address, so the address is stored together with information to which endpoint the application wants the notification to be sent.

High-level flow for `notifyDataReception` (Flow C)

1. C1: The network protocol plug-in receives the network-triggered command `receiveData` on `NetworkToNotificationPluginAdapter`.
2. C2: `SendDataPluginSouth` can be used to add additional information to the request before passing in on.
3. C3: `NetworkToNotificationPluginAdapter` performs the protocol translation.
4. C4: `StoreHelper` is used to examine if the request matches any stored `NotificationData`. If so, the information in `NotificationData` is retrieved. This information includes which application instance that the request resolves to and on which endpoint this application wants to be notified about the network triggered event.
5. C5: `NotificationCallbackFactory` is used to get a hold of an active `NotificationCallback` EJB to pass on the request to.
6. C6: The request is passed on to the `NotificationCallback` EJB.
7. C7: The request is passed on to an application.

Interfaces

The example Communication Service translates between an application-facing interface, defined in WSDL, see [Web Service Interface Definition](#) and a network interface, TCP/IP based, see [Network Interface Definition](#).

Web Service Interface Definition

Interface: SendData

This interface is a simple interface containing operations for sending data.

Operation: sendData

Send data to the network.

Input message: sendDataMessage

Part name	Part type	Optional	Description
data	xsd:string	N	The data to be sent to the target device
address	xsd:anyURI	N	Address of the target device. Example: tel:4154011234

Output message: sendDataResponse

Part name	Part type	Optional	Description
none			

Interface: NotificationManager

The Notification Manager Web Service is a simple interface containing operations for managing subscriptions to network triggered events.

Operation: startEventNotification

Start the subscription of event notification from the network.

Input message: startEventNotificationRequest

Part name	Part type	Optional	Description
correlator	xsd:string	N	Service unique identifier provided to set up this notification.
endPoint	xsd:string	N	Endpoint address. Endpoint of the application to receive notifications. Example: http://www.hostname.com/NotificationService/services/Notification
address	xsd:anyUR	N	Service activation number. Example: tel:4154567890

Output message: invokeMessageResponse

Part name	Part type	Optional	Description
none			

Operation: stopEventNotification

Stop the subscription of event notification from the network.

Input message: stopEventNotificationRequest

Part name	Part type	Optional	Description
correlator	xsd:string	N	Service unique identifier provided to set up this notification.

Output message: stopEventNotificationResponse

Part name	Part type	Optional	Description
none			

Interface: NotificationListener

The NotificationListener interface defines the methods that the Communication Service invokes on a Web Service that is implemented by an application.

Operation: notifyDataReception

Method used for receiving a notification.

Input message: notifyDataReceptionRequest

Part name	Part type	Optional	Description
correlator	xsd:string	N	Service unique identifier provided to set up this notification.
originatingAddress	xsd:anyURI	N	Address of the device where the data originated. Example: tel:4153083412
data	xsd:string		Data sent by the originating device.

Output message: notifyDataReceptionResponse

Part name	Part type	Optional	Description
none			

Network Interface Definition

sendDataToNetwork

This command sends data from the Communication Service to the network node.

Argument	Type	Description
fromAddress	String	The address from which the request is sent.
toAddress	String	The address to which the request shall be sent.
data	String	The data to send.

receiveData

This command sends data from the network node to the Communication Service.

Argument	Type	Description
fromAddress	String	The address from which the request is sent.
toAddress	String	The address to which the request shall be sent.
data	String	The data to send.

Directory Structure

Below is a description of the directory structure for the example Communication Service.

```

communication_service
+- build.properties
+- common.xml
+- build.xml
+- example
| +- common
| | +- build.xml
| | +- dist
| | | +- request_factory_skel
| | | +- tmp
| | | +- example.war
| | | +- example_callback.jar
| | | +- example_callback_client.jar
| | | +- example_service.jar
| | | +- resources
| | | | +- enabler
| | | | + facade
| | | +- src
| | | | +- com/<package name>Plugin
| | | | | +- ExceptionType.java
| | | | | +- NotificationManagerPluginFactory.java
| | | | | +- SendDataPluginFactory.java
| | | | | +- handlerconfig.xml
| | | | | +- weblogic.xml
| +- wsdl

```

Communication Service Example

```
| +- dist
| | +- com.acompany.plugin.example.netex.store_4.0.jar
| | +- example_enabler.ear
| | +- example_facade.ear
| +- plugins
| | +- nextex
| | | +- build.xml
| | | +- dist
| | | | +- example_netex_plugin.jar
| | | | +- com.acompany.plugin.example.nextex.store_4.0.0.0.jar
| | | +- build
| | | +- config
| | | | +- edr
| | | | | +- alarm.xml
| | | | | +- cdr.xml
| | | | | +- edr.xml
| | | | | +- alarm.xml
| | | +- instance_factory
| | | | +- instancemap
| | +- dist
| | | +- com.acompany.plugin.example.netex.store_4.0.jar
| | | +- example_netex_plugin.jar
| | +- src/com/acompany/plugin/example/netex/
| | | | +- context
| | | | +- management
| | | | +- notification
| | | | +- notificationmanager
```

```

| | | | +- senddata
| | | | +- store
| | | +- storage
| | | +- wlng-cachestore-config-extensions.xml

```

Directories for WSDL

Below is a list of WSDL files that define the application-facing interface and the Java representation of these in the plug-in.

Application-initiated traffic

\$PDS_HOME/example/communication_service/example/common/wsdl/service

```

example_common_faults.wsdl
example_common_types.xsd
example_data_send_interface.wsdl
example_data_send_service.wsdl
example_notification_manager_interface.wsdl
example_notification_manager_service.wsdl

```

Network-triggered traffic

\$PDS_HOME/example/communication_service/example/common/wsdl/callback

```

example_notification_interface.wsdl
example_notification_service.wsdl

```

Directories for Java Source

Below is a list of Java source directories for the [Communication Service Common](#) and the [Plug-in](#).

Communication Service Common

\$PDS_HOME/example/communication_service/example/common/src

```
com.acompany.example.plugin.ExceptionType  
com.acompany.example.plugin.NotificationManagerPluginFactory  
com.acompany.example.plugin.SendDataPluginFactory
```

Plug-in

\$PDS_HOME/example/communication_service/example/plugins/netex/src

```
com.acompany.plugin.example.netex.context.ContextTranslatorImpl  
com.acompany.plugin.example.netex.management.ConfigurationStoreHandler  
com.acompany.plugin.example.netex.management.ExampleMBean  
com.acompany.plugin.example.netex.management.ExampleMBeanImpl  
com.acompany.plugin.example.netex.management.Management  
com.acompany.plugin.example.netex.notification.north.NotificationHandlerNorth  
com.acompany.plugin.example.netex.notification.south.NetworkToNotificationPluginAdapter  
com.acompany.plugin.example.netex.notification.south.NetworkToNotificationPluginAdapterImpl  
com.acompany.plugin.example.netex.notificationmanager.north.NotificationManagerPluginNorth  
com.acompany.plugin.example.netex.senddata.north.SendDataPluginNorth  
com.acompany.plugin.example.netex.senddata.south.SendDataPluginSouth  
com.acompany.plugin.example.netex.senddata.south.SendDataPluginToNetworkAdapter  
com.acompany.plugin.example.netex.senddata.south.SendDataPluginToNetworkAdapterImpl  
com.acompany.plugin.example.netex.store.FilterImpl  
com.acompany.plugin.example.netex.store.NotificationData
```

```
com.acompany.plugin.example.netex.store.StoreHelper
com.acompany.plugin.example.netex.ExamplePluginInstance
com.acompany.plugin.example.netex.ExamplePluginService
```

Directories for resources

Only the Communication Service common components have associated resources. The resources are XML files that serve as deployment descriptors for the network tier EAR and the access tier EAR.

\$PDS_HOME/example/communication_service/example/common/resources/at/META-INF

Contains deployment descriptors for the access tier EAR file. These must be present in the META-INF directory of the EAR. See

http://edocs.bea.com/wls/docs100/programming/app_xml.html

```
application.xml
weblogic-application.xml
```

The code generation creates these files, and the build script takes care of the packaging.

\$PDS_HOME/example/communication_service/example/common/resources/nt/META-INF

Contains deployment descriptors for the network tier EAR file. These must be present in the META-INF directory of the EAR. See

http://edocs.bea.com/wls/docs100/programming/app_xml.html

```
application.xml
weblogic-application.xml
weblogic-extension.xml
```

The code generation creates these files, and the build script takes care of the packaging.

Directories for Configuration of Plug-in

\$PDS_HOME/example/communication_service/example/plugins/netex/config/edr

Sample entries to add in the EDR, CDR, and Alarm filters.

Communication Service Example

alarm.xml

cdr.xml

edr.xml

These serves as examples. Add the contents of these to the EDR configuration file. Use the **EDR Configuration Pane** as described in [Managing and Configuring EDRs, CDRs and Alarms](#) in the *System Administrator's Guide*.

\$PDS_HOME/example/communication_service/example/plugins/netex/instance_factory

Sample instance map for mapping of classes, interfaces, and abstract classes.

When using `com.bea.wlcp.wlng.api.util.InstanceFactory` to retrieve instances for a given interface, class, or abstract class, this mapping is referenced. The mapping can be overridden. See JavaDoc for `InstanceFactory` for details.

instancemap

\$PDS_HOME/example/communication_service/example/plugins/netex/storage

Sample store configuration file. Defines how the Storage service is used by the plug-in, store type, table names, query definitions, and get and set methods. See [StoreHelper](#), [FilterImpl](#), and [NotificationData](#).

wlng-cachestore-config-extensions.xml

Directories for Build and Configuration of Builds

\$PDS_HOME/example/communication_service/

build.properties

Defines the installation directory for Network Gatekeeper and for the Platform Development Studio.

common.xml

Defines properties, class paths, task definitions, and macros for the build.

build.xml

Main build file to build the Communication Service. This build file also contains targets for packaging deployable artifacts into the access and network tier.

\$PDS_HOME/example/communication_service/example/common

build.xml

Build file for the common parts of the Communication Service.

\$PDS_HOME/example/communication_service/example/plugins/netex

build.xml

Build file for the plug-in.

Directories for Classes, JAR, and EAR Files

\$PDS_HOME/example/communication_service/example/dist

Deployment artefacts for the Communication Service.

example_facade.ear

The part of the Communication Service that is deployed in the access tier.

example_enabler.ear

The part of the Communications Service that is deployed in the network tier.

\$PDS_HOME/example/communication_service/example/common/dist

JAR and WAR files for the common parts of the Communication Service.

example_callback_client.jar

example_callback.jar

example_service.jar

example.war

\$PDS_HOME/example/communication_service/example/common/dist/request_factory_skel

Auto generated source for skeleton classes extending
com.bea.wlcp.wlng.api.plugin.RequestFactory.

One class is generated per Service WSDL, that is per interface that defines application-initiated operations.

The classes are named <PreFix>PluginFactory, where <PreFix> is picked up from the WSDL binding in the WSDL file.

In the subdirectory that corresponds to the package name, the following classes are generated:

`NotificationManagerPluginFactory.java`

`SendDataPluginFactory.java`

These are generated as skeletons, but in the example they are adapted to the specific use cases.

\$PDS_HOME/example/communication_service/example/plugins/netex/dist

Contains individual JAR files comprises the plug-in.

`com.acompany.plugin.example.netex.store_4.0.jar`

Includes the schema file for the store used by the plug-in, packaged together with the classes for which instances are stored. This file must be put in `$DOMAIN_HOME/config/store_schema` on each server in the network tier. The server needs to be restarted if any changes have been done to the store schema or the classes referred to in the store schema.

`example_netex_plugin.jar`

The JAR for the plug-in.

\$PDS_HOME/example/communication_service/example/plugins/netex/dist/mbean_generationdir

Output directory for the MBean that has been processed by the `javadoc2annotation` ant task.

Classes

Below is a description of the classes and the methods defined in these classes:

- **Communication Service Common**
 - `ExceptionType`
 - `NotificationManagerPluginFactory`
 - `SendDataPluginFactory`
- **Plug-in Layer**
 - `ContextTranslatorImpl`
 - `ExamplePluginService`
 - `ConfigurationStoreHandler`
 - `ExampleMBean`

- ExampleMBeanImpl
- Management
- NotificationHandlerNorth
- NetworkToNotificationPluginAdapter
- NetworkToNotificationPluginAdapterImpl
- NotificationManagerPluginNorth
- SendDataPluginNorth
- SendDataPluginSouth
- SendDataPluginToNetworkAdapter
- SendDataPluginToNetworkAdapterImpl
- FilterImpl
- NotificationData
- StoreHelper

Communication Service Common

ExceptionType

Class.

Enumeration for exception types:

Defines:

- SERVICE_ERROR
- POLICY_ERROR

NotificationManagerPluginFactory

Class.

Extends RequestFactory.

Helper class that is used by the service EJB for two purposes:

- Creating routing information requested by the Plug-in Manager when routing the method call to a plug-in.

- Converting Exceptions, thrown either by the Plug-in Manager or by the plug-in, to Exceptions that are supported by the application-facing interface.

Note: This class needs to remain in this package and the class name must not be changed.

public void validateRequest(Method method, Object... args)

Validates the request to make sure that mandatory parameters are present. Operates on a Java representation of the Web Service call.

public RequestInfo createRequestInfo(Class<? extends Plugin> type, Method method, Object... args)

Used by the service EJB to extract routing data from the method call. The routing data is then given to the Plug-in Manager. This method returns the routing data in a RequestInfo object.

Returns a:

- AddressRequestInfo if the request contains an actual address that can be routed to a specific plug-in.
- CorrelatorRequestInfo if the request contains an correlator that relates to an operation that relates to states (to start or to stop something). Most often it is the starting and stopping of notifications that use a correlator.

public Throwable convertEx(Method method, Throwable e)

Called by the service EJB in order to convert Exceptions thrown by the Plug-in Manager and the Plug-in to Exceptions defined by the called method.

private Throwable convertEx(Method method, PluginException e)

Converts a PluginException to an Exception that can be thrown by the method called by the application.

Plug-in Layer

ContextTranslatorImpl

Class.

Implements interface com.bea.wlcp.wlng.api.plugin.context.ContextTranslator.

Responsible for setting any non-simple parameter into the RequestContext.

public void translate(Object param, ContextInfo info)

Puts the member variables of a complex data type into the ContextInfo.

Checks the interface type.

Gets the simple data types provided in the parameter param.

Puts each of the parameters into the ContextInfo object.

These parameters are provided in each subsequent EDR that is emitted in the request.

ExamplePluginService

Package: com.acompany.plugin.example.netex

Implements ManagedPluginService.

Initial point for the network protocol plug-in.

Defines the life-cycle for a plug-in service.

Also holds the data that is specific for the plug-in instance.

This class manages the life-cycle for the plug-in service, including implementing the necessary interfaces that make the plug-in deployable in Network Gatekeeper. It is also responsible for registering the north interfaces with the Plug-in Manager. At startup time it uses the InstanceFactory to create one instance of each plug-in service and at activation time it registers these with the Plug-in Manager. The InstanceFactory uses an instancemap to find out which class it should instantiate for each plug-in interface implementation. The instance map is found under the resource directory. It also has

public boolean isRunning()

Checks to see if the plug-in service is in running state.

public String[] getSupportedSchemes()

Returns a list of address schemes the plug-in supports.

public void init(String id, PluginPool pool)

Initializes the plug-in service with its ID and a reference to its plug-in pool.

public void doStarted()

When entering state Started, the plug-in instantiates a TimerManager.

public void doStopped()

No action.

public void doActivated()

No action.

public void doDeactivated()

No action.

public void handleSuspending(CompletionBarrier barrier)

The plug-in service does not handle graceful shutdown: it propagates the request to [public void handleForceSuspending\(\)](#).

public void handleForceSuspending()

When the plug-in is being forcefully suspended, the plug-in service iterates through all plug-in instances and calls [public void handleSuspending\(\)](#) on each.

public boolean isActive()

While there is a connection to the network node and the plug-in is in state ACTIVE/RUNNING this method must return true, in all other cases false. This method is invoked by the Plug-in Manager during route selection.

public ServiceType getServiceType()

Returns the type of the service. Used by the Plug-in Manager to route requests to a plug-in instance that can manage the type of request. The ServiceType is auto-generated based on the WSDL that defines the application-facing interfaces.

public String getNetworkProtocol()

Returns a descriptive name of the network protocol being used.

createInstance(String)

Creates a new plug-in instance.

ExamplePluginInstance

Package: com.acompany.plugin.example.netex.

Implements `ManagedPluginInstance`

Defines the life-cycle for a plug-in instance/

This class manages the life-cycle for the plug-in instance including implementing the necessary interfaces that make the plug-in an instance in Network Gatekeeper.

It is also responsible for instantiating classes that implement the traffic interfaces and for initializing stores to use and MBeans.

public String getId()

Returns the plug-in instance ID.

public void activate()

- Instantiates the classes implementing the `PluginNorth` interface:
 - [SendDataPluginNorth](#)
 - [NotificationManagerPluginNorth](#)
 - [NotificationHandlerNorth](#)
- Instantiates the class implementing the `PluginSouth` interface:
 - [SendDataPluginSouth](#)
- Instantiates the classes that implements the southbound and northbound adapter instances:
 - [NetworkToNotificationPluginAdapterImpl](#)
 - [SendDataPluginToNetworkAdapterImpl](#)
- Creates the network proxy:
- Registers the `PluginNorth` interfaces into the Plug-in Manager.
- Registers the `PluginSouth` interfaces into the Plug-in Manager.
- Registers the [NetworkToNotificationPluginAdapter](#) into the network proxy to be notified when a request arrives from the network node.
- Sets [NotificationHandlerNorth](#) to [NetworkToNotificationPluginAdapter](#) in order to forward request to the application.
- Sets the network proxy into the [SendDataPluginToNetworkAdapter](#) in order to send request to the network.

- Sets [SendDataPluginToNetworkAdapter](#) into [SendDataPluginNorth](#).
- Instantiates [ConfigurationStoreHandler](#).
- Instantiates [Management](#) and registers the plug-in into it.

private void rethrowServiceDeploymentException(Exception e)

Re-throws a ServiceDeploymentException if any other exception is encountered. The exception is wrapped in a ServiceDeploymentException.

public ConfigurationStoreHandler getConfigurationStore()

Returns a handle to the ConfigurationStore used by the plug-in instance. The ConfigurationStore was initiated in [public void activate\(\)](#).

public NetworkProxy getNetworkProxy()

Returns handle to the NetworkProxy. The NetworkProxy was initiated in [public void activate\(\)](#).

public void connect()

Connects to the network using NetworkProxy.

ConnectTimerTask

Inner class of [ExamplePluginService](#).

Extends java.util.TimerTask.

It has one method, run(), that tries to connect to the network node, if not connected. This class is instantiated and scheduled as a java.util.Timer in [public void handleResuming\(\)](#).

ConfigurationStoreHandler

Handles storage of configuration data using the StorageService.

A set of default settings are defined as static final variables. These are used to populate the ConfigurationStore with default values the first time the plug-in is deployed.

Takes the plug-in ID as a parameter. The plug-in ID is the key in the ConfigurationStore.

Uses ConfigurationStoreFactory to get a handle to the ConfigurationStoreService and gets the local ConfigurationStore that handles configuration data for the plug-in instance.

The plug-in only deals with configuration data that is unique for the instance in a specific server, so the store is fetched as outlined in [Listing 5-1](#).

Listing 5-1 Get a server-specific (local) ConfigurationStore

```
ConfigurationStoreFactory factory = ConfigurationStoreFactory.getInstance();
localConfigStore = factory.getStore(pluginId, LOCAL_STORE,
ConfigurationStore.STORE_TYPE_LOCAL);
```

If the plug-in uses a ConfigurationStore that is shared between the plug-in instances in the cluster, it must fetch that one as well, as outlined in [Listing 5-2](#)

Listing 5-2 Get a cluster-wide (shared) ConfigurationStore

```
ConfigurationStoreFactory factory = ConfigurationStoreFactory.getInstance();
sharedConfigStore = factory.getStore(pluginId, SHARED_STORE,
ConfigurationStore.ConfigurationStore.STORE_TYPE_SHARED);
```

After the ConfigurationStore is fetched, it is initialized with default values for the available configuration settings. These default values can be changed later on, using the MBeans, see [ExampleMBean](#).

```
public void setLocalInteger(String key, Integer value),
public Integer getLocalInteger(String key),
public void setLocalString(String key, String value), and
public String getLocalString(String key)
```

The methods above are used to set and get data to and from the ConfigurationStore. One set/get pair must be implemented per data type in the ConfigurationStore. It is only necessary to implement set/get methods for the data types actually used by the plug-in.

In the set methods, the parameter name/key is provided as the first parameter and the actual value is provided in the second parameter.

In the get methods, the parameter name/key is provided as the parameter and the actual value is returned.

ExampleMBean

Interface.

Management interface for the example simulator.

It defines the following methods:

- `public void setNetworkPort(int port) throws ManagementException;`
- `public int getNetworkPort() throws ManagementException;`
- `public void connect() throws ManagementException;`
- `public void disconnect() throws ManagementException;`
- `public boolean connected();`

Implemented by `ExampleMBeanImpl`.

All MBean methods should throw `com.bea.wlcp.wlng.api.management.ManagementException` or a subclass thereof if the management operation fails.

Management

Class.

Handles registration of the [ExampleMBean](#) in the MBean Server.

NotificationHandlerNorth

NotificationHandlerNorth()

Constructor.

Empty.

`public void deliver(String data, String destinationAddress, String originatingAddress)`

Delivers data originating from the network node to the application.

`NetworkToNotificationPluginAdapterImpl` calls this method upon a network triggered request.

The actual delivery is not done directly to the application. Instead it is done via the service callback client EJB which forwards the request to the service callback EJB. Both of these are generated during the build process.

First, the [NotificationData](#) associated with the destination address is fetched.

NotificationCallback, which is a generated class, is fetched using `private NotificationCallback getNotificationCallback()`.

NotifyDataReception, a generated class that is a Java representation of the operation defined in the callback WSDL is instantiated.

The correlator associated with the `NotificationData` is set on NotifyDataReception.

The data (payload) in the network triggered request is set on NotifyDataReception.

The originating address in the network-triggered request is converted to a URI and set on NotifyDataReception.

The endpoint associated with NotificationData is fetched.

A remote call is done to the method `notifyDataReception` on the Callback EJB in the access tier. The endpoint and NotifyDataReception are supplied as parameters.

private NotificationCallback getNotificationCallback()

Helper method to get the object representing the Callback EJB.

If the object is already retrieved it is returned, otherwise the NotificationCallbackFactory is used to get a new object. This is the preferred pattern.

Using the CallBackFactory ensures high-availability between the network tier and the access tier for network triggered requests.

The Callback is generated during the build process when the access tier is generated. Three files are generated per callback WSDL. The names are based on the interface name defined in the WSDL. The interface in the WSDL is *Notification*, so:

- the factory is named *NotificationCallbackFactory*.
- the implementation class is named *NotificationCallbackImpl*
- an interface is named *NotificationCallback*.

The classes are completely based on the WSDL file for the callback interface. The factory shall be used to retrieve the implementation class that implements the interface.

private NotificationData getNotificationData(String destinationAddress)

Helper method to fetch the NotificationData from the StoreHelper. The NotificationData is retrieved based on the key destination address.

NetworkToNotificationPluginAdapter

Interface

extends `PluginSouth`, `NetworkCallback`

Defines the interface between [NetworkToNotificationPluginAdapter](#) and the network node.

public void setNotificationHandler(NotificationHandlerNorth notificationHandlerNorth)

Sets the `NotificationHandler`.

NetworkToNotificationPluginAdapterImpl

Class.

Implements [NetworkToNotificationPluginAdapter](#).

public void setNotificationHandler(NotificationHandlerNorth notificationHandlerNorth)

Sets [NotificationHandlerNorth](#) in the class.

public String resolveAppInstanceId(ContextMapperInfo info)

From interface `com.bea.wlcp.wlng.api.plugin.PluginSouth`

Gives the plug-in an opportunity to add additional values to the `RequestContext` before the network-triggered requests is passed on to [public void](#)

[receiveData\(@ContextKey\(EdrConstants.FIELD_ORIGINATING_ADDRESS\) String fromAddress, @ContextKey\(EdrConstants.FIELD_DESTINATION_ADDRESS\) @MapperInfo\(C\) String toAddress, String data\).](#)

This method is called only once per network-triggered request. It is invoked after `resolveAppInstanceId(ContextMapperInfo)`, when the `RequestContext` for the current request has been rebuilt.

The default implementation is supposed to be empty.

`RequestContext` contains the fully rebuilt `RequestContext`.

`ContextMapperInfo` contains the annotated parameters in [public void receiveData\(@ContextKey\(EdrConstants.FIELD_ORIGINATING_ADDRESS\) String fromAddress, @ContextKey\(EdrConstants.FIELD_DESTINATION_ADDRESS\) @MapperInfo\(C\) String toAddress, String data\).](#)

```
public void receiveData(@ContextKey(EdrConstants.FIELD_ORIGINATING_ADDRESS)
String fromAddress, @ContextKey(EdrConstants.FIELD_DESTINATION_ADDRESS)
@MapperInfo(C) String toAddress, String data)
```

From NetworkCallback.

The network node invokes this method when a network-triggered events occurs.

The parameter:

- fromAddress is the address representing the originator of the request
- toAddress is the address representing the destination of the request.
- data contains the payload of the request.

The method is annotated with @Edr, so the method is woven with annotation EDR.

fromAddress and toAddress are annotated with @ContextKey, which means that they will be put in the current RequestContext under the key specified by the string in the argument of the annotation. As illustrated in [Listing 5-3](#), they are put in the RequestContext under the keys EdrConstants.FIELD_ORIGINATING_ADDRESS and EdrConstants.FIELD_DESTINATION_ADDRESS, respectively. These keys ensure that the values will be available in all subsequent EDRs emitted during this request.

toAddress is also annotated with @MapperInfo, which means that the value should be registered in ContextMapperInfo under the key specified by the string in the argument of the annotation. In [Listing 5-3](#), the key is C.

Listing 5-3 Annotation of network-triggered method

```
...
@Edr
public void receiveData(
    @ContextKey(EdrConstants.FIELD_ORIGINATING_ADDRESS)
    String fromAddress,
    @ContextKey(EdrConstants.FIELD_DESTINATION_ADDRESS)
    @MapperInfo(C)
    String toAddress,
```

```
String data) {  
...  

```

NotificationManagerPluginNorth

Class.

Implements NotificationManagerPlugin.

**public StartEventNotificationResponse
startEventNotification(@ContextTranslate(ContextTranslatorImpl.class)
StartEventNotification parameters)**

Starts a subscription for notifications on network-triggered requests.

The method is a Java representation of the application-facing operation startEventNotification, defined in the WSDL that was used as input for the code generation.

As illustrated in [Listing 5-4](#), the method is annotated with @EDR, and the parameter is put in the RequestContext using the annotation @ContextTranslate, since the parameter is a complex data type that requires traversal in order to resolve the simple data types. When using this annotation, the class is provided as an ID.

Listing 5-4 Annotations for startEventNotification

```
...  
@Edr  
public StartEventNotificationResponse startEventNotification(  
@ContextTranslate(ContextTranslatorImpl.class) StartEventNotification  
parameters)  
throws ServiceException {  
...  

```

In the operation, these parameters are included:

```
<xsd:element name="correlator" type="xsd:string"/>
```

```
<xsd:element name="endPoint" type="xsd:string"/>
<xsd:element name="address" type="xsd:anyURI"/>
```

The values of correlator and endPoint are put in NotificationData.

The application instance ID for the originator of the request, the application that uses the Web Services interface, is resolved from the RequestContextManager and put in NotificationData.

Using StoreHelper, NotificationData is put in the StorageService.

**public StopEventNotificationResponse
stopEventNotification(@ContextTranslate(ContextTranslatorImpl.class)
StopEventNotification parameters)stopEventNotification(StopEventNotification)**

Ends a previously started subscription for notifications on network-triggered requests.

The method is a Java representation of the application-facing operation `stopEventNotification`, defined in the WSDL that was used as input for the code generation.

The method is annotated in a similar manner to [public StartEventNotificationResponse startEventNotification\(@ContextTranslate\(ContextTranslatorImpl.class\) StartEventNotification parameters\)](#).

Using StoreHelper, NotificationData corresponding to the correlator provided in the requests is removed from the StorageService.

SendDataPluginNorth

Class.

Implements SendDataPlugin.

public void setPluginToNetworkAdapter(SendDataPluginToNetworkAdapter adapter)

Sets SendDataPluginToNetworkAdapter to be used for application-initiated requests.

**public SendDataResponse sendData(@ContextTranslate(ContextTranslatorImpl.class)
SendData parameters)**

Sends data to the network

The method is a Java representation of the application-facing operation `sendData`, defined in the WSDL that was used as input for the code generation.

The method is annotated in a similar manner to `public StartEventNotificationResponse startEventNotification(@ContextTranslate(ContextTranslatorImpl.class) StartEventNotification parameters)`.

Passes on the request to `SendDataPluginToNetworkAdapter`.

If there is a need to retry the request, this method re-throws a `PluginRetryException`, so the request can be retried by the service interceptors.

SendDataPluginSouth

Class.

implements `PluginSouth`.

public SendDataPluginSouth()

Constructor.

Empty.

public void send(NetworkProxy proxy, String address, String data)

Sends data to the network node.

Passes on the request to `sendDataToNetwork` using the `NetworkProxy`.

The method is annotated with `@Edr`.

public String resolveAppInstanceId(ContextMapperInfo info)

Empty implementation that returns null. This method has meaning, and is used, only in network-triggered requests.

The application instance ID is already known in the `RequestContext`, since the class only handles application-initiated requests.

public void prepareRequestContext(RequestContext ctx, ContextMapperInfo info)

From interface `com.bea.wlcp.wlng.api.plugin.PluginSouth`

Gives the plug-in an opportunity to add additional values to the `RequestContext` before the application-initiated requests is passed on to `public void send(NetworkProxy proxy, String address, String data)`.

Empty in this example. Normally all data about the request should be known at this point, so no additional data needs to be set.

SendDataPluginToNetworkAdapter

Interface.

Defines the interface between the plug-in and the network node for application-initiated requests.

SendDataPluginToNetworkAdapterImpl

Class.

public SendDataPluginToNetworkAdapterImpl()

Constructor.

Instantiates SendDataPluginSouth.

public void setNetworkProxy(NetworkProxy networkProxy)

Sets the NetworkProxy object. This is a remote object in the network node.

public void send(String address, String data)

Hands off the request to the network node using SendDataPluginSouth.

FilterImpl

Class.

Implements interface `com.bea.wlcp.wlng.api.storage.filter.Filter`.

This is the query filter used for the named store NotificationData.

Evaluates whether an entry in the named store NotificationData matches the filter. The filter is defined in XML, see [Store configuration](#).

public boolean matches(Object value)

Must be invoked after [public void setParameters\(Serializable ... parameters\)](#).

Returns true if the value provided in Object matches parameters[0], as set in [public void setParameters\(Serializable ... parameters\)](#).

public void setParameters(Serializable ... parameters)

Sets the query parameters for the filter.

The parameters are ordered as provided to the StoreQuery and it is the responsibility of the implementation to handle them in this order.

NotificationData

Class.

Implements Serializable

The data structure representing a notification. The notification is registered and de-registered by applications using the application-facing Web Services interfaces and represents a subscription for network-triggered events. The NotificationData is used for:

- Matching a network-triggered event with a subscription started by an application. The match is usually based on the destination address in the requests from the network.
- Resolving information on which application instance created the subscription, and the endpoint on which the application expects to be notified of the event.

NotificationData is stored using the storage service, normally using the invalidating cache storage provider for cluster-wide access and high performance.

Each of the attributes to be stored must have a corresponding set method and get method.

The class must be serializable.

public NotificationData()

Constructor.

Empty.

StoreHelper

Class.

Singleton.

Helper class for storing NotificationData using the StorageService.

public static StoreHelper getInstance()

Returns the single instance of StoreHelper.

public void addNotificationData(URI address, NotificationData notificationData)

Stores the NotificationData using the Storage Service.

The named store is retrieved using `private Store<String, NotificationData> getStore()`.

The NotificationData is put into the named store. The address is the key and the object is the value.

The named store is released. This should always be done in a finally{...} block.

public void removeNotificationData(String correlator)

Removes NotificationData using the StorageService.

The named store is retrieved using `private Store<String, NotificationData> getStore()`.

A Set of matching entries are returned using `private Set<Map.Entry<String, NotificationData>> getEntries(String correlator, Store<String, NotificationData> store)`.

If there are matching entries, all are removed using `private void removeEntries(Set<Map.Entry<String, NotificationData>> set, Store<String, NotificationData> store)`.

The named store is released. This should always be done in a finally{...} block.

public NotificationData getNotificationData(String destinationAddress)

Gets NotificationData using the StorageService

The named store is retrieved using `private Store<String, NotificationData> getStore()`.

The NotificationData that is keyed on destinationAddress is fetched from the store.

The named store is released. This should always be done in a finally{...} block.

private Store<String, NotificationData> getStore()

Gets a named stored from com.bea.wlcp.wlng.api.storage.StoreFactory.

private Set<Map.Entry<String, NotificationData>> getEntries(String correlator, Store<String, NotificationData> store)

Gets a java.util.Set of entries of NotificationData from a named store using the StorageService.

The query being used is a named query, com.bea.wlcp.wlng.plugin.example.netex.Query, defined in wlng-cachestore-config-extensions.xml.

private void removeEntries(Set<Map.Entry<String, NotificationData>> set, Store<String, NotificationData> store)

Removes a java.util.Set of entries of NotificationData using the StorageService. The NotificationData is removed from a named store.

ExamplePluginInstance

Class.

Implements `com.bea.wlcp.wlng.api.plugin.ManagedPluginInstance`.

Defines the life-cycle for a plug-in instance.

Also holds the data that is specific to the plug-in instance.

public ExamplePluginInstance(String id, ExamplePluginService parent)

Constructor.

The id is the plug-in instance ID, and the parent is the Plug-in service the of which the plug-in is an instance.

public String getId()

The plug-in instance returns the ID that it was instantiated with.

public void activate()

Called when the plug-in instance is activated, so the plug-in:

- Instantiates the traffic interfaces.
- Registers the traffic interfaces with the Plug-in Manager.
- Register callbacks between the interfaces.
- Initiates the Store.
- Instantiates and registers the MBean interface.

If the plug-in service is in state ACTIVE (RUNNING), `public void handleResuming()` is called.

public void handleResuming()

Connects to the network node.

If the connection fails, a timer is triggered to retry the connection setup.

public void deactivate()

Called when the plug-in instance is deactivated.

If the plug-in service is in state ACTIVE (RUNNING), `public void handleSuspending()` is called.

The call-back is unregistered from the network node.

The MBean is unregistered.

public void handleSuspending()

If existing, the timer associated with connection setup is cancelled.

The plug-in disconnects from the network node.

public List<PluginInterfaceHolder> getNorthInterfaces()/ public List<PluginInterfaceHolder> getSouthInterfaces()

Returns a list of the interfaces.

public boolean isConnected()

Returns true if there is a connection to the network node, that is if the plug-in instance is ready to accept traffic.

public int customMatch(RequestInfo requestInfo)

Checks which operation that is about to be invoked on the plug-in instance by introspection of the RequestInfo associated with request.

If the operation is StopEventNotification and the correlator provided is cached using the Storage service, the request must be sent to all instances of the plug-in, since the request depends on an earlier request (startNotification). MATCH_REQUIRED is returned.

If the operation is any other than StopEventNotification, the request is unrelated to any previous operation and any plug-in instance can be used. MATCH_OPTIONAL is returned.

private void rethrowDeploymentException(Exception e)

Re-throws a DeploymentException given another exception. The exception is wrapped in a DeploymentException.

public ConfigurationStoreHandler getConfigurationStore()

Gets the [ConfigurationStoreHandler](#).

ExamplePluginService

Class.

Implements com.bea.wlcp.wlng.api.plugin.ManagedPluginService.

Defines the life-cycle for a plug-in service.

Also holds the data that is specific for the plug-in instance.

public ExamplePluginService()

Constructor.

Empty.

public TimerManager getTimerManager()

Gets a handle to the TimerManager.

public boolean isRunning()

Checks if the plug-in service is in RUNNING state.

public String[] getSupportedSchemes()

Returns an array of supported address schemes.

public void init(String id, PluginPool pool)

Initializes the plug-in service with the ID and a reference to the plug-in pool.

The PluginPool holds all plug-in instances.

public void doStarted()

Instantiates a TimerManager to be used.

public void doStopped()/public void doActivated()/public void doDeactivated()

Empty implementation. Nothing to do here.

public void handleResuming()

Iterates over all plug-in instances using the PluginInstancePool and calls [public void handleResuming\(\)](#) on [ExamplePluginInstance](#)

public void handleSuspending(CompletionBarrier barrier)

The nature of the example network protocol is that it does not have connections to maintain.

Because it is possible to treat this event as in [public void handleForceSuspending \(\)](#) the request is passed on to that method.

public void handleForceSuspending ()

When the plug-in service is being forcefully suspended, the plug-in instances are disconnected from the network node immediately, without waiting for any in-flight requests to complete.

This is done by iterating over the `PluginInstancePool` and calling [public void handleSuspending\(\)](#) on [ExamplePluginInstance](#)

public ServiceType getServiceType()

Returns the service type, `com.acompany.example.servicetype.ExampleServiceType.type`. The type is automatically generated when the service EJB is generated.

public String getNetworkProtocol()

Returns the network protocol. A string used for informational purposes.

public ManagedPluginInstance createInstance(String pluginInstanceId)

Creates a new instance of the plug-in service. The ID for the new plug-in is supplied together with the object that created the instance (`this`).

Store configuration

The store configuration file `wlng-cachestore-config-extensions.xml` defines:

- Which data to store
- The get and set methods to retrieve and store the data
- The database table structure use to store the data
- Queries to perform on the store

[Listing 5-5](#) shows the store configuration file for the example Communication Service.

The configuration file defines:

- The store type ID: since the store type ID is prefixed with `wlng.db.wt` (`wlng.db.wt.es_example`), the store is a write-through cache.
- The table to be used: `es_example`
- The identifier for the store is a combination of the type of the key column (`java.lang.String`) and the type of the value column (`com.acompany.plugin.example.netex.store.NotificationData`). These are used when the

store is retrieved from the StoreFactory, see [private Store<String, NotificationData>getStore\(\)](#)

- The key column: address
- The value columns for the key:
 - correlator
 - endpoint
 - appinstance
- The get- and set methods for the value columns.
- The query to use when doing lookups in the store.

The configuration file, together with any non-complex data types must be packaged into a Jar and put in the directory `$DOMAIN_HOME/config/store_schema` so it can be accessed by the storage service.

Listing 5-5 Store configuration for the example Communication Service

```
<?xml version="1.0" encoding="UTF-8"?>

<store-config xmlns="http://www.bea.com/ns/wlng/30"
               xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
               xsi:schemaLocation="http://www.bea.com/ns/wlng/30
wlng-cachestore-config.xsd">

  <db_table name="es_example">

    <key_column name="address" data_type="VARCHAR(100)" />

    <value_column name="correlator" data_type="VARCHAR(100)">

      <methods>

        <get_method name="getCorrelator" />

        <set_method name="setCorrelator" />

      </methods>

    </value_column>

  </db_table>

</store-config>
```



```

<value_column name="endpoint" data_type="VARCHAR(255)">
  <methods>
    <get_method name="getEndPoint"/>
    <set_method name="setEndPoint"/>
  </methods>
</value_column>
<value_column name="appinstance" data_type="VARCHAR(100)">
  <methods>
    <get_method name="getApplicationInstance"/>
    <set_method name="setApplicationInstance"/>
  </methods>
</value_column>
</db_table>

<store type_id="wlng.db.wt.es_example" db_table_name="es_example">
  <identifier>
    <classes key-class="java.lang.String"
value-class="com.acompany.plugin.example.netex.store.NotificationData"/>
  </identifier>
  <index>
    <get_method name="address"/>
  </index>
</store>

<query name="com.bea.wlcp.wlng.plugin.example.netex.Query">
  <sql><![CDATA[SELECT * FROM es_example WHERE correlator LIKE ?]]></sql>
</query>

```

```
</store-config>
```

SLA Example

Below is an example SLA for the example Communication Service. There are examples of service provider group and application group SLAs in:

\$PDS_HOME\pte\resource\sla

Listing 5-6 Example SLA for the example Communication Service

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>

<Sla xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
applicationGroupID="default_app_group"
xsi:noNamespaceSchemaLocation="app_sla_file.xsd">

  <serviceContract>

    <startDate>2008-04-17</startDate>

    <endDate>2099-04-17</endDate>

    <scs>com.acompany.example.plugin.SendDataPlugin</scs>

    <contract/>

  </serviceContract>

  <serviceContract>

    <startDate>2008-04-17</startDate>

    <endDate>2099-04-17</endDate>

    <scs>com.acompany.example.plugin.NotificationManagerPlugin</scs>

    <contract/>

  </serviceContract>

  <serviceContract>

    <startDate>2008-04-17</startDate>

    <endDate>2099-04-17</endDate>
```

```
<scs>com.acompany.example.callback.NotificationCallback</scs>  
<contract/>  
</serviceContract>  
</Sla>
```

Communication Service Example

Container Services

This chapter provides a high-level description of Network Gatekeeper container services. It also provides an overview of other parts of the API available for the use of extension developers:

- [Container service APIs](#)
- [Class: InstanceFactory](#)
- [Class: ClusterHelper](#)
- [Service: EventChannel Service](#)
- [Plug-in](#)
- [Management](#)
- [EDR](#)
- [SLA Enforcement](#)
- [Service Correlation](#)
- [Parameter Tunneling](#)
- [Storage Services](#)
 - [ConfigurationStore](#)
 - [StorageService](#)
- [Shared libraries](#)

JavaDoc for the container API is available in the `$PDS_Home/doc/javadoc` directory of the Platform Development studio installation and on the Network Gatekeeper site at `edocs.bea.com`.

Container service APIs

The Network Gatekeeper container service APIs provide the basic infrastructure by which a Communication Service and the container services of Network Gatekeeper can communicate.

All APIs for inter-working with the container services are found in `com.bea.wlcp.wlng.api.*`.

In order for a network protocol plug of a Communication Service to interact with Network Gatekeeper it must be *deployable* in the context of Network Gatekeeper. Once it is deployable, it can have access to certain utility functions.

Table 6-1 Summary of the container services APIs

Package	Summary
<code>com.bea.wlcp.wlng.api.account</code>	Represents an application instance and the related accounts and groups and the states of the accounts.
<code>com.bea.wlcp.wlng.api.corba</code>	Factory to retrieve an ORB.
<code>com.bea.wlcp.wlng.api.edr.*</code>	Annotations, interfaces and classes used when annotating EDRs. Descriptor classes for alarms, EDRs, and CDRs. Helper classes for EDR listeners. See Annotations, EDRs, Alarms, and CDRs .
<code>com.bea.wlcp.wlng.api.event_channel</code>	Classes to publish and listen to events over cluster-wide event channels. See Service: EventChannel Service .
<code>com.bea.wlcp.wlng.api.interceptor</code>	Interfaces and classes for service interceptors. See Service Interceptors .
<code>com.bea.wlcp.wlng.api.management.*</code>	MBean helper classes. See Making Communication Services Manageable .
<code>com.bea.wlcp.wlng.api.plugin.*</code>	Plug-in related classes and interfaces. See Plug-in .

Table 6-1 Summary of the container services APIs

Package	Summary
com.bea.wlcp.wlng.api.servicecorrelation	Interface to implement if extending the existing service correlation mechanism. See Service Correlation .
com.bea.wlcp.wlng.api.statistics	Annotation for statistics.
com.bea.wlcp.wlng.api.storage	Interfaces and classes for the Storage Service. See Storage Services .
com.bea.wlcp.wlng.api.timers	Factory for using commonj.timers API.
com.bea.wlcp.wlng.api.util	Classes and interfaces for commonly used functions, for example ID generator, InstanceFactory, and clustering.
com.bea.wlcp.wlng.api.work	Factory for using commonj.work API.

Class: InstanceFactory

The Instance Factory is the mechanism used in Network Gatekeeper to retrieve instances of a given interface, class, or abstract class. You retrieve an instance of the Instance Factory using the public static method `getInstance()`. The factory itself has a single method:

`getImplementation(Class theClass)` - Retrieves a class that implements a given interface or extends a given class

The implementation to be used is located and used based on the following rules:

1. First, check the jar file's `instancemap`, a standard `java.util.Properties` file. Every jar file can have its own `instancemap`. The `instancemap` provides a list that maps a given interface, class, or abstract class to the preferred implementation of that functionality. See [Listing 6-1](#) for an example.

Note: The interface name used in the `instancemap` must be unique across all plug-ins for a given Service Enabler. It is not possible to use the same interface in two `instancemap` files belonging to two different plug-ins and still map them to two different implementations.

2. If a mapping is provided and the target class has a public constructor or static singleton method, instantiate it.

3. If there is no explicit mapping, or if there is no public constructor or static singleton method for a mapped class, instantiate an object named according to the following pattern:
`theClass.getClass().getName() + "Impl"` if this exists and has a public constructor or static singleton method.

Listing 6-1 Example instancemap file

```
com.bea.wlcp.wlng.MyInterface=com.bea.wlcp.wlng.MyImplementation
com.bea.wlcp.wlng.MyOtherInterface=com.bea.wlcp.wlng.MyOtherImplementation
```

For details see Javadoc for Package `com.bea.wlcp.wlng.api.util` Class `InstanceFactory`.

Class: ClusterHelper

`com.bea.wlcp.wlng.api.util.cluster.ClusterHelper`

Helper class for getting the JNDI Context for the network and access tier.

For details see Javadoc for Package `com.bea.wlcp.wlng.api.util.cluster` Interface `ClusterHelper`.

Service: EventChannel Service

This service is used to broadcast events to other Network Gatekeeper server instances and to register listeners for events originating in other Network Gatekeeper server instances.

Interface: EventChannel

Use this interface to broadcast events to other instances of Network Gatekeeper, and to register listeners for events originating in them. It is used, for example, in propagating changes of cached data. It is retrieved using the `com.bea.wlcp.wlng.api.event_channel.EventChannelFactory`.

An event has a name and a value, where the name is an identifier for the event and the value is any object implementing `java.io.Serializable`.

The following methods are available:

- `deactivateAllListeners()` - Deactivates all registered listeners.
- `publishEvent` - Publishes an event to all registered listeners.

- `publishEventToOneNode` - Publishes an event to one Network Gatekeeper instance.
- `registerEventListener` - Registers an `EventListener`.
- `unregisterEventListener` - Unregisters an `EventListener`.

Interface: `EventChannelListener`

This interface is used to receive events published using `EventChannel`.

The following method is available:

- `processEvent(String eventType, Serializable event, String source)` - Receives an event.

Plug-in

The `com.bea.wlcp.wlng.api.plugin.*` packages contain a range of interfaces and classes for use by the extension developer.

See [Communication Service Description](#).

Management

Base classes and annotations for giving the Network Gatekeeper Management Console or other JMX tools management access to Communication Services. See [Chapter 9, “Making Communication Services Manageable”](#) for more information. Also see the JavaDoc for the packages: `com.bea.wlcp.wlng.api.management.*`

EDR

See [Chapter 8, “Annotations, EDRs, Alarms, and CDRs.”](#) Also see the JavaDoc for the packages `com.bea.wlcp.wlng.api.edr.*`

SLA Enforcement

SLA enforcement operates on methods identified by the Java representation of the interface, and the operation on the application-facing interface for the Communication Service.

The content of the tag `<scs>` defined in the `<serviceContract>` tag in the SLA is the plug-in type for the plug-in.

An operation on the application-facing interface is represented in the rules according to the following scheme: <service name> and <operation name>.

Parameters in the operation are represented in the rules according to the following scheme:

arg<n>.<parameter name>

where <n> in arg<n> depends on the WSDL that defines the application-facing interface, normally this is arg0.

If the parameter in <parameter name > is

- a composed parameter, the notation is according to the Java Bean notation for that parameter.
- an enumeration, the notation is according to the Java-representation of that parameter, <parameter name >.<enumeration value>. The <enumeration value> is the String representation.

It is not possible to extend the SLA schema. The SLA schema represents a contract supported by Network Gatekeeper base product and is defined by the XSD. It is version controlled and corresponds to the product release that implements the contracted functionality. Extensions cannot restrict or expand the same XSD/contract for specific solutions. Other contracts or configuration mechanisms must be used for specific extensions. However, it is possible to customize the way requests are enforced and filtered using one or both of the following methods:

- If the enforcement or filtering concerns more than one specific plug-in, create an interceptor.
- If the enforcement or filtering concerns a specific plug-in, augment the plug-in itself to perform this function.

Service Correlation

It is often the case that service providers would like to be able to bundle what are to Network Gatekeeper separate services into a single unit for charging purposes. An end user could send an SMS to the provider requesting the location of the coffee shop closest to her current location. The application would receive the network-initiated SMS (one service), do a user location lookup on the customer (one service), and then send the customer an MMS with a map showing the requested information (one service). So three Network Gatekeeper services need to be grouped into a single service charging unit. To do this, Network Gatekeeper provides the framework for a Service Correlation service that uses a Service Correlation ID (SCID) to combine/correlate all the services.

- The Service Correlation ID is optional.
- The Service Correlation ID is captured in the CDRs and EDRs generated from WLNG.
- The Service Correlation ID is propagated as a String.
- The Service Correlation ID is propagated to and from the application in the SOAP header.

The SCID itself is provided either by the application or by an external mechanism that the Communication Service must provide (see [Interface: ExternalInvocation](#)). Network Gatekeeper does not check whether or not it is unique. The SCID is stored in WLS Work Context, so that it can be accessed by both the Access Tier and the Network Tier. The Service Correlation class registers itself as a `RequestContextListener`. When application-initiated request traffic enters the plug-in, the Service Correlation service takes the SCID from the Work Context and places it in the `RequestContext` object, where it will be available to the EDR service. When network-initiated request traffic is leaving the plug-in, the Service Correlation service takes the SCID from the `RequestContext` object and places it in the Work Context, where it can be retrieved by the SOAP Handler and passed along to the application.

Interface: ExternalInvocation

Because Service Correlator IDs may need to be stored across several invocations and a `RequestContext` object exists only for the lifetime of a single request, a Communication Service needs to create a way of storing and retrieving the SCIDs. This is done by implementing the `ExternalInvocation` interface. This interface has two methods: one stores the Service Correlation ID and one retrieves it. The implementor is free to modify the ID once it has been stored, or to use the Invocation object to create IDs in the first place.

When the Service Correlation service takes the SCID (should there be one) out of the Work Context of an application-initiated request, it automatically attempts to store it in an object of this type before putting the SCID in the `RequestContext`.

When a network-initiated request is leaving the plug-in, the Service Correlation service automatically attempts to retrieve an SCID from an object of this type, using the SCID (should there be one) it finds in the `RequestContext` object before it sets the Work Context. In this way, if the `ExternalInvocation` object has modified the SCID in any way, it is this modified version that is put in the Work Context and thus sent on to the application. The `ExternalInvocation` implementation class should have an empty public constructor or a static method that returns itself.

Class: ExternalInvokerFactory

This class is used by the Service Correlation service to locate and instantiate the correct `ExternalInvocation` object. It does this by using an `instancemap`. The `instancemap` entry should look like this:

```
com.bea.wlcp.wlmg.api.servicecorrelation.ExternalInvocation=myPackageStructure.myImplClass
```

where `myImplClass` is the `ExternalInvocation` implementation.

Class: ServiceCorrelation

This class manages the transport and storage of the Service Correlation ID across multiple service invocations.

Implementing the ExternalInvocation Interface

There are four basic steps in creating a custom service correlation:

1. Create a jar file that includes your code. For example:

Listing 6-2 Sample Custom Service Correlation

```
package myPackageStructure;

import com.bea.wlcp.wlmg.api.servicecorrelation.ExternalInvocation;

import
com.bea.wlcp.wlmg.api.servicecorrelation.ExternalInvocationException;

public class MyImplClass implements ExternalInvocation {

    public MyImplClass() {

    }

    public String pushServiceCorrelationID(String scID, String serviceName,
String methodName, String spID, String appID, String appInstGrp) throws
ExternalInvocationException {

        // your code here

        return scID;
    }
}
```

```

    }

    public String getServiceCorrelationID(String scID, String serviceName,
String methodName, String spID, String appID, String appInstGrp) throws
ExternalInvocationException {

        // your code here

        return scID;
    }

}

```

2. Create the instancemap. See [Class: ExternalInvocatorFactory](#).
3. Put the instancemap file in the JAR. This makes your custom service correlation available to the service interceptor InvokeServiceCorrelation.
4. Put the JAR file in \$DOMAIN_Home/lib.

Parameter Tunneling

Parameter tunneling is a feature that allows an application to send additional parameters to Network Gatekeeper and lets a plug-in use these parameters. This feature makes it possible for an application to tunnel parameters that are not defined in the interface that the application is using and can be seen as an extension to the application-facing interface.

The application sends the tunneled parameters in the SOAP header of a Web Services request.

The tunneled parameter can be retrieved in a plug-in by the key. The parameter is fetched from the RequestContext, using the method `getXParam(String key)`. If a value for the key can not be found, null is returned.

Listing 6-3 Get the value of the tunneled parameter 'aParameterName'.

```
RequestContext.getCurrent().getXParam("aParameterName");
```

If the same parameter is defined in the `<contextAttribute>` SLA tag, it should override the parameter tunneled from the application. This behavior, however, is defined per plug-in.

Storage Services

The storage services provided in Network Gatekeeper are of two types, described below:

- [“ConfigurationStore” on page 6-10](#)
- [“StorageService” on page 6-15](#)

ConfigurationStore

The Network Gatekeeper Core exposes a `ConfigurationStore` Java API that Communication Services can use to store simple configuration parameters instead of using JDBC and caching algorithms in each module.

Note: This utility is intended for configuration parameters only, not traffic data

All data stored in a `ConfigurationStore` are stored in a database table and cached in memory.

Below are the characteristics of a `ConfigurationStore`:

- It is a named store.
- Parameters stored in it must be initialized before they can be used.
- Stores can be either domain wide (shared) or limited to a single Network Gatekeeper server (local). The domain wide store type replicates all data changes to all servers in the cluster, while the local store type keeps a different view of the parameters on different servers and data changes affect only the view for that particular server.
- Parameters stored in a `ConfigurationStore` are persisted to database.
- Data in all `ConfigurationStores` are also cached in memory.
- Only one instance of each named `ConfigurationStore` is cached in memory per server.
- Updates to a cluster wide named `ConfigurationStore` is reflected in all cluster nodes.
- The named `ConfigurationStore` only supports parameters of type Boolean, Integer, Long, String, and Serializable.

Interfaces

The Java interface APIs are found in the package `com.bea.wlcp.wlng.api.storage`.

The entry point to configuration stores is through the `com.bea.wlcp.wlng.api.storage.configuration.ConfigurationStoreFactory` using the following method:

```
public abstract ConfigurationStore getStore(String moduleName, String name,
int storeType) throws ConfigurationException;
```

The `ConfigurationStore` service exposes an interface with the following features:

- Methods to initialize a the store with the following data types:

- Boolean,
- Integer,
- Long,
- String
- Serializable

A `ConfigurationStore` is initialized using a name in key/value pair. You get and set configuration parameters using the key.

- Methods to set and get the following data types:

- Boolean,
- Integer,
- Long,
- String
- Serializable

- Methods to add and remove listeners for notifications on updates. When a parameter has been updated in one instance of the `ConfigurationStore`, a notification is broadcast to all other instances of the `ConfigurationStore`.

[Listing 6-4](#) is an example of using the Configuration Store.

Listing 6-4 Example of a ConfigurationStoreHelper

```
package com.acompany.plugin.example.netex.management;

import com.bea.wlcp.wlmg.api.storage.configuration.*;

/**
 * Class used for handling the configuration store.
 *
 * @author Copyright (c) 2007 by BEA Systems, Inc. All Rights Reserved.
 */

public class ConfigurationStoreHandler {

    /**
     * Constants used for the values stored in the store.
     */

    public static final String KEY_NETWORK_HOST = "KEY_NETWORK_HOST";
    public static final String KEY_NETWORK_PORT = "KEY_NETWORK_PORT";

    /**
     * Constant to access either the local store. Note that these are
     * just names for the store.
     */

    private static final String LOCAL_STORE = "local";

    /**
     * Local configuration store instance.
     */

    private ConfigurationStore localConfigStore;

    /**
     * Constructor.
     *
     * @param pluginId The plugin id
     */
}
```



```

    * @throws ConfigurationException An exception thrown if the initialization
failed
    */

    public ConfigurationStoreHandler(String pluginId)
        throws ConfigurationException {

        ConfigurationStoreFactory factory = ConfigurationStoreFactory.getInstance();
        localConfigStore = factory.getStore(pluginId, LOCAL_STORE,
            ConfigurationStore.STORE_TYPE_LOCAL);

        // To obtain a shared configuration store, use
        ConfigurationStore.STORE_TYPE_SHARED

        localConfigStore.initialize(KEY_NETWORK_HOST, "localhost");
        localConfigStore.initialize(KEY_NETWORK_PORT, 5001);
    }

    /**
    * Sets an integer value in the local store.
    *
    * @param key The key associated with the value.
    * @param value The value to store.
    * @throws ConfigurationException An exception thrown if the operation failed
    */
    public void setLocalInteger(String key, Integer value)
        throws ConfigurationException {
        localConfigStore.setInteger(key, value);
    }
    /**

```

Container Services

```
* Gets an integer value from the local store.
*
* @param key The key associated with the value.
* @return The value associated with the key.
* @throws InvalidTypeException thrown if type is invalid.
* @throws NotInitializedException thrown if key value has not been
* initialized.
*/
public Integer getLocalInteger(String key)
    throws InvalidTypeException, NotInitializedException {
    return localConfigStore.getInteger(key);
}
/**
 * Sets a string value in the local store.
 *
 * @param key The key associated with the value.
 * @param value The value to store.
 * @throws ConfigurationException An exception thrown if the operation failed
 */
public void setLocalString(String key, String value)
    throws ConfigurationException {
    localConfigStore.setString(key, value);
}
/**
 * Gets a string value from the local store.
 *
 * @param key The key associated with the value.
```

```

* @return The value associated with the key.
* @throws InvalidTypeException thrown if type is invalid.
* @throws NotInitializedException thrown if key value has not been
* initialized.
*/
public String getLocalString(String key)
    throws InvalidTypeException, NotInitializedException {
    return localConfigStore.getString(key);
}
}

```

StorageService

The Storage Service is used for storing data that is not configuration-related, but related to the traffic flow through a Communication Service, in a cluster-wide store.

It provides mechanisms for:

- Store initialization

A store is created using the `StoreFactory` singleton class, by specifying either a key/value class pair where the value class should be a class that is unique to the Store (recommended), or a Store name.

- Basic Map usage

Since the Store interface extends the `java.util.Map` interface, it can be used as any other Map, and it is extended to be a cluster-wide view of the store.

- Named queries

In addition to the standard `java.util.Map` interface, Stores have support for a `StoreQuery` interface. The behaviors of these named queries are configured as part of the Storage Service configuration files. There is an option to define a cache filter and/or SQL query. If there is an index specified for the Store, this index can be used by implementing

the `IndexFilter` interface for the cache filter. The index is automatically used for SQL queries that can make use of these indexes.

- Store listener

The Store API has support for registering `StoreListeners`. These listeners get notified if the Storage Service decides to automatically remove Store entries (based on configuration parameters). It will not be notified if the extension itself removes entries from the Store.

- Cluster locking

Cluster wide locking can be done using the Store interface. This should be used if the same entry in a Store may be modified on multiple servers at the same time, to avoid getting errors due to concurrent modification when a transaction commits.

A Communication Service extension uses the `StorageService` through an API. The API functionality is implemented by a storage provider. Network Gatekeeper uses a write-through invalidating storage provider. Invalidating stores are backed by a database table. Other storage providers, supporting additional features, can be integrated but are not supported out-of-the box.

Extensions can use the `com.bea.wlcp.wlng.api.storage.Store` interface. This interface extends a `java.util.Map` interface and adds the following methods:

- `addListener`: Adds a listener for the store.
- `getQuery`: Gets a named query.
- `lock`: Takes a cluster-wide lock.
- `release`: Releases the current store instance.
- `removeListener`: Removes a registered listener.
- `unlock`: Unlocks a previously obtained cluster-wide lock.

The storage service use configuration files that defines the configuration for stores and the relationship between the cluster-wide store and the database table that backs the store. In each configuration file it is possible to define named queries towards the store. There is one configuration file per plug-in. Each configuration store configuration file shall, together its XSD and any complex data types stored, be created and packaged in a JAR file, in the directory `$DOMAIN_HOME/config/store_schema`. The configuration file must be named `wlng-cachestore-config-extensions.xml` and it must be present in the root of the JAR.

For details about the store configuration file, see the corresponding xsd:

```
com.bea.wlcp.wlng.storage_4.0.0.0.jar/wlng-cachestore-config.xsd in
$BEA_HOME/wlng400/modules.
```

A Store is retrieved from `com.bea.wlcp.wlng.api.storage.StoreFactory`, either by the name of the store or by the class names of the key/value names. How to retrieve the Store depends on how the store is configured.

The store interface needs to be released when no longer needed. The programming model is to retrieve the Store from the StoreFactory when the Store is used, and to release it once it has finished, using `try { .. } finally { store.release(); }`

Listing 6-5 Example: retrieve a store identified by key/value classes, operate on it, and release it.

```
Store<String, NotificationData> store =
StoreFactory.getInstance().getStore(String.class, NotificationData.class);

try {
    notificationData = store.put(address.toString(), notificationData);
} finally {
    store.release();
}
```

If it is a named store, it can also be retrieved by name as illustrated below.

Listing 6-6 Retrieving a store by name

```
Store<Serializable,Serializable> store =
StoreFactory.getInstance().getStore("A", this.getClass().getClassLoader());
```

Store configuration file

The configuration file `wlng-cachestore-config-extensions.xml` defines attributes of the store and relations between the store, the cache for the store, and the mapping to a database table. This part is used by extension developers.

In addition, the configuration file can contain a section with mapping information between a store, the provider it uses, and the factory for the storage provider. This section should not be used by extension developers.

The XSD for the configuration file is located in

`com.bea.wlcp.wlng.storage_4.0.0.0.jar/wlng-cachestore-config.xsd` in
`$BEA_HOME/wlng400/modules`.

There is one configuration file per plug-in. The file must be embedded in a JAR that contains the file itself and any complex data types used. The JAR must be stored in

`$DOMAIN_HOME/config/store_schema`.

Below is an example of a store configuration file for extensions.

Listing 6-7 Example of a store configuration file for extensions

```
<store-config>

  <db_table name="example_store_notification">

    <key_column name="address" data_type="VARCHAR(255)"/>

    <!-- bucket_column using default BLOB type -->

    <bucket_column name="notification_data_value"/>

    <value_column name="correlator" data_type="VARCHAR(255)">
      <methods>
        <get_method name="getCorrelator"/>
        <set_method name="setCorrelator"/>
      </methods>
    </value_column>

  </db_table>

  <store type_id="wlng.db.wt.example_store_notification">
```

```

        db_table_name="example_store_notification">
<identifier>
    <classes key-class="java.lang.String"
value-class="com.acompany.plugin.example.netex.notification.NotificationData"/
>
</identifier>
<index>
    <get_method name="getCorrelator"/>
</index>
</store>

<query name="com.bea.wlcp.wlng.plugin.example.netex.Query">
    <sql>
        <![CDATA[
SELECT * FROM example_store_notification WHERE correlator = ?
]]>
    </sql>

<filter-class>com.acompany.plugin.example.netex.store.FilterImpl</filter-class>
>
</query>
</store-config>

```

A store is defined between the elements **<store-config>** and **</store-config>**

Each Store has three sections:

- **store**: Defines the store.
- **db_table**: Defines the database table used to persist data in the store.

- **query:** Defines queries on the store. This is optional, only required if non-key queries are used with the store.

<store>

The **store** section defines the store itself. The attribute **type_id** defines the type of the store and a store type identifier. The ID must be mapped to a provider store mapping defined in `wlng-cachestore-config.xml`.

The name should always have the prefix `wlng.db.wt.` when using the storage provider in Network Gatekeeper. The prefix which indicates that it is a write-through cache, that is, data put in the store is always written to database without any delay.

The attribute **db_table_name** identifies the database definition to use.

store contains the following elements:

- **identifier:** Holds one **classes** element. This element defines the classes for the key and the value that defines the store. The class for the key is defined in the attribute **key-class** and the class for the value part is defined in the attribute **value-class**. If a named store is used, the name is given in the element **name**.
- **index:** Defines an index on the cache and one or more get methods. The methods maps to an index on the corresponding columns in the table and potentially a cache index if supported by the provider in use.

<db_table>

The **db_table** section defines the database table used to persist data in store. The attribute **name** defines the table name to use. This name must be the same as the **db_table_name** specified in the **store** section. It contains the following elements:

- **key_column:** Has the attributes **name** and **data_type**. The attribute **name** specifies the column name for the key and the attribute **data_type** specifies the SQL data type for the key.
- **multi_key_column:** Has the attributes **name** and **data_type**. The attribute **name** specifies the column name for *one* part of a *multi key column* and the attribute **data_type** specifies the SQL data type for *the part* of the key. The difference between **multi_key_column** and **key_column** is that **multi_key_column** supports 2 or more columns to be parts of the key, so **multi_key_column** can occur two or more times in the configuration file.
- **bucket_column:** Has the attribute **name**. This attribute specifies the name of the column for the value part of the store. By default this is a BLOB. There is an optional attribute

`data_type`, that can be used if other data types are used. This must be a Java to SQL supported data type mapping and corresponds to the data type in the value part of the store.

- **value_column**: Is used if attributes in the value part of the store should be stored in a separate column. The attribute **name** defines the name of the column and the `data_type` specifies the SQL data type for the column. **value_column** has the sub-element **methods**, which encloses the elements **get_method** and **set_method**. The sub-element **methods** defines the names of the set and get methods for the data stored in **value_column** and the set and get methods for the attribute of the object in the store.

Listing 6-8 Example of single key column configuration

```
...
<db_table name="single_key_store">
  <key_column name="sample_key_1" data_type="VARCHAR(30)">
    <methods>
      <get_method name="getSampleKey1"/>
      <set_method name="setSampleKey1"/>
    </methods>
  </key_column>
  <value_column name="sample_value" data_type="VARCHAR(30)">
    <methods>
      <get_method name="getSampleValue"/>
      <set_method name="setSampleValue"/>
    </methods>
  </value_column>
</db_table>
...
```

Listing 6-9 Example of multi key column configuration

```
...
<db_table name="combined_key_store">
  <multi_key_column name="sample_key_1" data_type="VARCHAR(30)">
    <methods>
      <get_method name="getSampleKey1"/>
      <set_method name="setSampleKey1"/>
    </methods>
  </multi_key_column>
  <multi_key_column name="sample_key_2" data_type="INT">
    <methods>
      <get_method name="getSampleKey2"/>
      <set_method name="setSampleKey2"/>
    </methods>
  </multi_key_column>
  <value_column name="sample_value" data_type="VARCHAR(30)">
    <methods>
      <get_method name="getSampleValue"/>
      <set_method name="setSampleValue"/>
    </methods>
  </value_column>
</db_table>
...
```

<query>

In addition to the standard `java.util.Map` interface, Stores have support for a `StoreQuery` interface. The behavior of these named queries are configured as part of the Storage Service configuration files.

The query section specifies a named query and a filter associated with the named query. The attribute name defines the name of the query. When using the storage service, the query is fetched using this name. The SQL query towards the database is defined in the element **sql**. The actual query is defined in the element `<![CDATA[.....]]>`.

The filter is a class that implements `com.bea.wlcp.wlng.api.storage.filter.Filter`, and the name of the class is defined in the element **filter-class**. The filter implements the method `setParameters`, and a `matches(...)` method.

The `setParameters` method maps the parameters to the filter class or a `PreparedStatement` `setObject` call ordered as the parameter array given. The filter class must implement the `matches` method in such a way that it will yield the same result as the SQL query specified.

Listing 6-10 Example of a named query

```
<query name="com.bea.wlcp.wlng.plugin.example.netex.Query">
    <sql>
        <![CDATA[
            SELECT * FROM example_store_notification WHERE correlator = ?
        ]]>
    </sql>

    <filter-class>com.acompany.plugin.example.netex.store.FilterImpl</filter-class>
</query>
```

Listing 6-11 Example of using the named query using a filter

```
StoreQuery<String, NotificationData> storeQuery =
store.getQuery("com.bea.wlcp.wlng.plugin.example.netex.Query");
```

```
storeQuery.setParameters(correlator);  
set = storeQuery.entrySet();
```

Listing 6-12 Example of a filter implementation

```
public class FilterImpl implements Filter {  
  
    /**  
     * The query parameters.  
     */  
    private Serializable[] parameters;  
  
    /**  
     * Default constructor.  
     */  
    public FilterImpl() {  
  
    }  
  
    /**  
     * Evaluate if a store entry matches the filter.  
     *  
     * @param value The store entry value to evaluate.  
     */  
    public boolean matches(Object value) {  
  
        if (parameters == null || value == null || parameters.length == 0) {
```

```

        return false;
    }

    if (value instanceof NotificationData) {
        String compareValue = ((NotificationData) value).getCorrelator();

        if (compareValue != null) {
            return compareValue.equals(parameters[0]);
        }
        return compareValue == parameters[0];
    }

    return false;
}

/**
 * Set query parameters. The parameters will be ordered as provided to the
 * StoreQuery and it is the responsibility of the implementation to handle
 * them in this order.
 *
 * @param parameters The query parameters to use.
 */
public void setParameters(Serializable ... parameters)
    throws StorageException {

    this.parameters = parameters;
}

```

```
}

```

```
}

```

<provider-mapping>

The **provider-mapping** section contains definitions of which storage provider a given type-id is mapped to. This section shall not be used unless a custom storage provider is used.

In the **type_id** attribute for **store_mapping type**, the same ID shall be used as when the store was defined. A best match (longest matching entry) is performed. A wildcard (*) can be used at the end of **type_id** to match the prefix.

The **<provider-name>** entry references the type of store being used, see “<providers>” on [page 6-27](#).

The **type_id** for the storage provider mapping in use is `wlng.db.wt.*`, which references the write-through provider.

There is another set of **type_id** attributes defined for **store_mapping**:

- `wlng.db.log.*`, which is used for internal purposes only.
- `wlng.db.wb.*`, which shall be used if the storage provider supports write-behind operations. The invalidating storage provider does not support write-behind operations, write-through will be used.
- `wlng.db.wt.*`, which shall be used if the storage provider supports write-through operations.
- `wlng.cache.*`, which shall be used if the storage provider supports cache-only operations. The invalidating storage provider does not support cache-only operations, write-through will be used.
- `wlng.local.*`, which is used for internal purposes only.

These store mapping types are present for internal and future use. All store mapping types (except for the internal `wlng.db.log.*`) are by default mapped to the keyword `invalidating` which represents the invalidating storage provider. This should not be changed unless a custom storage provider is used.

<providers>

The **providers** section contains mappings between the **provider-name** defined in the **provider-mapping** section and the factory class for the storage provider. This section should not be changed unless a custom storage provider is used.

Shared libraries

It is possible for multiple plug-ins to share common libraries, for example a third party library or custom code that can be shared.

If there are such parts, these should preferably not be packaged into the plug-in jar but instead be copied into the `APP-INF/lib` directory of the Communication Service network tier EAR. All classes in this directory are available for all of the plug-ins in the EAR.

Communication Service Description

This chapter provides a description of an Communication Service and its components:

- [High-level components](#)
- [Communication Service Common](#)
- [Plug-in](#)
- [Management](#)
- [SLA Enforcement](#)
- [Shared libraries](#)

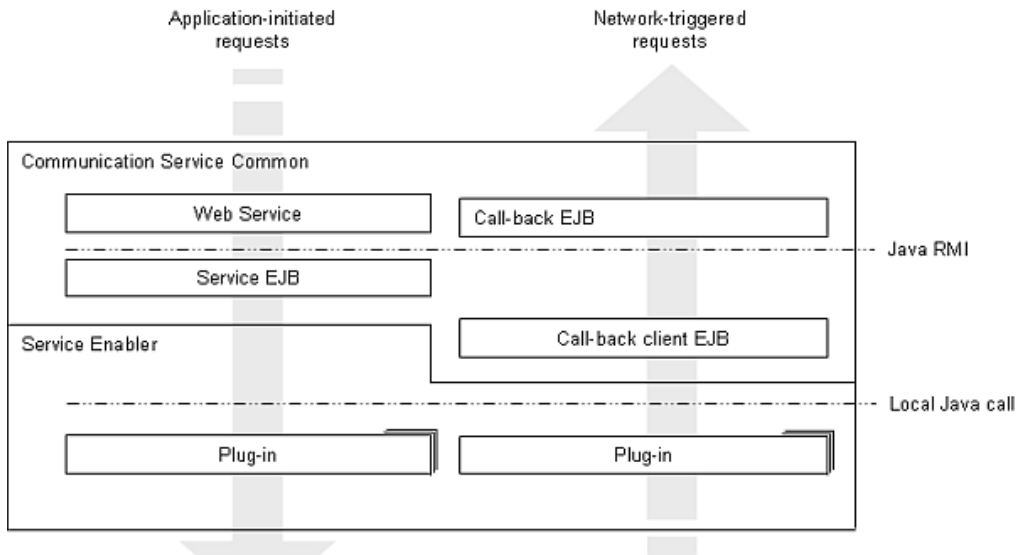
High-level components

From a component point of view, a Communication Service consists of a set of components:

- A Service Web Service
- A Service EJB
- A Callback EJB
- A Call-back client EJB
- A set of network protocol plug-ins

One set of the above handles application-initiated requests, while another handles network-triggered requests, as illustrated in [Figure 7-1](#). Some calls are remote since the modules may be deployed in separate clusters.

Figure 7-1 High-level component of a Communication Service



Communication Service Common

The Communication Service common parts are auto-generated based on one or more WSDLs. For application-initiated requests, these are referred to as service WSDLs, while for network triggered requests, they are referred to as callback WSDL files.

Based on the service WSDLs, the following common parts of a Communication Service is generated using the Eclipse wizard:

- Service Web Service
- Service EJB
- Call-back EJB
- Call-back client EJB

The Service Web Service implements the interfaces defined in the set of WSDL files that defines the Web Service for application-initiated requests.

The Web Service is packaged into one single WAR file. An example of this is Parlay X 2.1, which defines the following interfaces for application-initiated requests: `SendSms`, `ReceiveSms`, and `SmsNotificationManager`. The Service Web Service implements all the above interfaces and is packaged into one single WAR file for the Communication Service.

The Web Service makes a Java RMI call to the Service EJB which, using the Plug-in Manager, calls the appropriated plug-in instance. The operations defined between the Service Web Service and the Service EJB are Java realizations of the interfaces defined in the service WSDLs. The Service EJB is packaged into one single JAR file for the Communication Service.

The Callback EJB is a Web Services client that uses a Web Service implemented by an application. It uses the interfaces defined in the set of WSDL files that defines the Web Service for network-triggered requests, the callback WSDL files. The Web Service client is packaged into one single JAR file for the Communication Service.

The Callback EJB client is a client library that abstracts the remote call between the plug-in POJO and the Callback EJB and provides an invalidating cache of references to the remote object in order to support in-production redeployment of the EAR file for the access tier. The Callback EJB client is packaged into one single JAR file for the Communication Service.

Module	Description	North interface	South interface
Service Web Service	<p>Implements the interfaces defined in the set of WSDL files that defines the Web Service for application-initiated requests. Passes on the requests to the Service EJB. Any Service EJB of the same type can be chosen, regardless of which server it is deployed in. The requests are load-balanced across the different server instances.</p> <p>Packaged into one single WAR file.</p> <p>Deployed as a part of the access tier EAR for the Communication Service.</p> <p>The Service Web Service is transparent to an extension developer.</p>	SOAP/HTTP representation of the Service WSDLs	Java RMI representation of the Service WSDLs
Service EJB	<p>Accepts requests from the Service Web Service implementation and propagates them to the appropriate plug-in using the Plug-in Manager.</p> <p>The Service EJB is responsible for:</p> <ul style="list-style-type: none"> • Constructing the RequestInfo object. • Converting any exception caught to an exception that is defined in the Service WSDLs. <p>This functionality must be implemented in the PluginFactory class, which extends Class: RequestInfo.</p> <p>Packaged in one single JAR file.</p> <p>Deployed as a part of the network tier EAR.</p>	Java RMI representation of the Service WSDLs	Local Java representation of the Service WSDLs.

Module	Description	North interface	South interface
Callback EJB	<p>A Web Services client that uses a Web Service implemented by an application.</p> <p>Accepts requests from the Service callback client EJB and propagates them to an application.</p> <p>Packaged into one single JAR file for the Communication Service.</p> <p>Deployed as a part of the access tier EAR.</p>	SOAP/HTTP representation of the Service callback WSDLs.	Java RMI representation of the Callback WSDLs
Callback EJB client	<p>A client library that abstracts the remote call between the plug-in and the Callback EJB.</p> <p>Accepts requests from a plug-in and propagates them to the Callback EJB.</p> <p>It provides an invalidating cache of references to the remote object in order to support in-production redeployment of the EAR file for the access tier.</p> <p>Any Callback EJB of the same type can be chosen, regardless of which server it is deployed in. The requests are load-balanced across the different server instances.</p> <p>See Class: CallbackFactory and Interface: Callback.</p> <p>Packaged into one single JAR file for the Communication Service.</p> <p>Deployed as a part of the network tier EAR.</p>	Java RMI representation of the Service callback WSDLs.	Local Java representation of the Callback WSDLs.

Plug-in

The `com.bea.wlcp.wlmg.api.plugin.*` packages contain a range of interfaces and classes for use by the extension developer.

First of these is a set of interfaces that define the borders of a plug-in and related helper classes. These borders are used to apply aspects. See JavaDoc for `com.bea.wlcp.wlng.plugin`

Plug-in Service and Plug-in Instance

A plug-in service is a JEE application that implements `com.bea.wlcp.wlng.api.plugin.ManagedPluginService`. It has:

- A life-cycle, defined in `com.bea.wlcp.wlng.api.plugin.PluginServiceLifecycle`.
- A registry, defined in `com.bea.wlcp.wlng.api.plugin.PluginService`.
- a factory to create plug-in instances, defined in `com.bea.wlcp.wlng.api.plugin.PluginInstanceFactory`

The plug-in instance is a class that implements `com.bea.wlcp.wlng.api.plugin.ManagedPluginInstance`. It has:

- A life-cycle defined in `com.bea.wlcp.wlng.api.plugin.PluginInstanceLifecycle`.
- A set of `PluginNorth` and `PluginSouth` interfaces that it implements. These interfaces are defined by the application-facing interfaces and the network-facing interfaces.
- A registry, defined in `com.bea.wlcp.wlng.api.plugin.PluginInstance`. This registry holds the list of the registered interfaces.
- Logic that examines the data in a request and determines if the instance can handle it or not. The interface for this logic is defined in `com.bea.wlcp.wlng.api.plugin.PluginInstance`.
- Logic that maintains the state of a connection. The interface for this logic is defined in `com.bea.wlcp.wlng.api.plugin.PluginInstance`.

Plug-in routing and registration with the Plug-in Manager is done by the plug-in instance. It is the plug-in instance that is part of the traffic flow.

Life-cycle management is performed on the plug-in service.

States

A plug-in service is in one of a distinct set of states:

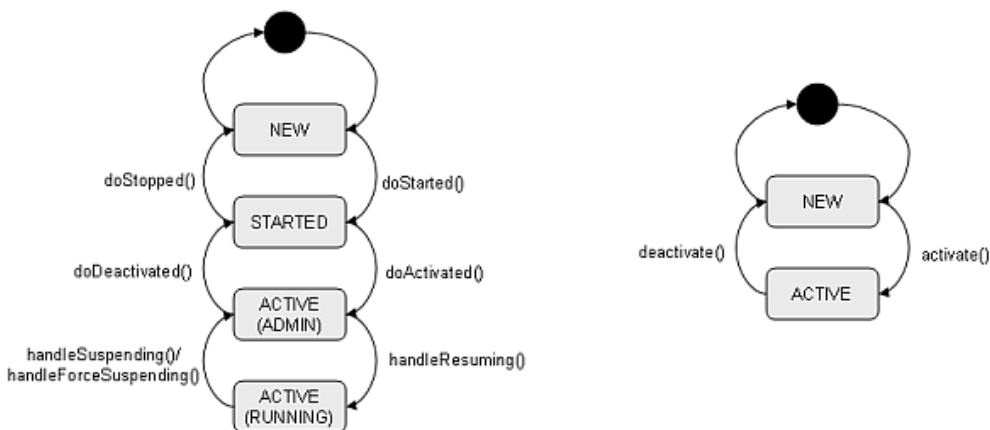
- NEW
- STARTED

- ACTIVE (ADMIN)
- ACTIVE (RUNNING)

The plug-in instance is in one of the following states:

- NEW
- ACTIVE

Figure 7-2 States for a plug-in service (left) and a plug-in instance (right)



The state transitions in Table 7-1 are triggered by either the start-up sequence of the server the plug-in is deployed in or an explicit deployment of the plug-in using either the weblogic.Deployer, see <http://edocs.bea.com/wls/docs100/deployment/wldeployer.html>, or the administration console, see:

- <http://edocs.bea.com/wls/docs100/ConsoleHelp/taskhelp/deployment/DeployApplicationsAndModules.html>
- <http://edocs.bea.com/wls/docs100/ConsoleHelp/taskhelp/deployment/UpdateApplication.html>
- <http://edocs.bea.com/wls/docs100/ConsoleHelp/taskhelp/deployment/RemoveAnApplicationOrModuleFromADomain.html>.

Note: All deployments are one EAR level, which means that individual plug-ins are not target, but all plug-ins within the EAR are affected.

Table 7-1 Plug-in service state transitions

Transition	Triggered by	Descriptions
init	Deployment or startup.	The plug-in service has been created and initialized. The only method that will be called in this state is doStarted()
doStarted	Deployment or startup	The plug-in service should perform as much initialization as possible without be externally visible. Examples include: retrieve configuration data, create internal objects, and initialize stores.
doActivated	Deployment or startup	The plug-in service should continue activation and become visible, for example register MBeans, without accepting traffic.
handleResuming	Deployment or startup.	The plug-in service should order all plug-in instances to establish connections with the network node, if applicable, and accept traffic.
handleSuspending	Graceful undeployment/re deployment/stop That is, invoking weblogic.Deployer with -graceful	The plug-in service should order the plug-in instance to reject new traffic, but to continue processing of in-flight work. A com.bea.wlcp.wlng.api.plugin.CompletionBarrier is provided in the request. When all in-flight work has been processed, the plug-in should get the com.bea.wlcp.wlng.api.plugin.CompletionBarrierCallback from the CompletionBarrier and call completed() on the CompletionBarrierCallback.
handleForceSuspending	Forced undeployment/re deployment/stop That is invoking weblogic.Deployer with -retiretimeout	The plug-in service should order the plug-in instance to reject new traffic and to discard in-flight work.
doDeactivated	Undeployment.	The plug-in service should deactivate itself, unregister any MBeans and become invisible.
doStopped	Undeployment.	The plug-in service should perform cleanup and be available for garbage collection.

The state transitions in [Table 7-2](#) are triggered by either the start-up sequence of the server the plug-in instance is created in or an explicit creation of a new instance using the Plug-in manager: see [Managing and Configuring the Plug-in Manager](#) in the System Administrator's Guide.

Table 7-2 Plug-in instance state transitions

Transition	Triggered by	Descriptions
activate	Creation of the plug-in instance using the Plug-in Manager MBean.	<p>The plug-in instance has been created. Depending on the state of the plug-in service, the plug-in instance should take the appropriate action. If the plug-in service is in state:</p> <ul style="list-style-type: none"> ACTIVE (ADMIN), the plug-in instance shall: <ul style="list-style-type: none"> – instantiate and register the PluginNorth and call-back interfaces with the Plug-in Manager. – instantiate and register the PluginSouth interfaces with the Plug-in Manager. – instantiate any ConfigurationStore. – register the MBean for the instance. ACTIVE (RUNNING), the plug-in shall: <ul style="list-style-type: none"> – connect to the network node, if a connection-oriented protocol is used. – register call-backs with the network node, if any.
deactivate	Destruction of the plug-in instance using the Plug-in Manager MBean.	<p>The plug-in instance shall:</p> <ul style="list-style-type: none"> • de-register any call-backs with the network node. • disconnect from the network node, if connected. • de-register the MBean for the instance. • cancel any timers.

The Plug-in Manager maintains a pool of plug-in instances. This pool is provided to the plug-in when `init()` is called. This pool can be used to iterate over in order to propagate events related to state transitions in the plug-in service to the plug-in instances.

The Plug-in Manager has a registry of all PluginNorth and PluginSouth interfaces, and it is the responsibility of the plug-in instance to register these interfaces with the Plug-in Manager. The Plug-in Manager use these when routing it to an appropriate plug-in instance. The Plug-in

Manager queries the plug-in instance for information in order to make a routing decision. A plug-in instance maintains:

- A list of PluginNorth interfaces
- A list of PluginSouth interfaces
- Whether the plug-in instance has a connection to the network node.
- Custom pattern matching, where the plug-in examines the request and marks the plug-in instance as either a
 - mandatory
 - optional
 - requiredtarget for the request.

The plug-in service maintains a:

- Service type, used by all plug-in instances to generate EDRs, CDRs, and Statistics.
- List of supported address schemes, used by the Plug-in Manager when taking a routing decision.

PluginPool

A collection of PluginInstances. The pool is populated when creating a plug-in instance using the PluginInstanceFactory.

- Using the pool, the plug-in service can list plug-in instances and get a plug-in instance by its plug-in instance ID.

Interface: Plugin

Superinterface for [Interface: PluginNorth](#), [Interface: PluginNorthCallBack](#), and [Interface: PluginSouth](#).

It must be implemented by all classes that handles application-triggered requests from the service EJB to the plug-in. There must be one class per interface.

PluginNorth defines the entry-point for application-initiated requests and is one of the borders at which aspects are woven.

It must be implemented by any plug-in that handles network-triggered requests, either new requests or notifications. `PluginNorthCallback` defines the limit between the plug-in and the service callback EJB and further on to an application.

Interface: PluginNorth

All interfaces in the plug-in that implement the traffic interfaces defined in the service WSDLs must implement this interface. A list of the implementations is maintained in the class that implements in [Interface: ManagedPluginInstance](#). Statistics aspects are applied for classes that implement this interface.

Interface: PluginNorthCallback

All interfaces in the plug-in that implement the traffic interfaces defined in the service callback WSDLs must implement this interface.

Interface: PluginSouth

This interface must be implemented by the plug-in. Defines the south border of a plug-in, that is the network-facing border.

It contains methods used to rebuild the object defined by [Interface: RequestContext](#) for network-initiated requests, using information from the object defined by [Interface: ContextMapperInfo](#), and methods for resolving which application instance the request belongs to.

When a network triggered request arrives at the plug-in, the usual pattern is to correlate the request with a previous subscription for notifications.

By extending `PluginSouth` in the class that implements the request, aspects that call the method

```
public String resolveAppInstanceId(ContextMapperInfo)
```

are applied.

It is the responsibility of the plug-in instance to extract the information provided in the request and to resolve the application instance that matches this data as a part of the rebuilding of the `RequestContext`. This is done using the [Context Aspect](#).

After resolving the application instance, the method

```
public void prepareRequestContext(RequestContext ctx, ContextMapperInfo
info)
```

is called. In the implementation of this method, the plug-in instance has the option to add additional data to the object defined by [Interface: RequestContext](#).

Interface: ManagedPluginService

The interface a plug-in service must implement.

It extends the interfaces `PluginService`, `PluginInstanceFactory` and `PluginServiceLifecycle`.

Interface: PluginService

The interface that defines the plug-in service when registered in the Plug-in Manager.

It defines a set of attributes that must be defined by implementing the following methods:

- `getNetworkProtocol()`, returns a descriptive name for the supported network protocol. For example "SMPP v3.4"
- `getServiceType()`, returns a `ServiceType`. See [Class: ServiceType](#).
- `getSupportedSchemes()`, returns a list of supported address schemes. This is a String array of URI schemes: for example "tel", "mailto", and "sip".

Interface: PluginInstanceFactory

Factory that allows a plug-in service to create plug-in instances.

Defines the method:

```
ManagedPluginInstance createInstance(String pluginInstanceId)
```

The plug-in service is responsible for creating an instance of the class implementing [Interface: ManagedPluginInstance](#) when this method is invoked. The method is triggered by the method `createPluginInstance` on the Plug-in Manager MBean.

Interface: PluginServiceLifecycle

Defines the life-cycle for a plug-in service. See [States](#).

Interface: ManagedPluginInstance

Must be implemented by a plug-in instance.

It extends the interfaces `PluginInstance` and `PluginInstanceLifecycle`.

Interface: PluginInstance

Defines the plug-in instance when registered in the Plug-in Manager.

The plug-in instance is responsible for:

- maintaining a list of north interfaces that the plug-in implements.
- maintaining a list of south interfaces that the plug-in implements.

Both of these lists are arrays of `PluginInterfaceHolder`.

The lists shall be returned when `getNorthInterfaces()` and `getSouthInterfaces()` are invoked, respectively.

The plug-in instance is also responsible for implementing `customMatch(RequestInfo requestInfo)`. In this request, the plug-in instance examines the `RequestInfo` object and decides if the plug-in instance can be used to serve the request. By returning:

- **MATCH_OPTIONAL**: Indicates that the request can be served by any plug-in instance. The request is completely stateless.
- **MATCH_REMOVE**: The request cannot be served. This situation can occur, for example, if a plug-in service does not implement the method being invoked or if the request relates to a previous request which is known only to a subset of the plug-in instances in the cluster.
- **MATCH_REQUIRED**: The request must be served by the plug-in instance. This situation can occur, for example, if the request relates to a previous request which is known only to a subset of the plug-in instances in the cluster.

Only these constants can be returned.

The plug-in instance is also responsible for maintaining information on the connection status with the network node it is connected to by returning `True` or `False` when `isConnected()` is invoked.

All methods in this interfaces are invoked by the Plug-in Manager when selecting a plug-in instance to route the request to.

Interface: PluginInstanceLifecycle

Defines the life-cycle for a plug-in service. See [States](#).

Class: RequestFactory

The Request Factory is used to perform application-initiated request processing both before and after a request is processed in the plug-in. Each Communication Service must have one implementation of the RequestFactory per each application-facing interface, named according to the pattern: `<myinterfacename>.PluginFactory`. A skeleton for the factory is generated by the Eclipse plug-in.

The RequestFactory has two main functions:

- Packages routing information contained in the request into a `RequestInfo` object that the Plug-in Manager uses to select an appropriate plug-in to process the request. See below for more information on `RequestInfo` objects.

Note: In order to support sendlists which target multiple plug-ins, the Request Factory implementation must support three methods that are not required for non-sendlist based plug-ins:

- `createRequestInfos`: allows the creation of multiple `RequestInfo` objects. Each instance of a `RequestInfo` object is matched to a plugin. For example if an SMS message request is sent to 3 addresses, the factory should create an array of 3 `AddressRequestInfo` objects.
- `createPartialRequest`: splits a request into multiple requests sent to different plug-ins
- `mergeResults`: merges the results reported back by multiple plug-ins into a single result.

For more information, see the RequestFactory JavaDoc

Note: Plug-ins are invoked in sequence and if one of them fails the whole request is considered a failure. In this case, an exception is thrown and the transaction is rolled back.

- Translates any exceptions thrown in the plug-in (or the underlying network) into a form that can be sent back to the application.

Class: CallbackFactory

This class is used by a plug-in instance to get an implementation of [Interface: Callback](#). There is one CallbackFactory per interface defined in the callback WSDLs.

The naming pattern is `com.acompany.example.callback.<interface name>CallbackFactory`

The implementation of the interface is fetched using the following pattern:

Listing 7-1

```
import com.acompany.example.callback.NotificationCallback;
import com.acompany.example.callback.NotificationCallbackFactory;
...
private NotificationCallback cachedNotificationCallback = null
....
private NotificationCallback getNotificationCallback() {
    if(cachedNotificationCallback == null) {
        cachedNotificationCallback =
            NotificationCallbackFactory.getInstance().create();
    }
    return cachedNotificationCallback;
}
```

Interface: Callback

This interface is used by a plug-in to propagate a network-triggered request from the plug-in to the callback EJB. The interface defines a Java representation of the methods defined in the callback WSDLs. There is one interface per interface defined in the callback WSDLs.

The naming pattern is `com.acompany.example.callback.<interface name>Callback`

Class: RequestInfo

The object created by the RequestFactory to hold information from the application-initiated request. There are four sub-classes of `RequestInfo` that can be used depending on the request:

- `AddressRequestInfo`, if the request contains an address.
- `CorrelatorRequestInfo`, if the request contains a correlator.

- `RegistrationIdentifierRequestInfo`, if the request contains a registration identifier.

`RequestIdentifierRequestInfo`, if the request contains a request identifier.

Class: `ServiceType`

This is an abstract utility class that each plug-in must implement. An object of this type is passed to the Plug-in Manager when the plug-in registers itself, so that the Plug-in Manager can query for service type.

Aspects takes care of making this service type available in the request thread of each plug-in. The service type is used by various services, including the `EdrService`.

Table 7-3 Existing `ServiceTypes`

ServiceType	Plugin
<code>AccessServiceType</code>	Access
<code>ThirdPartyCallServiceType</code>	Third-party Call
<code>CallNotificationServiceType</code>	Call Notification
<code>SmsServiceType</code>	Sms
<code>MultimediaMessagingServiceType</code>	Mms
<code>TerminalLocationServiceType</code>	Terminal Location
<code>AudioCallServiceType</code>	Audio Call
<code>PresenceServiceType</code>	Presence

Interface: `ContextMapperInfo`

This interface defines a `ContextMapperInfo` object. When network-initiated traffic enters the plug-in from the network-facing (south) side, aspects take any annotated arguments from the network call that will be needed by the plug-in for correlation purposes and places them in this very short-lived object. Arguments are stored by key, defined when the annotation is set, that make it possible to retrieve a particular value. So if an argument is annotated with `@MapperInfo(C)`, its value can be retrieved using the key “C”. Methods in the plug-in that need

to retrieve these arguments in order to perform a mapping (for example, associating a notification with the session ID of the request that established it) can use this object. The `PluginSouth` interface includes one such method, `resolveAppInstanceId`.

Interface: RequestContext

Defines a `RequestContext` object. A `RequestContext` object is available in all Communication Services for both application-initiated and network-initiated requests. It contains contextual information about the request, including the service provider account ID, application account ID, and application instance of the application that initiated either the request or the notification, as well as the session ID.

Class: ManagedPlugin

Deprecated. Use `ManagedPluginService` instead.

Allows the plug-in to register itself in the Plug-in Manager. See [Class: AbstractManagedPlugin](#).

Class: AbstractManagedPlugin

Deprecated. Use `ManagedPluginService` instead.

Extends `ManagedPlugin`, implements `ServiceDeployable`. It makes the plug-in deployable as a service in Network Gatekeeper, and assists in registering the plug-in with the Manager. See the `com.bea.wlcp.wlmg.api.plugin.common` package JavaDoc for details.

Management

These are base classes and annotations for giving the Network Gatekeeper Management Console or other JMX tools management access to Communication Services. See [Chapter 9, “Making Communication Services Manageable”](#) for more information. Also see the JavaDoc for the packages: `com.bea.wlcp.wlmg.api.management.*`

SLA Enforcement

SLA enforcement operates on methods identified by the Java representation of the interface, and the operation of the application-facing interface for the Communication Service

The content of the tag `<scs>` defined in the `<serviceContract>` tag in the SLA is the plug-in type for the plug-in.

An operation on the application-facing interface is represented in the rules according to the following scheme: <service name> and <operation name>.

Parameters in the operation are represented in the rules according to the following scheme:

arg<n>.<parameter name>

where <n> in arg<n> depends on the WSDL that defines the application-facing interface, normally this is arg0.

If the parameter in <parameter name > is

- a composed parameter, the notation is according to the Java Bean notation for that parameter.
- an enumeration, the notation is according to the Java-representation of that parameter, <parameter name >.<enumeration value>. The <enumeration value> is the String representation. See [Using the Platform Test Environment](#) for information about the SLA Editor.

Shared libraries

It is possible for multiple plug-ins to share common libraries: for example, a third party library or custom code that can be shared.

If there are such parts, these should preferably not be packaged into the plug-in jar but instead be copied into the `APP-INF/lib` directory of the Communication Service EARs that utilizes this shared library. All jars in this directory are available for each of the plug-ins in the EAR.

Annotations, EDRs, Alarms, and CDRs

The following section describe aspects and generation of EDRs, alarms, CDRs, and statistics:

- [“About aspects and annotations” on page 8-2](#)
- [“How aspects are applied” on page 8-2](#)
- [“Context Aspect” on page 8-3](#)
- [“EDR Generation” on page 8-6](#)
 - [“Exception scenarios” on page 8-8](#)
 - [“Adding data to the RequestContext” on page 8-9](#)
 - [“Trigger an EDR programmatically” on page 8-12](#)
 - [“EDR Content” on page 8-13](#)
 - [“RequestContext and EDR” on page 8-20](#)
- [“Categorizing EDRs” on page 8-22](#)
 - [“The EDR descriptor” on page 8-22](#)
- [“Check-list for EDR generation” on page 8-33](#)
- [“Alarm generation” on page 8-37](#)
 - [“Trigger an alarm programmatically” on page 8-38](#)
 - [“Alarm content” on page 8-39](#)

- “CDR generation” on page 8-41
 - “Triggering a CDR” on page 8-41
 - “Trigger a CDR programmatically” on page 8-41
 - “CDR content” on page 8-42
 - “Out-of-the box (OOTB) CDR support” on page 8-47
- “Extending Statistics” on page 8-48

About aspects and annotations

Aspects allow developers to manage cross-cutting concerns in their code in a straight forward and coherent way. Aspects in Network Gatekeeper (pointcuts, advice, etc.) are written in the AspectJ 1.5.3 annotation style. There is already support for editing annotations in many modern IDEs, and aspects are simply set up as annotated classes.

How aspects are applied

All aspects are applied at build time by weaving the byte code of previously complied Java packages. Minimal reflection is used at runtime to make aspect based decisions.

Different aspect types are applicable to different Network Gatekeeper modules. In general there are two categories of aspects:

- Those restricted to the code for the traffic flow
- Those that can be applied to other packages.

Note: In this case, traffic flow is defined to include only plug-in implementations.

Traffic aspects are subdivided into two categories:

- Those that are always applied
- Those that are controlled using annotations.

Only statistics aspects are always applied because they are used in to enforce licensing. Traffic aspects are applied to North and South boundaries of a plug-in as well as to the internal processing of the plug-in.

Annotations are used to control the aspects that are not always applied for each plug-in. These annotations are defined as part of the functional areas that a given set of aspect implements. They allow the plug-in to communicate with the aspects as well as to customize their behavior.

Context Aspect

The Context aspect is woven at compile time, using `PluginSouth` as a marker.

While requests coming from the north interface have a valid context (with attributes like Service provider account ID, application Account ID, and so on) any events triggered by the network and entering a plug-in's south interface do not have a valid context.

The Context aspect solves this problem by rebuilding the context as soon as a south interface method is invoked: after this aspect is executed, a valid context will be available for any subsequent usages, such as the EDR aspect. All methods inside a class implementing the interface `PluginSouth` are woven by the Context aspect.

The Context aspect requires the following in order to correctly weave the south interface methods and be able to rebuild the context:

- Each Plugin must explicitly register its north and south interfaces.
- Each south interface must implement the `resolveAppInstanceId()` and `prepareRequestContext()` methods of the `PluginSouth` interface.
- North interfaces must implement `PluginNorth` and south interfaces must implement `PluginSouth`.

The following rules apply for methods in classes that implement `PluginNorth`:

- The default behavior is that EDRs are triggered only for exceptions and callbacks to EJBs in the access tier (Service Callback EJB)
- If a method is annotated with `@NoEdr`, no EDRs will be generated. It overrides the default behavior.
- If a method is annotated with `@EDR`, 2 EDRs will be generated:
 - When entering the method
 - When exiting the method.

The following rule applies for methods in classes that implement `PluginSouth`:

- Methods that perform requests to the network may have a parameter annotated with `@MapperInfo` in order to be able to rebuild the `RequestContext` when the response to the request arrives from the network. The annotated parameter must be used as a key to resolve the application instance ID using some plug-in specific lookup.

- Methods must implement `resolveAppInstanceId(ContextMapperInfo info)` in `PluginSouth` and return the application instance ID that corresponds to the original request to the network.

The ways of doing this are plug-in-specific, but normally a network triggered request is tied to an application instance in a store that is managed by the plug-in. The store used for context mapping may be a local cache or a cluster wide store, depending on whether responses are known to always arrive on the same plug-in instance, or if they can arrive to a plug-in on another server in the cluster.

Example:

1. An application sends a request to the network and an ID for this request is either supplied by the network or generated by the plug-in. At this point the originator of the requests, the application instance, is known since the request originated from an application.
2. The plug-in puts the application instance ID and the ID for the request into a store.
3. At a later stage, when a response to the original requests arrives to the plug-in, the method `resolveAppInstanceId()` is called by aspects.
4. In this method, the plug-in must perform a lookup in the store of the application instance related to that request and return the application instance ID to the aspect.
5. The aspect authenticates the application instance with the container and puts the application instance ID in the `RequestContext`.
6. The method in the plug-in receives the request from the network and the `RequestContext` contains the application instance ID.

In the example below the method `deliver(...)` is a request from the underlying network. The `destinationAddress` is annotated to be available by the aspect that handles network-triggered requests associated with this request represented by constant `C`.

`NotificationHandler` handles the store for notifications and supplies all necessary parameters to the store.

Listing 8-1 Application initiated request

```
protected static final String C = "destinationAddress";

@Edr

public void deliver(String data,
```

```

        @ContextKey(EdrConstants.FIELD_DESTINATION_ADDRESS)
        @MapperInfo(C) String destinationAddress,
        @ContextKey(EdrConstants.FIELD_ORIGINATING_ADDRESS) String
originatingAddress,
        String nwTransactionId)
    throws Exception {

        notificationHandler.deliver(data, destinationAddress, originatingAddress,
nwTransactionId);

    }

```

When a network triggered event occurs, the aspect calls `resolveApplicationInstanceGroup(...)` in `PluginSouth` and the plug-in looks up the application instance using any argument available in `ContextMapperInfo` that can help the plug-in to resolve this ID from `ContextMapperInfo`, using `info.getArgument(C)`. The application instance ID is returned to the aspect and the execution flow continues in the plug-in, with a `RequestContext` that contains the application instance ID, session ID and so on.

Listing 8-2 Rebuilding RequestContext

```

protected static final String C = "destinationAddress";

public String resolveAppInstanceGroupdId(ContextMapperInfo info) {

    String destinationAddress = (String) info.getArgument(C);

    NotificationData notificationData = null;

    try {
        notificationData =
StoreHelper.getInstance().getNotificationData(destinationAddress);
    } catch (StorageException e) {

```

```
        return null;
    }

    if (notificationData == null) {
        return null;
    }

    return notificationData.getAppInstanceGroupId();
}
```

Below are the steps you have to take to make your plug-in compliant with the Context aspect:

- Make sure to register all your PluginSouth objects before registering your plug-in in the Plug-in Manager.
- Make sure to implement the `resolveAppInstanceGroupId()` method for each PluginSouth instance.
- Annotate each parameter in south object methods that you need to have when aspect calls back the `resolveAppInstanceGroupId()` or the `prepareRequestContext()` methods. All the annotated parameters will be available in the `ContextMapperInfo` parameter. The aspect needs to have them annotated to be able to store them into the `ContextMapperInfo` object.

EDR Generation

EDRs are generated in the two following ways:

- automatically using aspects at given points in the traffic execution flow in a plug-in.
- manually anywhere in the code using the `EdrService`.

EDRs should be generated at the plug-in boundaries (north and south), using the `@Edr` annotation to ensure that the boundaries are covered. Additional Edrs can be added elsewhere in the plug-in if needed: for example for CDRs.

For extensions, the EDR ID should be in the range 500 000 to 999 999.

EDRs are generated automatically by an aspect in the following locations in the plug-in:

- Before and after any method annotated with @Edr
- Before and after any callback to an EJB
- After any exception is thrown

Note: Note that aspects are not applied outside the plug-in.

Table 8-1 Manual annotation for EDRs

Trigger	When	Modifiers restrictions	What is woven
method	before executing	public method only	only in methods annotated with @Edr
method	after executing	public method only	only in methods annotated with @Edr
method-call	before calling	any method	only for method call to a class implementing the PluginNorthCallback interface (EJB callback)
method-call	after calling	any method	only for method call to a class implementing the PluginNorthCallback interface (EJB callback)
exception	after throwing	any method	any exception thrown except in methods annotated with @NoEdr

The following values are always available in an EDR when it is generated from an aspect:

- class name
- method name
- direction of the request (south, north)
- position (before, after)
- interface (north, south, other, null)
- source (method, exception)

Exception scenarios

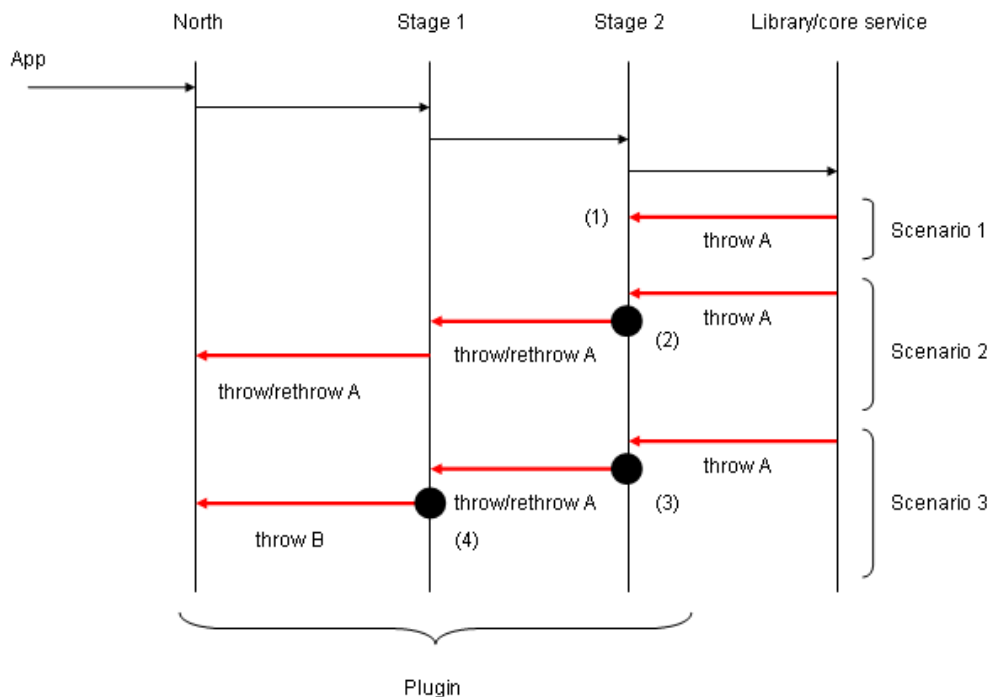
Exceptions are automatically woven by the aspect.

Some limitations apply:

- The aspect will catch only exceptions that are thrown by a plug-in method.
- The aspect will not catch an exception that is thrown by a library and caught by the plug-in.
- If the same exception is re-thrown several times, the aspect will only trigger an EDR once, for the first instance of the exception.

The diagram illustrates typical scenarios when a library (or core service) throws an exception in the plug-in.

Figure 8-1 Exception scenarios



Scenario 1:

The plug-in method in Stage 2 simply catches the exception but does not re-throw it or throw another exception. Since it just consume the exception, the aspect will not trigger an EDR.

Scenario 2:

The plug-in method in Stage 2 lets the exception A propagate (or re-throw exception A).

In this case, the aspect triggers an EDR after the method in stage 2. Since the same exception A (the same exception instance object) is propagated (or re-thrown), only the first method triggers an EDR.

Scenario 3:

This scenario is almost identical to scenario 2 except that the method in stage 1 is not throwing the exception A but another exception, named B. In this case, because B is not the same instance as A, the aspect will trigger another EDR after the method in stage 1.

Adding data to the RequestContext

In addition to the default values, an EDR also contains all the values put into the RequestContext using the `putEdr()` method.

Listing 8-3 Example to add values to and EDR using RequestContext

```
...
RequestContext ctx = RequestContextManager.getCurrent();
// this value will be part of any EDRs generated in the current request
ctx.putEdr("address", "tel:1234");
// this value will NOT be part of any EDRs since ctx.put(...) is used
ctx.put("foo", "bar");
...
```

Note: Common key names are defined in the class `com.bea.wlcp.wlmg.api.edr.EdrConstants`.

Using translators

When a parameter is a more complex object, it is possible to specify a translator that will take care of extracting the relevant information from this parameter.

The annotation is `@ContextTranslate`.

For example, the following method declares:

- The first (and only) parameter should be translated using the specified translator `ACContextTranslator`
- The returned object should also be translated using the specified translator `ACContextTranslator`

Listing 8-4 Using a translator

```
...  
  
@Edr  
  
public @ContextTranslate(ACContextTranslator.class) PlayTextMessageResponse  
playTextMessage(@ContextTranslate(ACContextTranslator.class) PlayTextMessage  
parameters) {  
  
    ...  
  
    return response;  
  
}  
  
...
```

The Translator is a class implementing the `ContextTranslator` interface.

Listing 8-5 Example of a Translator

```
public class ACContextTranslator implements ContextTranslator {  
  
    public void translate(Object param, ContextInfo info) {  
  
        if(param instanceof PlayTextMessage) {  
  
            PlayTextMessage msg = (PlayTextMessage) param;  
  

```

```

        info.put("address", msg.getAddress().toString());
    } else if(param instanceof PlayTextMessageResponse) {
        PlayTextMessageResponse response = (PlayTextMessageResponse) param;
        info.put("correlator", response.getResult());
    } ...
}
}

```

The ContextTranslator class specified in the @ContextTranslate annotation is automatically instantiated by the aspect when needed. It is however possible to explicitly register it using the ContextTranslatorManager.

Listing 8-6 Example of registering a context translator

```
ContextTranslatorManager.register(ACContextTranslator.class.getName(), new
ACContextTranslator());
```

Below is a summary of annotations to use.

Table 8-2 Annotations

Name	Type	Description
@ContextKey	Annotation	Specifies that an argument must be put into the current RequestContext under the name provided in this annotation
@ContextTranslate	Annotation	Same as @ContextKey but for complex argument that need to be translated using a translator (implementing the ContextTranslator interface).
ContextTranslator	Interface	Interface used by static translators to translate complex object.

Trigger an EDR programmatically

Network Gatekeeper triggers EDRs automatically in all plug-ins where aspects have been applied. It is also possible to trigger EDRs explicitly. In this case, you will have to manually create and trigger the EDR by following these steps:

1. Create an `EdrData` object
2. Trigger the EDR using the `EdrService` instance

Below is an example of triggering an EDR from inside a plug-in.

Listing 8-7 Trigger an EDR programmatically

```
public class SamplePlugin {

    // Get the EdrDataHelper like a logger

    private static final EdrDataHelper helper =
EdrDataHelper.getHelper(SamplePlugin.class);

    public void doSomething() {

        ...

        // Create a new EdrData using the EdrDataHelper class to allow
        // the WLNG to automatically populate some fields

        EdrData data = helper.createData();

        // Since we are creating the EdrData manually,
        // we have to provide the mandatory fields.

        // Note that the EdrDataHelper will provide most of them

        data.setValue(EdrConstants.FIELD_SOURCE,
EdrConstants.VALUE_SOURCE_METHOD);

        data.setValue(EdrConstants.FIELD_METHOD_NAME, "doSomething");

        // Log the EDR

        EdrServiceFactory.getService().logEdr(data);

        ...
    }
}
```

```
}  
  
}
```

EDR Content

The following table describes the content of an EDR. It describes which values are mandatory, who is responsible for providing these values, and other information.

Legends:

- A: Automatically provided by the WLNG
- H: Provided if the `EdrDataHelper createData` API is used to create the `EdrData` (which is the recommended way)
- M: Provided manually in the `EdrData`
- X: Provided in the EDR descriptor.
- C: Custom filter. Use the element `<attribute>` to specify a custom filter.

Note: EDRs triggered by aspects will have all the mandatory fields provided by the aspect.

Table 8-3 EDR content

Name	Description	Filter tag name
EdrId	<p>To get the ID, use <code>getIdentifier()</code> in <code>EdrConfigDescriptor</code>. This value is provided in the EDR descriptor.</p> <p>Provider INSIDE plug-in: X Provider OUTSIDE plug-in: X Mandatory: Yes</p>	C
ServiceName	<p>The name (or type) of the service.</p> <p>Fields in <code>EdrConstants</code>: <code>FIELD_SERVICE_NAME</code></p> <p>Provider INSIDE plug-in: H Provider OUTSIDE plug-in: M Mandatory: Yes</p>	C
ServerName	<p>The name of the Network Gatekeeper server.</p> <p>Fields in <code>EdrConstants</code>: <code>FIELD_SERVER_NAME</code></p> <p>Provider INSIDE plug-in: H Provider OUTSIDE plug-in: H Mandatory: Yes</p>	C
Timestamp	<p>The time at which the EDR was triggered (in ms since midnight, January 1, 1970 UTC)</p> <p>Fields in <code>EdrConstants</code>: <code>FIELD_TIMESTAMP</code></p> <p>Provider INSIDE plug-in: A Provider OUTSIDE plug-in: A Mandatory: Yes</p>	C

Table 8-3 EDR content

Name	Description	Filter tag name
ContainerTransaction Id	<p>The WebLogic Server transaction ID, if available.</p> <p>Fields in EdrConstants: FIELD_CONTAINER_TRANSACTION_ID</p> <p>Provider INSIDE plug-in: H</p> <p>Provider OUTSIDE plug-in: H</p> <p>Mandatory: No</p>	C
Class	<p>Name of the class that triggered the EDR.</p> <p>Fields in EdrConstants: FIELD_CLASS_NAME</p> <p>Provider INSIDE plug-in: H</p> <p>Provider OUTSIDE plug-in: H</p> <p>Mandatory: Yes</p>	<class>
Method	<p>Name of the method that triggered the EDR.</p> <p>Provider INSIDE plug-in: M</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: Yes</p>	<name> inside <method> or <method> inside <exception>
Source	<p>Indicates the type of source that triggered the EDR.</p> <p>Fields in EdrConstants: FIELD_SOURCE</p> <p>Values in EdrConstants: VALUE_SOURCE_METHOD, VALUE_SOURCE_EXCEPTION</p> <p>Provider INSIDE plug-in: M</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: Yes</p>	<method> or <exception>

Table 8-3 EDR content

Name	Description	Filter tag name
Direction	<p>Direction of the request.</p> <p>Fields in EdrConstants: FIELD_DIRECTION</p> <p>Values in EdrConstants: VALUE_DIRECTION_SOUTH, VALUE_DIRECTION_NORTH</p> <p>Provider INSIDE plug-in: M</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	<direction>
Position	<p>Position of the EDR relative to the method that triggered the EDR.</p> <p>Fields in EdrConstants: FIELD_POSITION</p> <p>Values in EdrConstants: VALUE_POSITION_BEFORE, VALUE_POSITION_AFTER</p> <p>Provider INSIDE plug-in: M</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	<position>
Interface	<p>Interface where the EDR is triggered.</p> <p>Fields in EdrConstants: FIELD_INTERFACE</p> <p>Values in EdrConstants: VALUE_INTERFACE_NORTH, VALUE_INTERFACE_SOUTH, VALUE_INTERFACE_OTHER</p> <p>Provider INSIDE plug-in: M</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	<interface>

Table 8-3 EDR content

Name	Description	Filter tag name
Exception	<p>Name of the exception that triggered the EDR.</p> <p>Fields in EdrConstants: FIELD_EXCEPTION_NAME</p> <p>Provider INSIDE plug-in: M</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	<p><name> inside</p> <p><exception></p>
SessionId	<p>Session ID.</p> <p>Fields in EdrConstants: FIELD_SESSION_ID</p> <p>Provider INSIDE plug-in: H</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	C
ServiceProviderId	<p>Service provider account ID.</p> <p>Fields in EdrConstants: FIELD_SP_ACCOUNT_ID</p> <p>Provider INSIDE plug-in: H</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	C
ApplicationId	<p>Application account ID.</p> <p>Fields in EdrConstants: FIELD_APP_ACCOUNT_ID</p> <p>Provider INSIDE plug-in: H</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	C

Table 8-3 EDR content

Name	Description	Filter tag name
AppInstanceId	Application instance ID. Fields in EdrConstants: FIELD_APP_INSTANCE_GROUP_ID Provider INSIDE plug-in: H Provider OUTSIDE plug-in: M Mandatory: No.	C
OrigAddress	The originating address with scheme included (for example “tel:1234”). Fields in EdrConstants: FIELD_ORIGINATING_ADDRESS Provider INSIDE plug-in: M Provider OUTSIDE plug-in: M Mandatory: No	C

Table 8-3 EDR content

Name	Description	Filter tag name
DestAddress	<p>The destination address(es) with scheme included (For example “tel:1234”). See Using send lists.</p> <p>Fields in EdrConstants: FIELD_DESTINATION_ADDRESS</p> <p>Provider INSIDE plug-in: M</p> <p>Provider OUTSIDE plug-in: M</p> <p>Mandatory: No</p>	C
<custom>	<p>Any additional information put into the current RequestContext using the putEdr() API will end up in the EDR.</p> <p>Fields in EdrConstants: -</p> <p>Provider INSIDE plug-in: -</p> <p>Provider OUTSIDE plug-in: -</p> <p>Mandatory: No</p>	C

Using send lists

If more than one address need to be stored in the DestAddress field, use the following pattern. Both patterns described below can be used.

Listing 8-8 Pattern to store one single or multiple addresses in field destination directly on EdrData.

```

EdrData data = ...;

// If there is only one address
data.setValue(EdrConstants.FIELD_DESTINATION_ADDRESS, address);

// If there are multiple addresses
data.setValues(EdrConstants.FIELD_DESTINATION_ADDRESS, addresses);

```

If you are using the current `RequestContext` object, simply store a `List` of addresses. The `EdrDataHelper` will automatically take care of converting this to a `List` of `Strings` in the `EdrData`.

Listing 8-9 Pattern to store one single or multiple addresses in field destination using `RequestContext`.

```
RequestContext ctx = RequestContextManager.getCurrent();

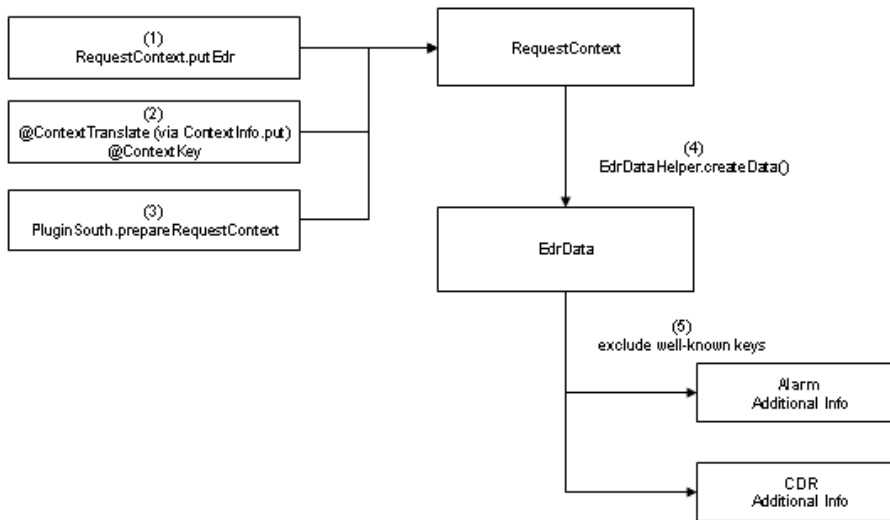
// If there is only one address
ctx.putEdr(EdrConstants.FIELD_DESTINATION_ADDRESS, address);

// If there are multiple addresses
URI[] addresses = ...;

ctx.putEdr(EdrConstants.FIELD_DESTINATION_ADDRESS, Arrays.asList(addresses));
```

RequestContext and EDR

The following diagram shows how and where information for the EDR is added to the `RequestContext` and how it finally ends up in the additional info column of the alarm and CDR databases.

Figure 8-2 RequestContext and EDR

There are 3 ways of putting information in the `RequestContext` that will end up in the EDR (more precisely in the `EdrData` object):

- Using the `putEdr()` API of the `RequestContext`
- Using the `@ContextKey` or `@ContextTranslate` annotation. In the case of the `@ContextTranslate` annotation, the information that will end up in the `RequestContext` will be what is put into the `ContextInfo` object.
- Any information put in the `RequestContext` parameter of the `PluginSouth.prepareRequestContext()` method.

When an EDR is created, the `EdrDataHelper` (which is the recommended way to create the EDR) will populate the `EdrData` with all the key/value pairs found in the `RequestContext`.

When the `EdrService` writes the alarm or CDR additional information content into the database, it will use all the `EdrData` key/value pairs EXCEPT a set of well-known keys that are either not relevant or already included in other columns of the database, see [“Alarm content” on page 8-39](#) and [“CDR content” on page 8-42](#).

Categorizing EDRs

Only one type of EDR exists: alarms and CDRs are subsets of this EDR type. In order to categorize the flow of EDRs as either a pure EDRs, an alarm or a CDR, the EDR service uses 3 descriptors:

- EDR descriptor contains descriptors that describe pure EDRs.
- Alarm descriptor contains descriptors that describe EDRs that should be considered alarms.
- CDR descriptor contains descriptors that describe EDRs that should be considered CDRs.

These XML descriptors can be manipulated using the **EDR Configuration Pane** as described in [Managing and Configuring EDRs, CDRs and Alarms](#) in the *System Administrator's Guide*. File representations of these must be included in `edrjmslistener.jar` if using external EDR listeners.

The EDR descriptor

Each descriptor contains a list of EDR descriptors that define an EDR as a pure-EDR, as an alarm or as a CDR.

Table 8-4 EDR descriptors.

Descriptor	Descriptor	Description
EDR	<edr...>	Defines which EDRs are pure EDRs
Alarm	<alarm...>	Defines which EDRs are alarms
CDR	<cdr...>	Defines which EDRs are CDRs

The descriptor is composed of two parts:

- The <filter> element: this is the filter
- The <data> element: this part is used to attach additional data with the EDR if it is matched by the <filter> part

The following table describes the elements allowed in the <filter> part:

Table 8-5 Elements allowed in <filter> part of an EDR descriptor.

Source	Filter	Min occurs	Max occurs	Description
<method>		0	unbounded	Filter EDR triggered by a method
	<name>	0	unbounded	Name of the method that triggered the EDR
	<class>	0	unbounded	Name of the class that triggered the EDR
	<direction>	0	2	Direction of the request
	<interface>	0	3	Interface where the EDR has been triggered
	<position>	0	2	Position relative to the method that triggered the EDR
<exception>		0	unbounded	Filter EDR triggered by an exception
	<name>	0	unbounded	Name of the exception that triggered the EDR
	<class>	0	unbounded	Name of the class where the exception was thrown
	<method>	0	unbounded	Name of the method where the exception was thrown
	<direction>	0	2	Direction of the request
	<interface>	0	3	Interface where the EDR has been triggered
	<position>	0	2	Position relative to the method that triggered the EDR
<attribute>		0	unbounded	Filter EDR by looking at custom attribute

Table 8-5 Elements allowed in <filter> part of an EDR descriptor.

Source	Filter	Min occurs	Max occurs	Description
	<key>	1	1	Name of the key
	<value>	1	1	Value

The following table describes the values allowed for each element of the <filter> part:

Table 8-6 Values allowed in each element of the <filter> part.

Source	Filter	Allowed values	Comment
<method>	<name>	“returntype nameofmethod([args])”	Method name. The arguments can be omitted with the parenthesis. See Special characters below.
	<class>	“fullnameofclass”	Fully qualified class name. See Special characters below.
	<direction>	“south”, “north”	
	<interface>	“north”, “south”, “other”	
	<position>	“before”, “after”	
<exception> >	<name>	“fullnameofexceptionclass”	Fully qualified exception class name. See Special characters below.
	<class>	“fullnameofclass”	Fully qualified class name where the exception was triggered. See Special characters below.
	<method>	“returntype nameofmethod([args])”	Method name. The arguments can be omitted with the parenthesis. See Special characters below.
	<direction>	“south”, “north”	

Table 8-6 Values allowed in each element of the <filter> part.

Source	Filter	Allowed values	Comment
	<interface>	“north”, “south”, “other”	
	<position>	“before”, “after”	
<attribute>	<key>	“astring”	
	<value>	“astring”	

Special characters

The filter uses special characters to indicate more precisely how to match certain values.

Using * at the end of a method, class or exception name matches all names that match the string specified prior to the * (that is, what the string starts with).

Note: The usage of any of these characters disables the caching of the filter containing them. To avoid a performance hit, using the other way of matching is strongly encouraged.

Table 8-7 Example filters

To match on	Use the filter
All sendInfoRes methods with one argument of type int.	<pre><method> <name>void sendInfoRes(int)</name> ... </method></pre>
All methods starting with sendInfoRes regardless of the arguments.	<pre><method> <name>void sendInfoRes</name> ... </method></pre>

Table 8-7 Example filters

To match on	Use the filter
All methods starting with void sendInfo.	<method> <name>void sendInfo*</name> ... </method>
All class names beginning with com.bea.wlcp.wlng.plugin	<method> <class>com.bea.wlcp.wlng.plugin*</class> ... </method>

Values provided

The exact value in these fields depends on who triggered the EDR. If the aspect triggered the EDR, then the name of the method (with return type and parameters) or the fully qualified name of the class/exception is indicated. If the EDR is manually triggered from the code, it is up to the implementer to decide what name to use. Here are some examples of fully qualified method/class names as specified by the aspect:

Example methods:

```
SendSmsResponse sendSms(SendSms)

void receivedMobileOriginatedSMS(NotificationInfo, boolean,
SmsMessageState, String, SmsNotificationRemote)

TpAppMultiPartyCallCallback reportNotification(TpMultiPartyCallIdentifier,
TpCallLegIdentifier[], TpCallNotificationInfo, int)
```

Example Class:

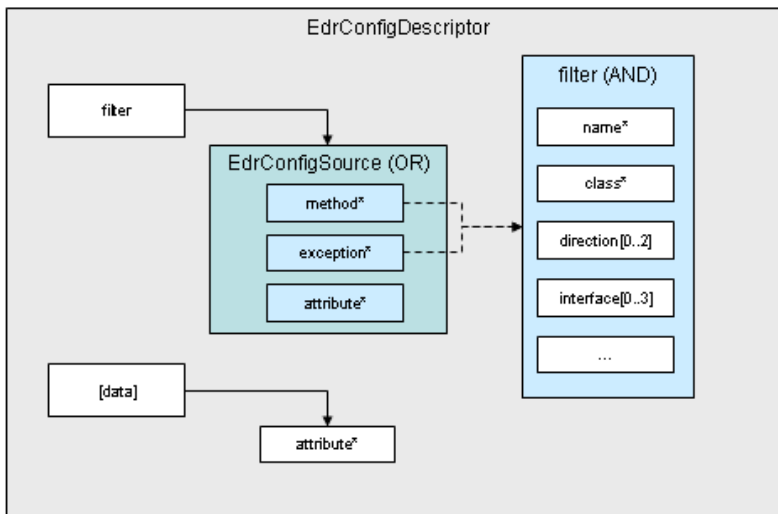
```
com.bea.wlcp.wlng.plugin.sms.smpp.SMPPManagedPluginImpl
```

Boolean semantic of the filters

The following diagram shows briefly how the filter works:

- The EdrConfigSource elements are the following: <method>, <exception> or <attribute>. They are combined using OR.
- The filter elements of each EdrConfigSource are combined using AND. However, if the same filter is available more than once (e.g. multiple class names), they are combined with OR.

Figure 8-3 Filter mechanism



Example filters

Example 1: filter

The following filter will categorize EDRs as pure EDRs with an id of 1000 when the following conditions are met:

- The class where the method triggered the EDR is `com.bea.wlcp.wlng.plugin.AudioCallPlugin` or any subclass of it.
- AND the request is southbound (direction = south)
- AND the interface where the EDR was trigger is north
- AND the EDR has been triggered after the method has been executed (position = after)

Listing 8-10 Example 1: filter

```
<edr id="1000" description="...">
  <filter>
    <method>
      <class>com.bea.wlcp.wlng.plugin.AudioCallPlugin</class>
      <direction>south</direction>
      <interface>north</interface>
      <position>after</position>
    </method>
  </filter>
</edr>
```

Example 2: Alarm filter

The following filter will categorize EDRs into alarms when the following conditions are met:

- The exception is the class `com.bea.wlcp.wlng.plugin.PluginException` or a subclass of it.
- OR the name of the exception starts with `org.csapi.*`. Since “*” is used, the matching will not be performed using the class hierarchy but only using a pure string matching.

The alarms descriptor has a `<alarm-group>` element that is used to group alarms by service/source: this group id and each individual alarm id is used to generate the OID of SNMP traps.

Listing 8-11 Example 3: filter

```
<alarm-group id="104" name="parlayX" description="Parlay X alarms">>
<alarm id="1000" severity="minor" description="Parlay X exception">
  <filter>
    <exception>
      <name>com.bea.wlcp.wlng.plugin.PluginException</name>
```

```

        <name>org.csapi*</name>

    </exception>

</filter>

</alarm>

</alarm-group>

```

Example 3: Alarm filter

The following filter will categorize EDRs into alarms when the following conditions are met:

- The exception is the class `com.bea.wlcp.wlng.plugin.PluginException` or a subclass of it
- OR the name of the exception starts with “org.csapi”. String matching is used.
- AND the exception was triggered in a class whose name starts with `com.bea.wlcp.wlng.plugin`
- AND the request is northbound (`direction = north`) when the exception was triggered

If the filter determines that the EDR is an alarm, the following attributes are available to the alarm listener (they are defined in the `<data>` part):

- `identifier = 123`
- `source = wlng_nt1`

Listing 8-12 Example 3: filter

```

<alarm id="1000" severity="minor" description="Parlay X exception">
  <filter>
    <exception>
      <name>com.bea.wlcp.wlng.plugin.PluginException</name>
      <name>org.csapi*</name>
      <class>com.bea.wlcp.wlng.plugin*</class>
      <direction>north</direction>
    </exception>
  </filter>
</alarm>

```

```
</exception>
</filter>
<data>
  <attribute key="identifier" value="123"/>
  <attribute key="source" value="wlng_nt1"/>
</data>
</alarm>
```

Example 4: filter

The following filter (for example purposes only) will categorize EDRs into pure EDRs with the id 1002 when the following conditions are met:

- The name of the method that triggered the EDR starts with “void play” AND the class is com.bea.wlcp.wlng.plugin.AudioCallPluginNorth or a subclass of it AND the EDR was triggered after executing this method.
- OR the name of the method that triggered the EDR is “String getMessageStatus” AND the class is 'com.bea.wlcp.wlng.plugin.AudioCallPluginNorth' or a subclass of it AND the EDR was triggered before executing this method.
- OR the name of the exception that triggered the EDR starts with com.bea.wlcp.wlng.bar AND the exception was triggered in a plug-in north interface
- OR the name of the exception that triggered the EDR starts with com.bea.wlcp.wlng.plugin.exceptionA AND the exception was triggered in a class whose name starts with com.bea.wlcp.wlng.plugin.classD AND the exception was triggered in a method whose name starts with void com.bea.wlcp.wlng.plugin.methodA AND the exception was triggered in a plugin north interface
- OR the EDR contains an attribute with key attribute_a and value value_a
- OR the EDR contains an attribute with key attribute_b and value value_b

Listing 8-13 Example 4: filter

```
<edr id="1002">
```



```

<filter>
  <method>
    <name>void play*</name>
    <class>com.bea.wlcp.wlng.plugin.AudioCallPluginNorth</class>
    <position>after</position>
  </method>
  <method>
    <name>String getMessageStatus</name>
    <class>com.bea.wlcp.wlng.plugin.AudioCallPluginNorth</class>
    <position>before</position>
  </method>
  <exception>
    <name>com.bea.wlcp.wlng.bar*</name>
    <interface>north</interface>
  </exception>
  <exception>
    <name>com.bea.wlcp.wlng.plugin.exceptionA</name>
    <class>com.bea.wlcp.wlng.plugin.classD</class>
    <method>void com.bea.wlcp.wlng.plugin.methodA</method>
    <interface>north</interface>
  </exception>
  <attribute key="attribute_a" value="value_a"/>
  <attribute key="attribute_b" value="value_b"/>
</filter>
</edr>

```

Example 5: filter with corresponding code for manually triggering a matching EDR

The following example shows a manually triggered EDR with its corresponding filter. The EDR is triggered using these lines.

Listing 8-14 Example 5: Trigger the EDR

```
// Declare the EdrDataHelper for each class

private static final EdrDataHelper helper =
EdrDataHelper.getHelper(MyClass.class);

public void myMethodName() {

    ...

    // Create a new EdrData. Use the EdrDataHelper class to allow the WLNG to
    automatically populate some fields

    EdrData data = helper.createData();

    // Because we are creating the EdrData manually, we have to provide the
    mandatory fields

    data.setValue(EdrConstants.FIELD_SOURCE, EdrConstants.VALUE_SOURCE_METHOD);
    data.setValue(EdrConstants.FIELD_METHOD_NAME, "myMethodName");
    data.setValue("myKey", "myValue");

    // Log the EDR

    EdrManager.getInstance().logEdr(data);

    ...
}
```

This EDR can be filtered using the following filter (note the various way of identifying this EDR):

Listing 8-15 Example: Filter 5

```

<edr id="1003">
  <filter>
    <!-- Match both method name and class name -->
    <method>
      <name>myMethodName</name>
      <class>com.bea.wlcp.wlng.myClassName</class>
    </method>
    <!-- OR match only the method name (looser than matching also the class
name) -->
    <method>
      <name>myMethodName</name>
    </method>
    <!-- OR match only the classname (looser than matching also the method
name) -->
    <method>
      <class>com.bea.wlcp.wlng.myClassName</class>
    </method>
    <!-- OR match only the custom attribute -->
    <attribute key="myKey" value="myValue"/>
  </filter>
</edr>

```

Check-list for EDR generation

Below is a list of steps to make your plug-in able to take advantage of the aspect EDR:

- Make sure to register all your PluginNorth (and south) objects within the ManagedPlugin before registering in the PluginManager.

- Annotate all the methods you want to be woven using the `@Edr` annotation.
- Annotate the specific arguments you want to see in the EDR for each annotated methods. Use either `@ContextKey` or `@ContextTranslate` depending on the kind of argument.
- Add to the EDR descriptor all the EDRs you are triggering, either manually or with the `@Edr` annotation. This is the only way to customize alarms and CDRs.
- If external EDR listeners, CDR, and alarms are used, the file `edrjmslistener.jar` needs to be updated on all the listeners. Add the contents of the EDR descriptors to `edr.xml`, CDR descriptor to `cdr.xml`, and alarm descriptor to `alarm.xml`. The xml files resides in the directory `edr` in `edrjmslistener.jar`.

Frequently Asked Questions about EDRs and EDR filters

Question: Is it possible to specify both exception and method name in the filter section?

Listing 8-16 Example: method name and exception in a filter.

```
<filter>
    <method>
        <name>internalSendSms</name>
    </method>
    <exception>

<name>com.bea.wlcp.wlng.plugin.sms.smpp.TooManyAddressesException</name>
    </exception>
</filter>
```

Answer

Yes, make sure that the `<method>` element is before the `<exception>` element. Otherwise the XSD will complain.

Q: Is it possible to specify multiple method names?

Answer

Yes.

Q: In some places I have methods re-throwing an exception. Is it possible to have only one of the methods generate the EDR and map that edr to an alarm?

Listing 8-17 Re-throwing an exception

```
myMethodA() throws MyException{
    myMethodB();
}

myMethodB() throws MyException{
    myMethodC();
}

myMethodC() throws MyException{
    ...
    //on error
    throw new MyException("Exception text..");
}
```

Answer

In this case, only the first exception will be caught by aspect. Or more precisely, they will all be caught by aspect but will only trigger an EDR for the first one, but not for the re-thrown ones (if they are the same, of course). So you don't need to use the `@NoEdr` annotation for `myMethodA` and `myMethodB`.

Q: Will aspect detect the following exception?

Listing 8-18 Example exception

```
try{
    throw new ReceiverConnectionFailureException(message);
}catch(ReceiverConnectionFailureException connfail){
    //EDR-ALARM-MAPPING
}
```

Answer

This exception will NOT be detected by the aspects. If you need to generate an EDR you will have to either manually create an EDR or call a method throwing an exception.

Q: Will EDRs for exceptions also work for private methods?

Answer

Yes, EDRs can work for any method.

Q: Will exceptions be disabled with the @NoEdr annotation?

Answer

Yes, with the @NoEdr annotation you will not get any EDRs, not even for exceptions.

Q: How can data from the current context be included in an alarm?

Can an alarm be generated in a request with more than 12 destination addresses? How can information be added to the alarm about how many addresses that were included in the request?

It is possible to specify some info in the alarm descriptor with something like

```
<data>
    <attribute key="source" value="thesource"/>
</data>
```

. Can something be put in the RequestContext using the putEdr method and then get it into the alarm in some way?

Answer

Yes, add custom information by putting this information into the current RequestContext, as show below.

```
RequestContext ctx = RequestContextManager.getCurrent();
ctx.putEdr("address", "tel:1234");
```

This value is part of any EDRs generated in the current request.

The information will be available in the database in the additional_info column. Make sure you are putting in only relevant information.

Q: Is it possible to specify classname in the filtering section?

Answer

Yes, use the <class> tag inside <method> or <exception> in the filter.

```
<filter>
  <exception>
    <class>com.y.y.z.MyClass</class>
    <name>com.x.y.z.MyException</name>
  </exception>
</filter>
```

Alarm generation

An alarm is a subset of an EDR. To generate an alarm, generate an EDR, either using one generated in aspects or programmatically and define the ID, and the descriptor of the alarm in the alarm descriptor.

The alarm ID, severity, description and other kind of attributes are defined in the alarm descriptor, see [“The EDR descriptor” on page 8-22](#). For extensions, the alarm ID shall be in the range 500 000 to 999 999.

Note: The alarm filter that provides the *first* match in the alarm descriptor is used for triggering the alarm.

There are two ways to trigger an alarm:

- Use an existing EDR that is generated in the plug-in and add its descriptor to the alarm descriptor.
- Programmatically trigger an EDR and add its descriptor in both the alarm descriptor file and the EDR descriptor. Make sure the ID of the alarm is unique and that the description is the same as in the EDR descriptor.

Trigger an alarm programmatically

Trigger an EDR as described in [“EDR Content” on page 8-13](#). Then specify in the alarm descriptor the corresponding alarms.

Listing 8-19 Example code to trigger an alarm

```
private static final EdrDataHelper helper =
    EdrDataHelper.getHelper(MyClass.class);

...

EdrData data = helper.createData();

data.setValue(EdrConstants.FIELD_SOURCE, EdrConstants.VALUE_SOURCE_METHOD);

data.setValue(EdrConstants.FIELD_METHOD_NAME, "com.bea.wlcp.wlng.myMethod");

data.setValue("myAdditionalInformation", ...);

EdrManager.getInstance().logEdr(data);

...
```

The corresponding entry in the alarm descriptor that matches this EDR is shown below.

Listing 8-20 Alarm descriptor

```
<alarm id="2006"
      severity="major"
      description="Sample alarm">
  <filter>
```



```

<method>
  <name>com.bea.wlcp.wlng.myMethod</name>
  <class>com.bea.wlcp.wlng.myClass</class>
</method>
</filter>
</alarm>

```

Alarm content

Below is a list of the information provided in alarms.

Table 8-8 Alarm information for alarm listeners, also stored in DB

Field	Comment
alarm_id	Unique ID for the alarm. Automatically provided by the EdrService.
source	Service name emitting the alarm. Automatically provided by the EdrService.
timestamp	Timestamp in milliseconds since midnight, January 1, 1970 UTC. Automatically provided by the EdrService.
severity	Severity level. Defined in the alarm. descriptor.
identifier	The alarm identifier. Defined in the alarm descriptor. The column in the database will always contain the identifier defined in the alarm descriptor.

Table 8-8 Alarm information for alarm listeners, also stored in DB

Field	Comment
alarm_info	<p>The alarm information or description.</p> <p>Defined in the alarm descriptor.</p>
additional_info	<p>Automatically provided by the EdrService.</p> <p>Not valid for backwards compatible alarm listeners.</p> <p>The format of this field was changed when Network Gatekeeper 3.0 was introduced. Each entry is formatted as:</p> <p>key=value\n</p> <p>Similar to the Java properties file.</p> <p>All the custom key/value pairs found in the EdrData except these are present (EdrConstants if not specified):</p> <ul style="list-style-type: none"> • FIELD_TIMESTAMP • FIELD_SERVICE_NAME • FIELD_CLASS_NAME • FIELD_METHOD_NAME • FIELD_SOURCE • FIELD_DIRECTION • FIELD_POSITION • FIELD_INTERFACE • FIELD_EXCEPTION_NAME • FIELD_ORIGINATING_ADDRESS • FIELD_DESTINATION_ADDRESS • FIELD_CONTAINER_TRANSACTION_ID • FIELD_CORRELATOR • FIELD_SESSION_ID • FIELD_SERVER_NAME • ExternalInvokerFactory.SERVICE_CORRELATION_ID • FIELD_BC_EDR_ID • FIELD_BC_EDR_ID_3 • FIELD_BC_ALARM_IDENTIFIER • FIELD_BC_ALARM_INFO

CDR generation

A CDR is a subset of an EDR. To generate a CDR, generate an EDR and define the ID of the EDR in the CDR descriptor.

Triggering a CDR

There are two ways to trigger a CDR:

- Use an existing EDR that is generated in the plug-in and add its description to the CDR descriptor.
- Programmatically trigger an EDR and add its description to the CDR descriptor.

Trigger a CDR programmatically

If none of the existing EDR is appropriate for a CDR, you can programmatically trigger an EDR that will become a CDR. See the section, [“Trigger an EDR programmatically” on page 8-12](#) for information on how to create and trigger an EDR. Specify in the CDR descriptor the description necessary for this EDR to be considered a CDR.

Listing 8-21 Example, triggering a CDR

```
private static final EdrDataHelper helper =
    EdrDataHelper.getHelper(MyClass.class);

...

EdrData data = helper.createData();

data.setValue(EdrConstants.FIELD_SOURCE, EdrConstants.VALUE_SOURCE_METHOD);

data.setValue(EdrConstants.FIELD_METHOD_NAME,
    "com.bea.wlcp.wlng.myEndOfRequestMethod");

// Fill the required fields for a CDR

data.setValue(EdrConstants.FIELD_CDR_START_OF_USAGE, ...);

...

EdrManager.getInstance().logEdr(data);
```

...

The description, in the CDR descriptor, that matches this EDR is shown below.

Listing 8-22 Filter to match the EDR

```
<cdr>
  <filter>
    <method>
      <name>com.bea.wlcp.wlng.myEndOfRequestMethod</name>
      <class>com.bea.wlcp.wlng.myClass</class>
    </method>
  </filter>
</cdr>
```

CDR content

In addition to the EDR fields, the following table lists the specific fields used only for CDRs.

Table 8-9 Fields in EdrConstants specific for CDRs.

Field in EdrConstants	Comment
FIELD_CDR_SESSION_ID	
FIELD_CDR_START_OF_USAGE	
FIELD_CDR_CONNECT_TIME	
FIELD_CDR_END_OF_USAGE	
FIELD_CDR_DURATION_OF_USAGE	
FIELD_CDR_AMOUNT_OF_USAGE	

Table 8-9 Fields in EdrConstants specific for CDRs.

Field in EdrConstants	Comment
FIELD_CDR_ORIGINATING_PARTY	
FIELD_CDR_DESTINATION_PARTY	Same pattern applies as for send lists, see “Using send lists” on page 8-19.
FIELD_CDR_CHARGING_INFO	

The CDR content is aligned toward the 3GPP Charging Applications specifications. As a result the database schema has been changed to accommodate these ends and to facilitate future extensions.

Legends:

- NU: Not used
- NC: New column in DB
- RC: Renamed column in DB

Table 8-10 Content in database

Field	Comment	DB
transaction_id	Unique id for the CDR. Provided automatically by the EDR service.	x
service_name	name of the service Provided automatically by the EDR service.	x
service_provider	the service provider account ID Provided automatically by the EDR service.	x
application_id	the application account ID (was user_id in 2.2)	RC
application_instance_group_id	the application instance ID.	NC
container_transaction_id	id of the current user transaction Provided automatically by the EDR service.	NC

Table 8-10 Content in database

Field	Comment	DB
server_name	name of the server that generated the CDR. Provided automatically by the EDR service.	NC
timestamp	in ms since midnight, January 1, 1970 UTC	NC
service_correlation_id	Service Correlation ID. Provided automatically by the EDR service.	NC
charging_session_id	Id that correlates requests that belong to one charging session as defined by the plug-in. Was 'session_id' in 2.2. Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR.	x
start_of_usage	The date and time the service capability module started to use services in the network (in ms since midnight, January 1, 1970 UTC) Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR.	x
connect_time	The date and time the destination party responded (in ms since midnight, January 1, 1970 UTC). Used for call control only. Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR.	x
end_of_usage	The date and time the service capability module stopped using services in the network (in ms since midnight, January 1, 1970 UTC). Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR	x

Table 8-10 Content in database

Field	Comment	DB
duration_of_usage	The total time the service capability module used the network services (in ms) Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR	x
amount_of_usage	Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR.	x
originating_party	The originating party address with scheme included (e.g. "tel:1234") Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR.	x
destination_party	the originating party address with scheme included (e.g. "tel:1234"). Additional addresses are stored in the additional_info field.	x
charging_info	The charging service code from the application. Plug-in specific. Plug-in needs to put the value into the RequestContext of the request that will trigger the CDR.	x
additional_info	Additional information provided by the plug-in	x
revenue_share_percentage	Not used.	NU
party_to_charge	Not used.	NU
slee_instance	Not used.	NU
network_transaction_id	Not used.	NU
network_plugin_id	Not used.	NU
transaction_part_number	Not used.	NU
completion_status	Not used.	NU

Additional_info column

The EDR populates the additional_info column of the DB with all the custom key/value pairs found in the EdrData except the ones listed below.

Excluded keys (EdrConstants if not specified):

- FIELD_SERVICE_NAME
- FIELD_APP_INSTANCE_GROUP_ID
- FIELD_SP_ACCOUNT_ID
- FIELD_CONTAINER_TRANSACTION_ID
- FIELD_SERVER_NAME
- FIELD_TIMESTAMP
- ExternalInvokerFactory.SERVICE_CORRELATION_ID
- FIELD_CDR_SESSION_ID
- FIELD_CDR_START_OF_USAGE
- FIELD_CDR_CONNECT_TIME
- FIELD_CDR_END_OF_USAGE
- FIELD_CDR_DURATION_OF_USAGE
- FIELD_CDR_AMOUNT_OF_USAGE
- FIELD_CDR_ORIGINATING_PARTY
- FIELD_CDR_DESTINATION_PARTY
- FIELD_CDR_CHARGING_INFO
- FIELD_CLASS_NAME
- FIELD_METHOD_NAME
- FIELD_SOURCE
- FIELD_DIRECTION
- FIELD_POSITION

- FIELD_INTERFACE
- FIELD_EXCEPTION_NAME
- FIELD_ORIGINATING_ADDRESS
- FIELD_DESTINATION_ADDRESS
- FIELD_CORRELATOR
- FIELD_APP_ACCOUNT_ID
- FIELD_SESSION_ID
- FIELD_BC_EDR_ID
- FIELD_BC_EDR_ID_3
- FIELD_BC_ALARM_IDENTIFIER
- FIELD_BC_ALARM_INFO

Two keys not present in the EdrData are added to additional_info.

Table 8-11 Keys not present in EdrData, but added in additional_info

Key	Description
destinationParty	If a send list is specified as the destination party, the first address will be written in the destination_party field of the DB and the remainder of the list will be written under this key name
oldInfo	Any backwards compatible additional info is available

The format of the additional_info field is formatted as:

```
key=value\n
```

similar to the Java properties file.

Out-of-the box (OOTB) CDR support

It is difficult to come up with a CDR generation scheme which will fulfill the requirements of all customers. Network Gatekeeper generates a default set of CDRs which can be customized by re-configuring the CDR descriptor.

The guiding principle for deciding when to generate CDRs is:

- Generate a CDR when you are 100% sure that you have completely handled the service request

In other words, after the last method, in a potential sequence of method calls, returns.

For network-triggered requests this means that you should a CDR at the south interface after the method has returned back to the network. For application-triggered requests generate a CDR at the north interface after the method has returned to the Network Tier SLSE.

Extending Statistics

Aspects are also used to generate statistics. To add a new statistic type to your Communication Service requires two steps:

1. You must add a new statistic type, using the `addStatisticType` operation in the Management Console. For more information, see “Managing and Configuring Statistics and Transaction Licenses” in the *System Administration Guide*.
2. The statistics aspect is automatically applied to all public methods at `PluginNorth`. By default extension Communication Services generate information identified with the transaction type `TRANSACTION_TYPE_EXTENSION`. To generate more specific types, annotate your code with `@Statistics(id=<My_Statistics_Type>)`

For extensions, the statistics ID shall be in the range 1000 to 2250.

Making Communication Services Manageable

Once you have created your extension Communication Service, any OAM functions that you have designed - read/write attributes and/or operations - must be exposed in a way that allows them to be accessed and manipulated, either through the Network Gatekeeper Console extension, or through other management tools. The following chapter provides a description of the mechanism that Network Gatekeeper uses to accomplish this.

Overview

WebLogic Network Gatekeeper uses the Java Management Extensions (JMX) 1.2 standard, as it is implemented in JDK 1.5. The JMX model consists of three layers, Instrumentation, Agent, and Distributed Services. As an Communication Service developer, you work in the Instrumentation layer. You create managed beans (MBeans) that expose your Communication Service management functionality as a management interface. These MBeans are then registered with the Agent, the Runtime MBean Server in the WebLogic Server instance, which makes the functionality available to the Distributed Services layer, management tools like the Network Gatekeeper Management Console. Finally, because configuration information needs to be persisted, you store the values you set using Network Gatekeeper's Configuration Store, which provides a write-through database cache. In addition to persisting the configuration information, the cache also provides cluster-wide access to the data, updating a cluster-wide store whenever there is a change in globally relevant configuration data.

For more information on the JMX model in general in relation to WebLogic Server, see [Developing Manageable Applications with JMX](#).

Create Standard JMX MBeans

Creating standard MBeans is a three step process.

1. [Create an MBean Interface](#)
2. [Implement the MBean](#)
3. [Register the MBean with the Runtime MBean Server](#)

Configuration settings should be persisted, see [Use the Configuration Store to Persist Values](#).

Create an MBean Interface

The first thing you need to do is to create an interface file that describes getter and setter methods for each class attribute that is to be exposed through JMX (getter only for read-only attributes; setter only for write-only) and a wrapper operation for each class method to be exposed. The attribute names should be the case-sensitive names that you wish to see displayed in the UI of the Console extension.

- For each read-write attribute define a *get* and *set* method that follows this naming pattern: *get*<Attribute name>, *set*<Attribute name> where <Attribute name> is a case-sensitive name that you want to expose to JMX clients.
- For each read-only attribute define only an *is* or a *get* method. For each write-only attribute, define only a *set* method.
- The Javadoc will be rendered in the console as a description of an attribute or operation. It will render exactly as in the Javadoc. For example:

```
/**
 * Connects to the simulator
 * @throws ManagementException An exception if the connection failed
 */
public void connect() throws ManagementException;
```

Will render as:

Operations

Select An Operation:

Connects to the simulator
@throws ManagementException An exception if the connection failed

- Any internal operation or attribute should be annotated with `@Internal` annotation. This attribute or method will not be shown in the console. Example:

```
@Internal
```

```
public String resetStatistics();
```

- Indicate optional parameters for the operation by `@OptionalParam` annotation. In the JavaDoc for the operation, explicitly specify which parameters are optional. Example:

```
/**
 * Gets the alarms matching the specified criteria from the database
 * @param Identifier EDR Identifier
 * @Param Source server name (optional)
 * @Param Severity 0 - Critical, 1- Major, 2 -Minor
 * @Param maxEntries max number of entries
 */
AlarmData[] getAlarms(long identifier,
                      @OptionalParam('source')String source,
                      int severity,
                      int maxEntries) throws ManagementException;
```

The interface should be named `<ServiceName>MBean.java`. The interface for the example Communication Service provided with the Platform Development Studio is named `ExampleMBean.java`.

Implement the MBean

Once you have defined the interface, it must be implemented.

You must name your class `<ServiceName>MBeanImpl.java`, based on the interface name. The implementation must extend `WLNGMBeanDelegate`. This class takes care of setting up notifications and MBean descriptions and all MBean implementation classes must extend it. All MBean implementations must also be public, non-abstract classes and have at least one public constructor. The MBean implementation for the example Communication Service provided with the Platform Development Studio is named `ExampleMBeanImpl.java`.

- MBean implementation must be a public, non abstract class
- MBean must have at least one public constructor
- MBean must implement its corresponding MBean interface and extend `WLNGMBeanDelegate`

Register the MBean with the Runtime MBean Server

The MBean must be registered with the Runtime MBean Server in the local WebLogic Server instance. Network Gatekeeper provides a proxy class for MBean registration:

```
com.bea.wlcp.wlng.api.management.MBeanManager
```

The MBean implementation is registered using an `ObjectName`, and a `DisplayName`:

```
registerMBean(Object mBeanImpl, ObjectName objectName, String displayName)
```

Construct the `ObjectName` using:

```
constructObjectName(String type, String instanceName, HashMap properties)
```

There should be no spaces in the `InstanceName` or `Type`. Object names are case-sensitive

If the MBean is a regular MBean, use the conventions as illustrated in [Table 9-1](#)

Table 9-1 MBean ObjectName

The ObjectName convention for extensions	
type	Fully qualified MBean Name. <MBeanObj>.class.getName()
instanceName	Unique name that identifies the instance of the MBean. For example, it can be obtained from serviceContext.getName() The unique name of the MBean. If this is a plug-in that potentially is used on the same server with multiple plug-in instances this should be unique per plug-in instance. It is recommended to use managedPlugin.getId().
properties	HashMap that contains objectName key and value pairs ObjectNameConstants class has set of constants that can be used as keys. Null for non-hierarchical MBeans.

Example:

```
com.bea.wlcp.wlng:AppName= wlng_nt_sms_px21#4.0,InstanceName=
Plugin_px21_short_messaging_smpp,
Type=com.bea.wlcp.wlng.plugin.sms.smpp.management.SmsMBean
```

If the MBean is a MBean that should be the child of a regular MBean, use the conventions as illustrated in [Table 9-2](#)

Table 9-2 MBean ObjectName with hierarchy

The ObjectName convention for extensions	
type	Fully qualified MBean Name of the parent MBean. <Parent MBeanObj>.class.getName()
instanceName	Unique name that identifies the instance of the parent MBean.
properties.key=ObjectNameConstants.LEVEL1_INSTANCE_NAME	properties.value is a unique name that identifies the instance of the MBean

Table 9-2 MBean ObjectName with hierarchy

The ObjectName convention for extensions	
properties.key=ObjectNameConstants.LEVEL1_TYPE	Fully qualified MBean Name: <MBeanObj>.class.getName()
properties.key=ObjectNameConstants.LEVEL2_INSTANCE_NAME	properties.value is a unique name that identifies the instance of the MBean
properties.key=ObjectNameConstants.LEVEL2_TYPE	Fully qualified MBean Name: <MBeanObj>.class.getName()

Example:

A child MBean, for example HeartBeatConfiguration, can register with the same Level1InstanceName for all instances of the Plug-in (since it is a child, its MBean name depends on the parent’s instance:

```
com.bea.wlcp.wlng:AppName= wlng_nt_sms_px21#4.0,InstanceName=
Plugin_px21_short_messaging_smpp,
Type=com.bea.wlcp.wlng.plugin.sms.smpp.management.SmsMBean,Level1InstanceName=HeartBeatManager,Level1Type=com.bea.wlcp.wlng.heartbeat.management.HeartbeatMBean

com.bea.wlcp.wlng:AppName= wlng_nt_multimedia_messaging_px21#4.0,InstanceName
Plugin_px21_multimedia_messaging_mm7, Type=
com.bea.wlcp.wlng.plugin.multimediamessaging.mm7.management.MessagingManagementMBean,Level1InstanceName=HeartBeatManager,Level1Type=com.bea.wlcp.wlng.heartbeat.man
agement.HeartbeatMBean
```

Use the Configuration Store to Persist Values

The Network Gatekeeper Configuration Store API provides a cluster-aware write-through database cache. Parameters stored in the Configuration Store are both cached in memory and written to the database. The store works in two modes: Local and Global. Values stored in the Local store are of interest only to a single server instance, whereas values stored in the Global store are of interest to all servers cluster-wide. Updates to a value in the Global store update all cluster nodes. The example Communication Service provides a handler class,

`ConfigurationStoreHandler`, that gives an example of both usages of the Configuration Store API.

Note: The configuration store supports only Boolean, Integer, Long, and String values.

Making Communication Services Manageable

Using the Platform Test Environment

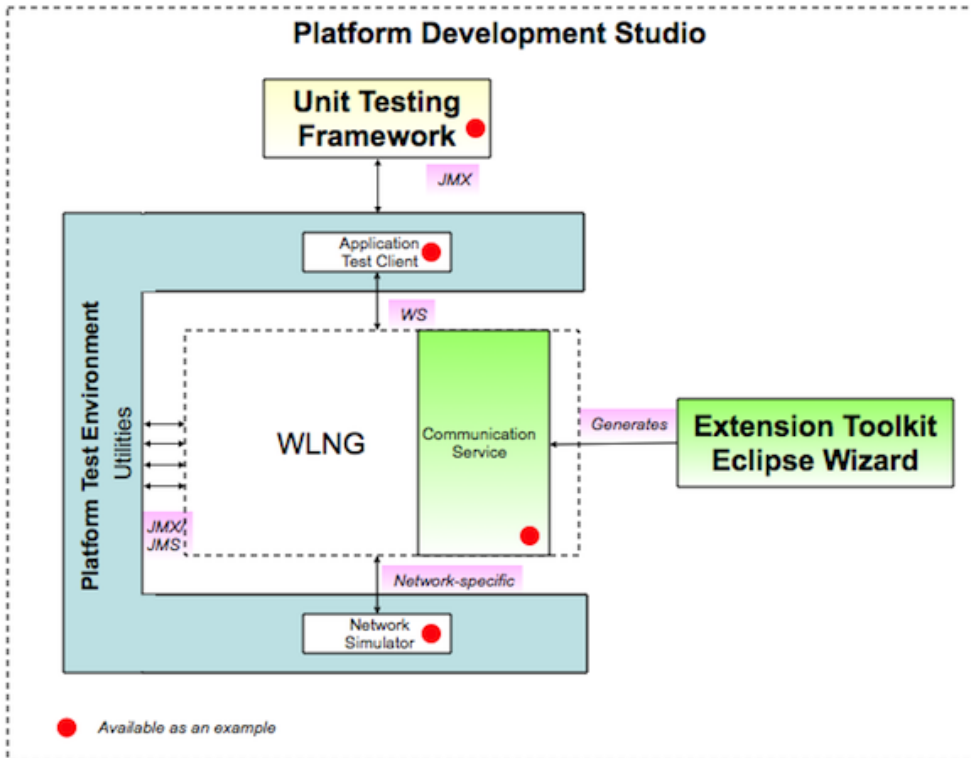
Testing is a key part of the development cycle. Network Gatekeeper provides an entire suite of testing tools to help you develop your extensions quickly and efficiently. This chapter introduces the Platform Test Environment (PTE). It consists of:

- [Overview](#)
- [Installing and Running the Platform Test Environment](#)
- [Navigating the Platform Test Environment GUI](#)
- [Extending the Platform Test Environment](#)
- [Using the Unit Test Framework \(UTFW\) with the Platform Test Environment](#)

Overview

The Platform Test Environment is a key part of the Platform Development Studio.

Figure 10-1 The Platform Test Environment in Context



The Platform Test Environment is a flexible, powerful tool, consisting of:

- Application service test clients for most out-of-the-box communication services
- PRM test clients for many operations covered by the Partner Relationship Management interfaces
- Network simulators for most node types supported by the out-of-the-box communication services
- Dual mode support:
 - Standalone with a Java Swing-based GUI

- Console, in which the PTE's functionality can be accessed using JMX, as, for example, from a unit test
- MBean browser for performing Network Gatekeeper management tasks
- A JMS-based EDR/CDR/Alarm listener
- JNDI browser
- Database browser for interacting with the database
- Real-time duration test graphing
- SLA editor
- Embedded TCP Monitor
- Easily extendable architecture
 - An example application test client for use with the example communication service
 - An example network simulator for use with the example communication service
 - A set of SPIs that allows your modules to interact with the PTE
- A framework for building unit tests, including:
 - A base test class, derived from JUnit
 - Mechanisms that simplify connecting to the Platform Testing Environment
 - An example test case for use with the example communication service

Installing and Running the Platform Test Environment

The Platform Test Environment is automatically installed when you install Network Gatekeeper. In standard installations, it is found in the `<bea_home>/wlng_pds400/pte` directory. Before you use the PTE, you must have:

- Installed WLNG 4.0
- Used the `setWLSEnv` script in `<bea_home>/wlng400/server/bin` or set the equivalent path so that you have access to the Ant 1.6.5 installation that comes with WLS.
- Created a `driver` directory under `<bea_home>/wlng_pds400/pte/lib` and copied your JDBC driver into it.

To start the PTE in GUI mode, either type `'ant run'` in a command window, or, if you are using Windows, double-click the `run.cmd` file in the PTE directory.

To start the PTE in console mode, type `'ant console'` in a command window, or, if you are using Windows, double-click the `console.cmd` file in the PTE directory.

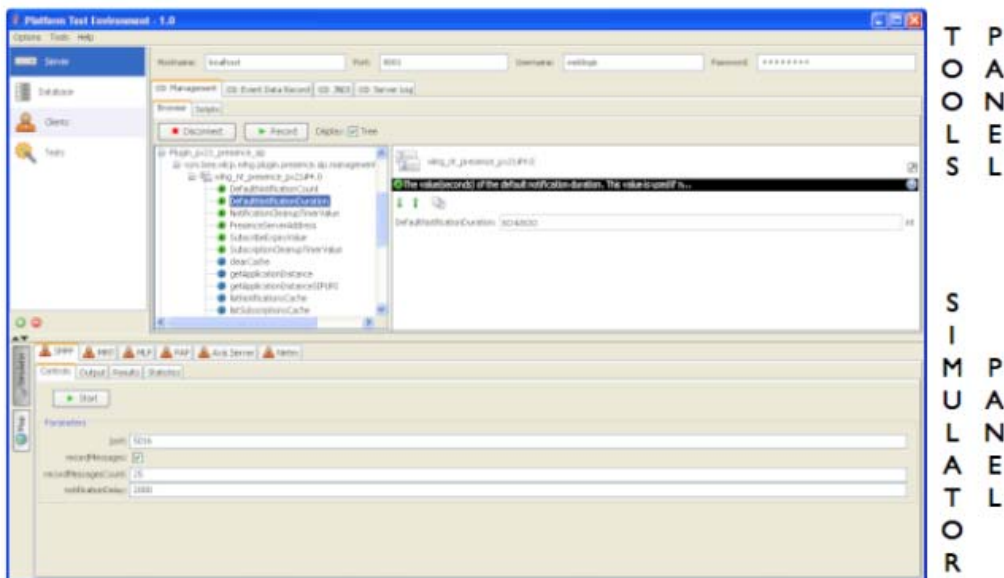
WARNING: Compatibility between the settings of this version of the PTE and any future versions is *not* guaranteed.

Navigating the Platform Test Environment GUI

The Java Swing-based GUI provides an easy to use access point to the many parts of the Platform Test Environment. Any changes made to the GUI are saved on exit. The GUI consists, broadly, of:

- [The Tools Panel](#), the upper panel
- [The Simulator Panel](#), the lower panel

Figure 10-2 The Main PTE GUI



The Tools Panel

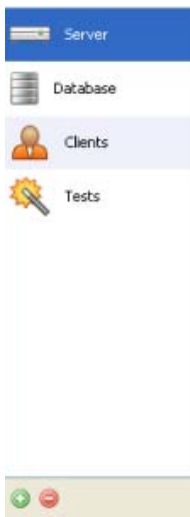
The Tools Panel is divided into two main sections:

- [The Tool Selector Panel](#)
- [The Tool Action Panel](#)

The Tool Selector Panel

Use this panel to select the tool you wish to use:

Figure 10-3 The Tool Selector



Notice the Plus and Minus icons in the lower left. You can use these to create multiple versions of the tools. For example, you could create a server tool to correspond to each server instance you are running.

The Tool Action Panel

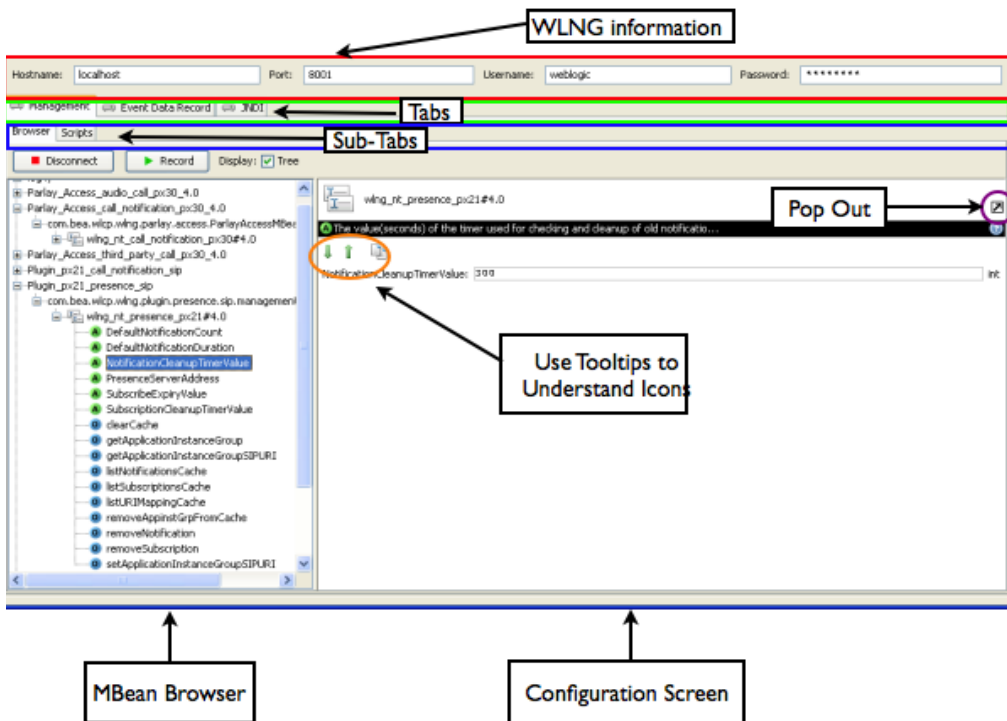
The Tool Action Panel displays the actions you can do in that particular tool.

- [The Server Tool](#)

- The Database Tool
- The Clients Tool
- The Tests Tool

The Server Tool

Figure 10-4 The Server Tool Action Panel with the Management Tab selected



At the top of the panel, you specify the server with which you wish to interact and your administrative username and password.

Below that are the three main tabs:

- **Management:** This tab lets you perform management tasks on Network Gatekeeper.

- **Browser:** The sub-tab lets you browse the MBeans on the server you have chosen. You can use this to make changes in the configuration of Network Gatekeeper instead of opening up the Management Console. Use the **Connect** button to connect to the server. Use your mouse to traverse the MBean list in the left column. Fill in your desired information in the right column. You can use the Pop Out icon on the right to have the same information displayed in a separate pop up window.

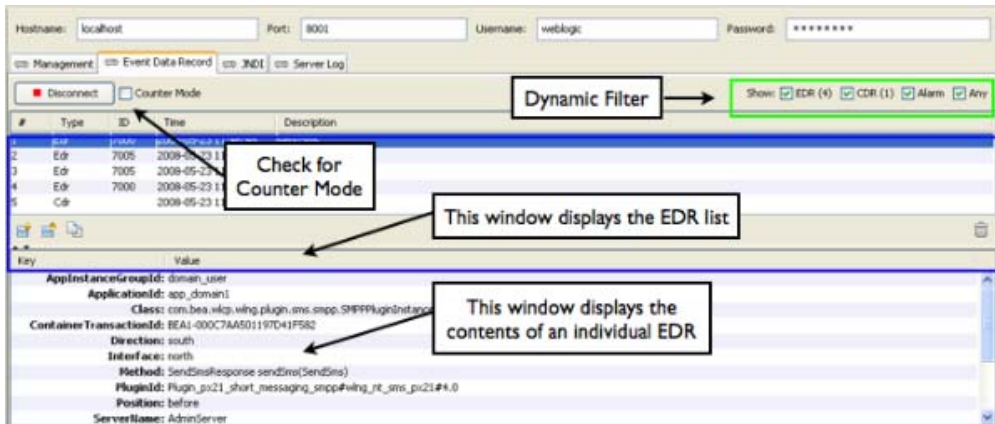
Note: For a detailed description of Network Gatekeeper management tasks, see the [System Administration Guide](#) and [Managing Accounts and SLAs](#), separate documents in this set.

If you want to record a script so that you can automate the management tasks you need to do more than once, use the **Record** button. Once you have begun recording, simply do the tasks you wish to automate. When you are finished, click the **Recording** button to stop. You will get a prompt asking you to name the script. The script will automatically be saved when you close the PTE.

Note: Many windows display convenience functions in the form of icons, as is shown above in the orange circle. Hover your mouse over the icon, and a tooltip explaining its function will appear.

- **Scripts:** This sub-tab displays a list of all the scripts you have recorded in the past. To play a script back, select the name you gave it when you created it, and click the **Run** button.
- **Event Data Record:** Use this tab to monitor EDRs, CDRs, and Alarms

Figure 10-5 The Event Data Record Tab



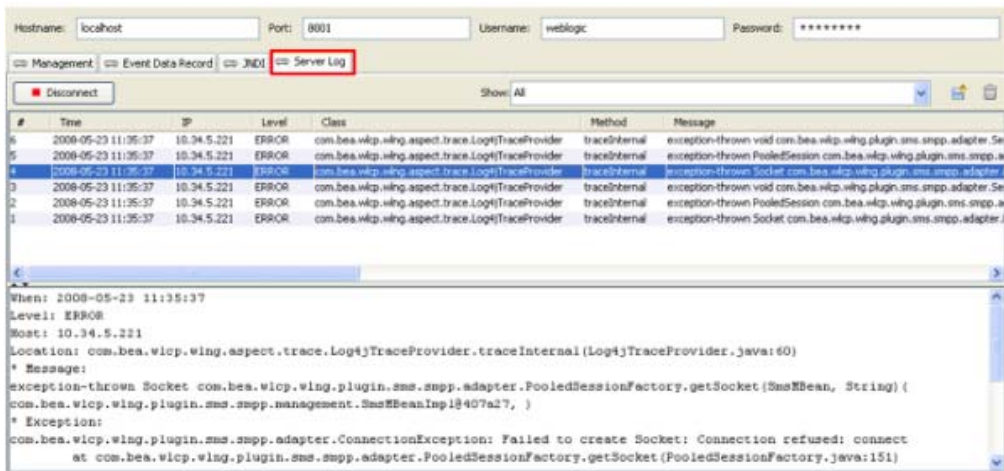
Click **Connect** to establish the JMS connection between the PTE EDR (and CDRs and Alarms) listener and the server instance. Select the sort of records you are interested in receiving using the Dynamic Filter, shown above in the green box. The list of EDRs displays in the window outlined in blue. Selecting a particular EDR in the list causes the contents of that record to be shown in the box outlined in red.

To export, import, or copy to the clipboard a list of EDRs, use the convenience icons in the bottom left of the EDR list window.

To save memory in situations where you are expecting a large number of EDRs, check Counter Mode. This will count the number of records, but will not display the contents.

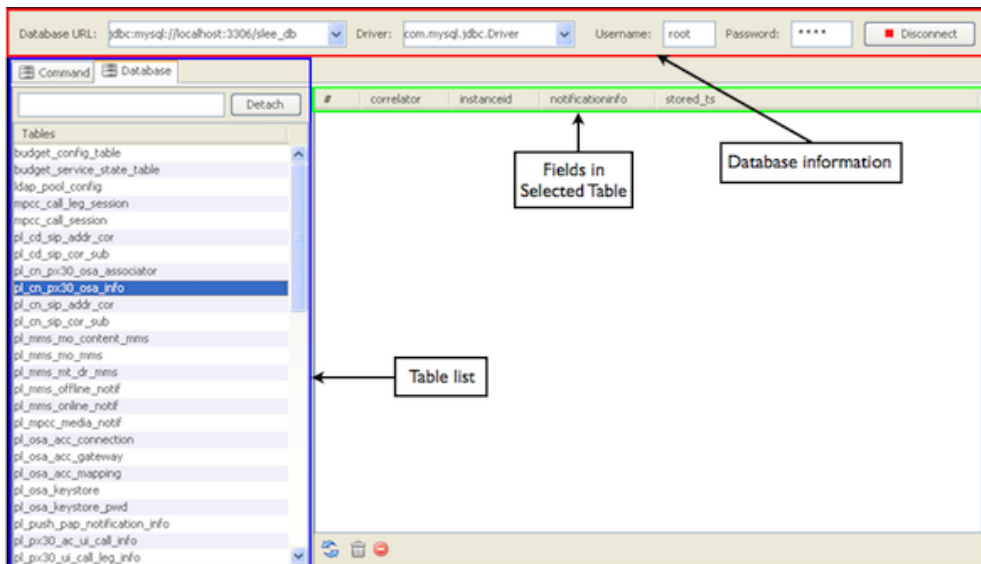
- **JNDI:** Use this tab to browse the JNDI tree.
- **Server Log:** Use this tab to browse the logs.

Figure 10-6 Server Log Tab



The Database Tool

Figure 10-7 The Database Tool Action Panel with the Database Tab selected



The Database tool lets you scan your database tables and manipulate them directly.

At the top of the panel you enter your database information, including your database username and password. Click the **Connect** button to connect to the database.

Below that are the two main tabs:

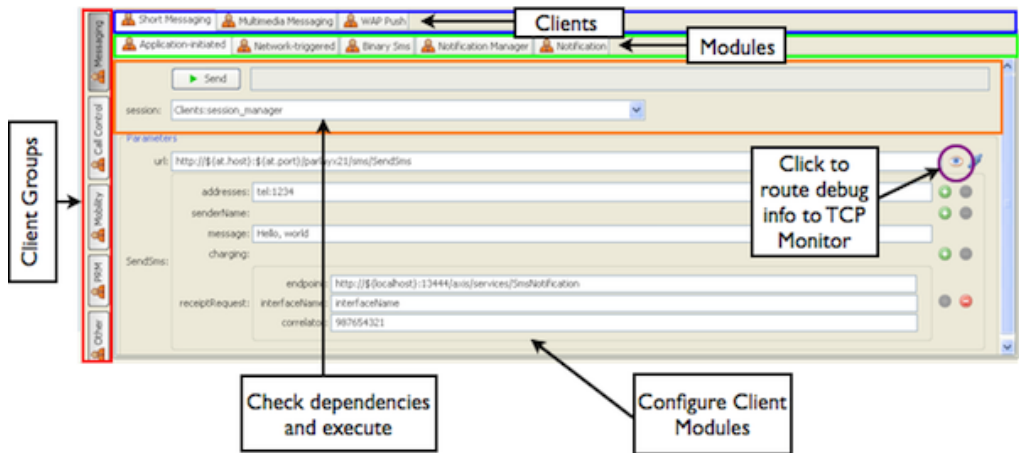
- **Database:** This tab allows you to scan your database tables, and to see the data in them. If there is a table you wish to monitor more closely, enter the name in the dialog box at the top of the left column and click **Detach**. That table then appears individually as an additional tab.
- **Command:** This tab allows you to enter any SQL command you wish directly to the database.

Figure 10-8 The Command Tab



The Clients Tool

Figure 10-9 The Clients Tool Action Panel with the Short Messaging Client Tab selected



The Clients Tool Action Panel is the most complex of the UIs for the Platform Test Environment. The display is divided into three hierarchical groups:

- **Client Groups:** Use the Clients Group column, outlined in red on the left above, to select the functional group of clients you are interested in manipulating. Your choices are:
 - Messaging
 - Call Control
 - Mobility
 - PRM (clients for the Partnership Relationship Management interfaces)
 - Other (clients for Session Management, Subscriber Profile, and the example communication service included with the PTE)
- **Clients:** Use the Clients tabs, outlined in blue above, to select the set of clients you are interested in manipulating. In the context of the PTE, a *client* is made up of *modules*. Each module represents one operation belonging to a set of interfaces. So, in the example above, the client that is selected represents the functionality offered by the Parlay X 2.1 Short Messaging set of interfaces.

Note: This particular client also happens to include a module that belongs to the Extended Web Services Binary SMS interface. This is because the two sets of interfaces share

a common network node and are bundled together in the same .EAR file in Network Gatekeeper.

- **Modules:** Use the Modules tabs, outlined in green above, to select the module you are interested in manipulating. A module represents a client that executes a single operation from an interface. So, in the example above, the module shown as **Application-initiated**, represents the single operation, `SendSms`. If you wanted to test the operation `SendSmsLogo`, you would need to create an additional module.

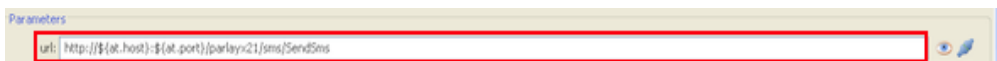
Once you have selected the module you wish to use, the display shows two windows. In the lower window, you configure the client module, setting any required parameters. In the upper right corner you will notice a small eye icon. Clicking this will route debug information to the TCP Monitor.

In the upper window, outlined in orange, dependencies are shown. In this case, the operation requires that the client acquire a session ID before sending the command. If you attempt to execute a command and the dependencies have not been set up, the PTE will offer to open the requisite modules for you. See [Figure 10-14](#) below for more information on running sessions.

Once the dependencies are taken care of, you can simply click the **Send** button if you wish to execute the operation at that moment. You can also choose to string several tests together into an automated set. See the [The Tests Tool](#) section below for more information.

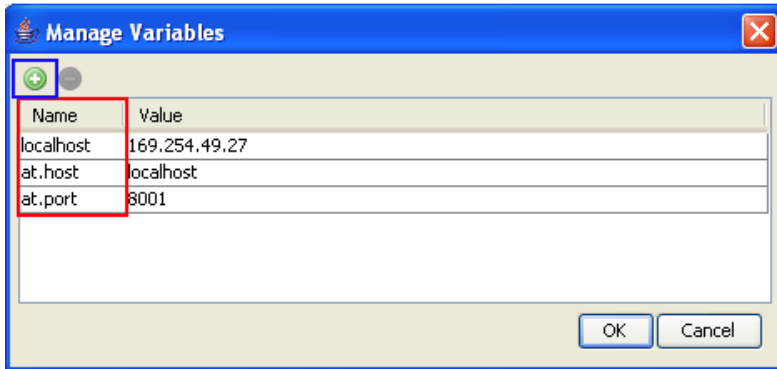
Note: In some module configuration windows, you will see URLs written out with variable values.

Figure 10-10 Variables in URLs



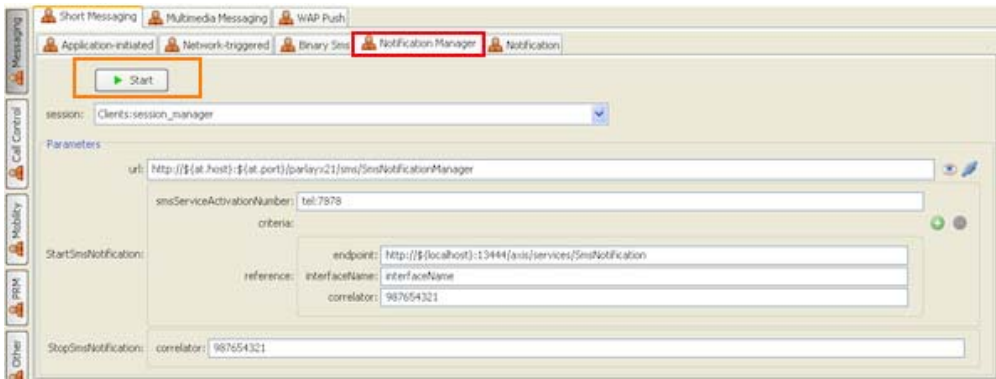
To set those variables, or to add others that are of use to you as you create and run tests, use the **Manage Variables** window. Click **Tools -> Manage Variables**. The default variables are shown outlined in red. To add a variable, click the plus button, outlined in blue. When you are finished editing, click the **OK** button.

Figure 10-11 The Manage Variables Window



Most client modules are stateless. When you click **Send**, the operation is executed and it completes. But some modules have state. They are started, and they run until they are stopped. The **Short Messaging Notification Manager** is such a module. As shown in Figure 10-12, it has a **Start** button, rather than a **Send** button.

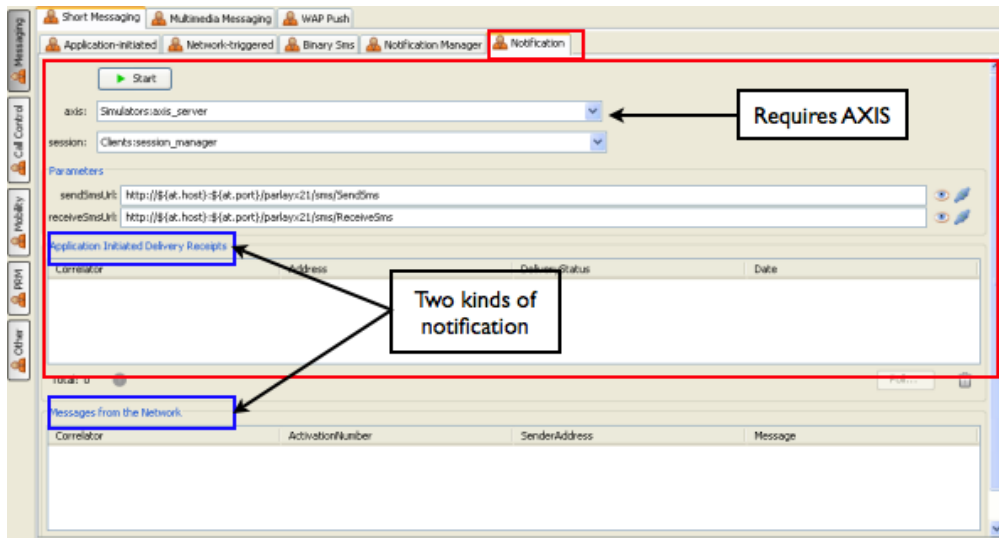
Figure 10-12 A Stateful Module



When you set up notifications, that is, when you tell Network Gatekeeper that your client is interested in receiving asynchronous messages from the network, the client must provide a web service to which the notifications can be delivered. These client-based web services also show up in the PTE GUI as client modules. In Figure 10-13 below, the **Notifications** tab is selected. This module runs the web service to which both kinds of Short Messaging notifications can be returned: Delivery Receipts for Application-initiated messages and actual SMSes sent from the

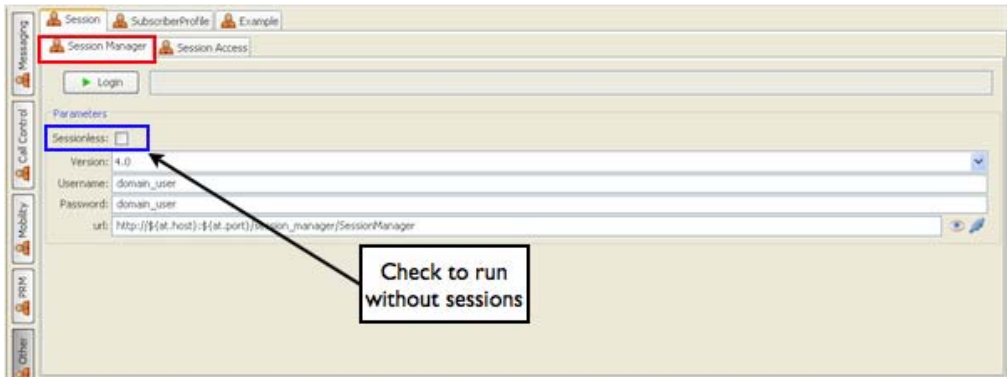
network to the client application. Notice the dependency on the Axis web server. It must be running for the web service to function.

Figure 10-13 The Client Notification Web Service



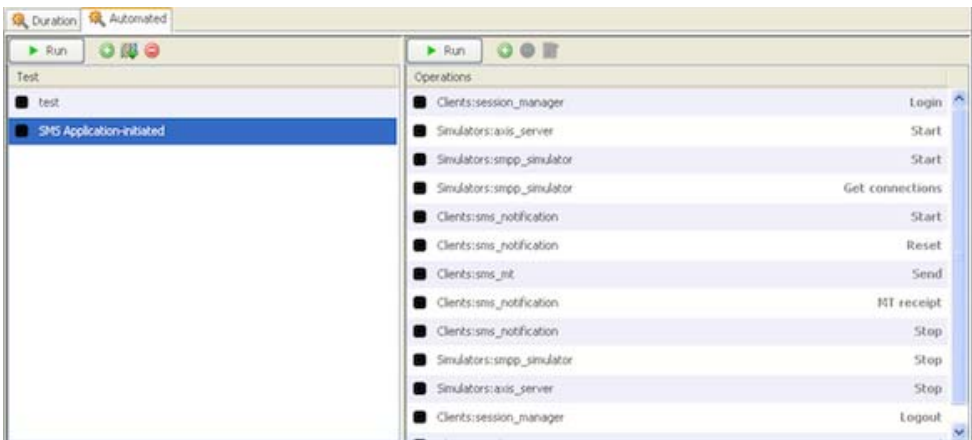
Finally there is the question of session management. The default setting in Network Gatekeeper is to require applications to start a session and get a Session ID before they send traffic through the system. But this requirement is configurable in Network Gatekeeper, and so the PTE makes it possible to turn the session requirement on and off. By selecting the Other client group and the Session client, the Session Manager module, you can simply check the Sessionless option, shown below in [Figure 10-14](#), and your clients will not be required to acquire or use a Session ID in order to run traffic.

Figure 10-14 Turning Sessions Off






The Tests Tool

Figure 10-15 Creating an Automated Test Sequence



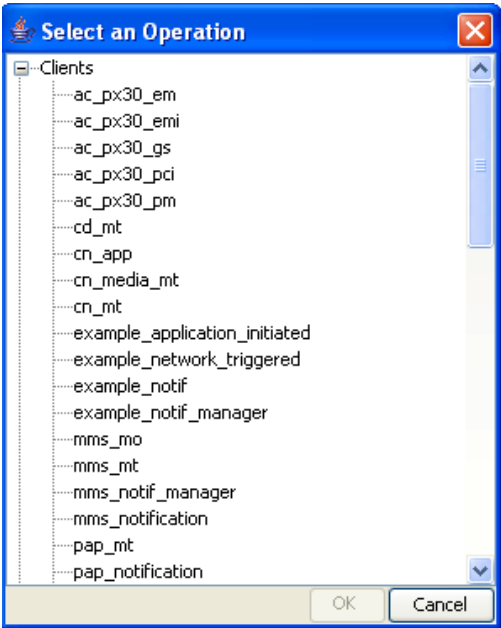
- Automated:** While it is possible to execute a single operation from the client configuration screen, seen above in [Figure 10-9](#), it is usually the case that you will want to string together a set of actions, including setting up client and simulator modules, and executing operations, into a single automated test sequence. The PTE makes this simple to do. To create a test, use the icons at the top of the left column, as shown in [Table 10-1](#):

Table 10-1 Creating tests

Icon	Function
	Adds a new test. When you click it, you will be prompted to give your test a name.
	Allows you to select from a set of predefined tests. This can be useful for understanding test flows
	Deletes the selected test.

Once you have created your test, click the plus icon at the top of the right column to add operations. The **Select an Operation** window opens, as in [Figure 10-16](#) below.

Figure 10-16 The Select an Operation Window

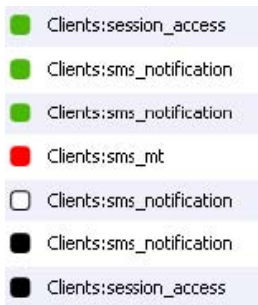


This window shows all operations available in every module in the PTE: clients, simulators, and even duration tests. Select your first operation, and then continue adding until you have completed the desired test sequence.

Note: If you have multiple clients that might be able to perform a particular operation, a popup window will appear and allow you to choose the one you wish to use.

Each test sequence that you create is automatically persisted when the PTE is shut down, so that you only need to create a test once. To run a single test from the GUI, click the **Run** button on the top of the right column. To run the entire test sequence from the GUI, click the **Run** button on the top of the left column. The status of the tests is indicated by the color of the box next to the individual test item names: see [Figure 10-17](#) below.

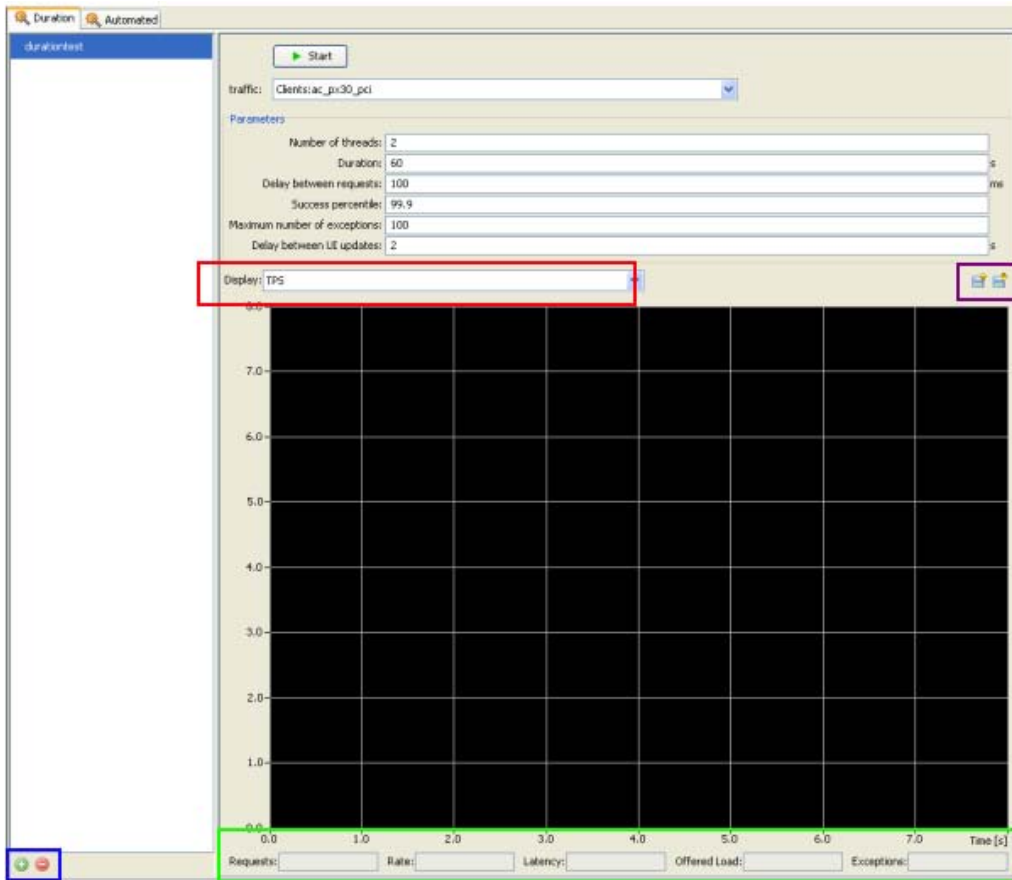
Figure 10-17 A Running Test



Green boxes indicate success; red boxes indicate failure; white boxes indicate an in-progress test; and black boxes indicate tests that have not yet run.

- **Duration:** In addition to functional testing, it is also important to see behavior over time. The PTE also makes it easy to create duration tests and includes a real-time graphing display. See [Figure 10-18](#) below.

Figure 10-18 Duration Tests



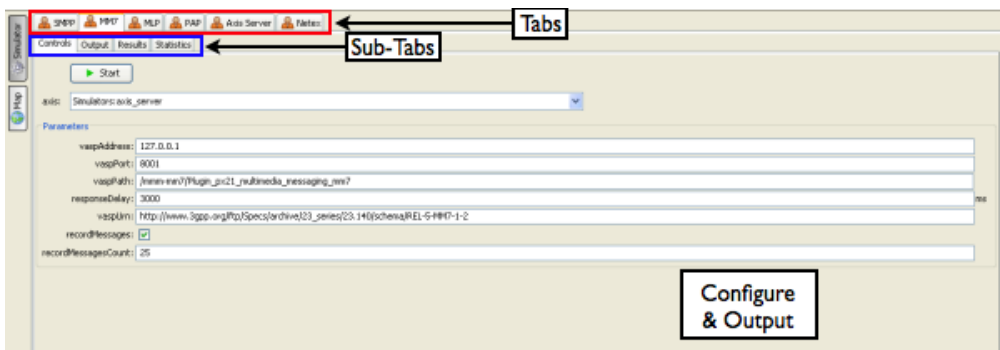
Create a new test by clicking the plus icon in the lower left corner. You are prompted for a name for the test. Configure the test in the upper portion of the right column. Select the type of traffic you wish to run, based on the client type, from the dropdown **Traffic** menu. Select what you wish to see graphed (Transactions Per Second, Exceptions, or Latency) in the **Display** dropdown menu outlined in red in the graphic above. Current statistics appear in the boxes at the bottom of the graph, outlined in green.

Make sure the appropriate simulator is running and start the test by clicking the **Start** button. The test runs in the background, so it is possible to run multiple tests in parallel.

Because duration test results are *not* saved across PTE sessions, you can choose to export results to be saved in a file and then import them back into the tool later, using the icons outlined in purple above.

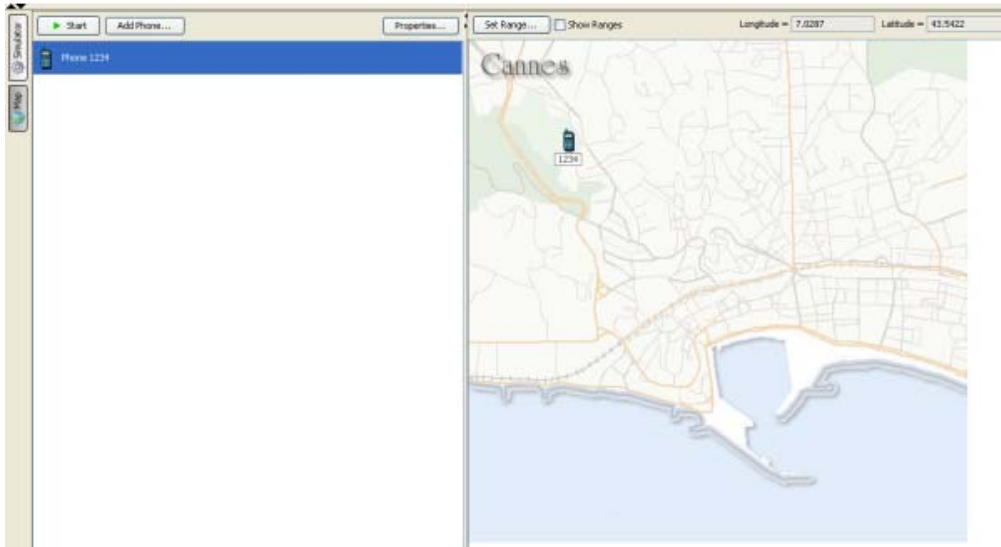
The Simulator Panel

Figure 10-19 The Simulator Panel with MM7 Selected



Like the Clients Tool, the Simulator panel is set up as a hierarchy. On the extreme left there are two buttons: **Map** and **Simulator**. Under the **Simulator** button is a set of tabs and sub-tabs. The tabs list the available simulator modules, including a simulator for the example communication service (Netex), and a separate tab and module for the Axis Server, which is required to run traffic over HTTP based protocols like MM7. See [Figure 10-19](#) above. Under the row of tabs is the row of sub-tabs. The number of sub-tabs depends on the module selected. In all cases, there is a **Control** tab in which you can set up any necessary configurations. This area is also where the **Start** button is for each of the modules. The other tabs may allow you to see the actual content of a message or show you the statistics associated with traffic.

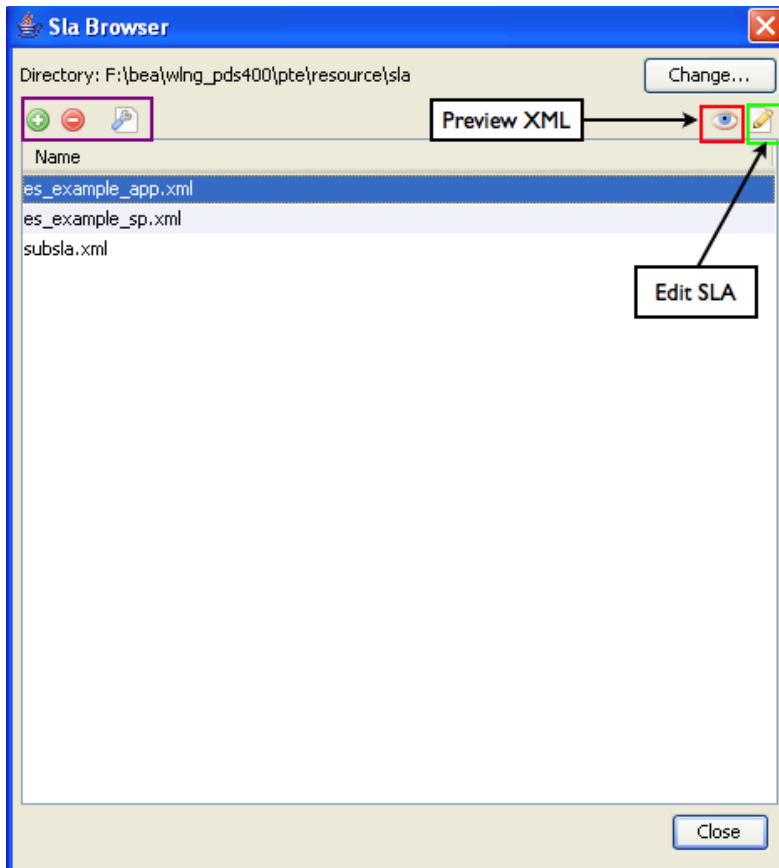
Figure 10-20 The Maps Panel



The Maps panel is a variant of a tool which was originally developed as part of the Application Developers SDK. It provides a map on which you can place phone terminals. This offers visual support for testing Parlay X 2.1 SMS, MMS, and Terminal Location traffic.

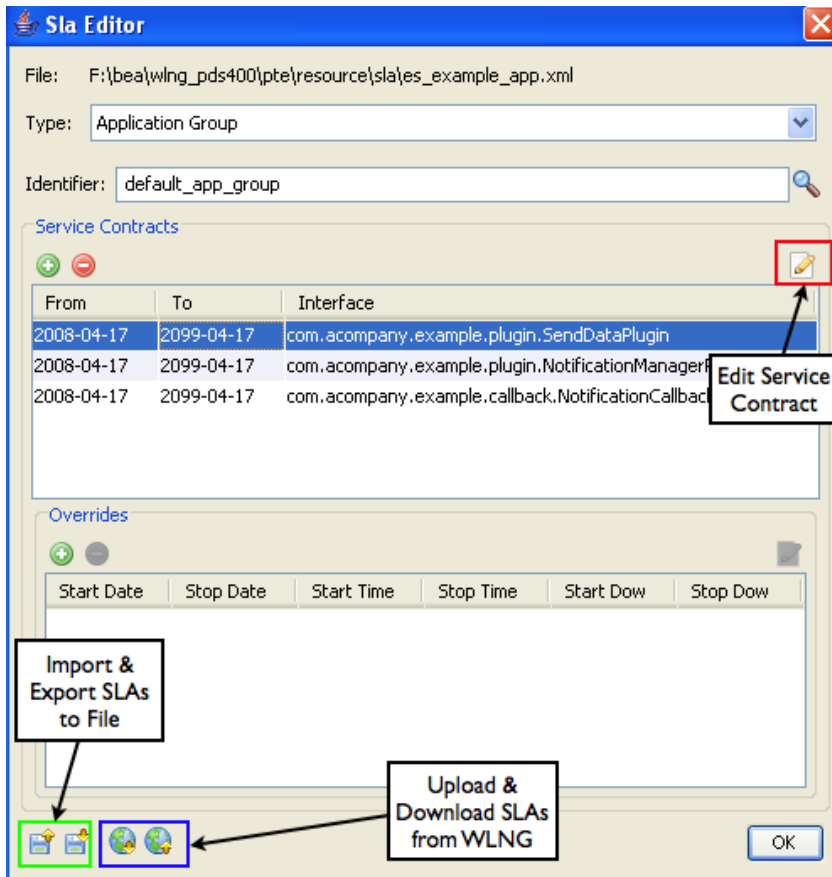
The SLA Editor

Managing large XML files can be difficult, particularly in a test environment where you may wish to change small details multiple times for various iterations of testing. To help you manage your SLAs, the PTE ships with an SLA editor, which manages the tags and validation so that you can focus on setting appropriate values. To access the SLA editor, first make sure you have selected the Server Tool and are connected to the server. Then click **Tools -> Manage SLAs** in the Menu Bar. The SLAs are fetched from the file system and the SLA Browser window opens. See [Figure 10-21](#).

Figure 10-21 The SLA Browser


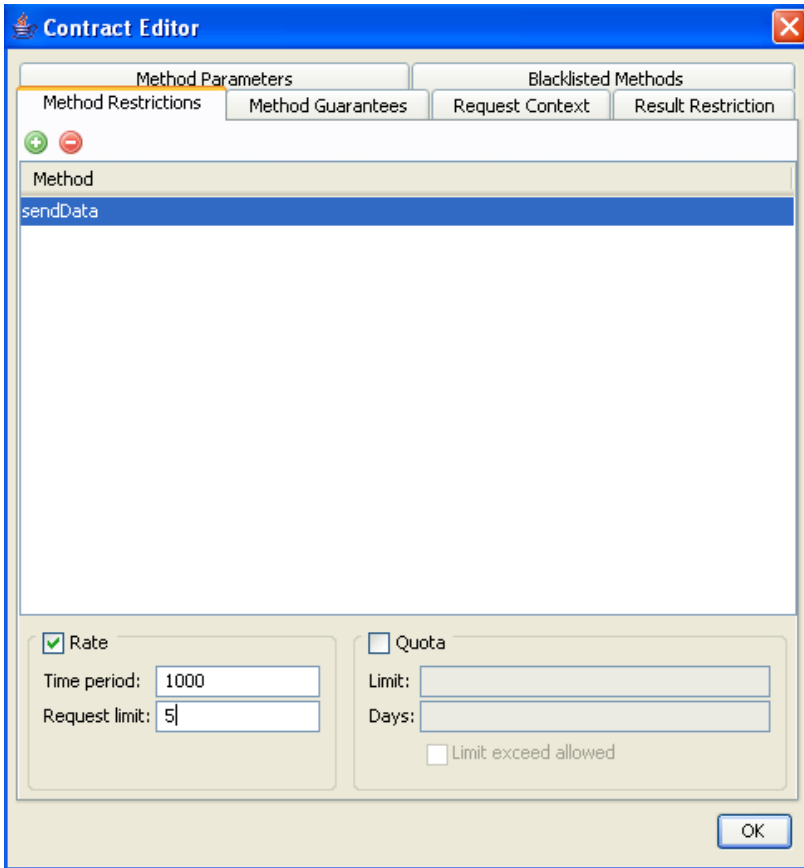
The SLA files that were fetched from the file system are listed in the main window. If you want to search for other files that may be on your system, click the **Change...** button and a file browser appears. To create an entirely new SLA, click the Plus button in the upper left corner, outlined in purple. You can also delete or rename SLAs using buttons in the same area.

To edit an SLA, select the one you are interested in and then click the pencil icon in the upper right corner, outlined in green. The **SLA Editor** window opens. See [Figure 10-22](#).

Figure 10-22 The SLA Editor

Select the SLA type using the dropdown menu and specify the group identifier. If you wish to import a different SLA from the file system, or to save changes out, use the import and export icons on the bottom left. To upload the SLAs to, or download them from, the repository in the running instance of WLNG, use the icons outlined in blue on the bottom left.

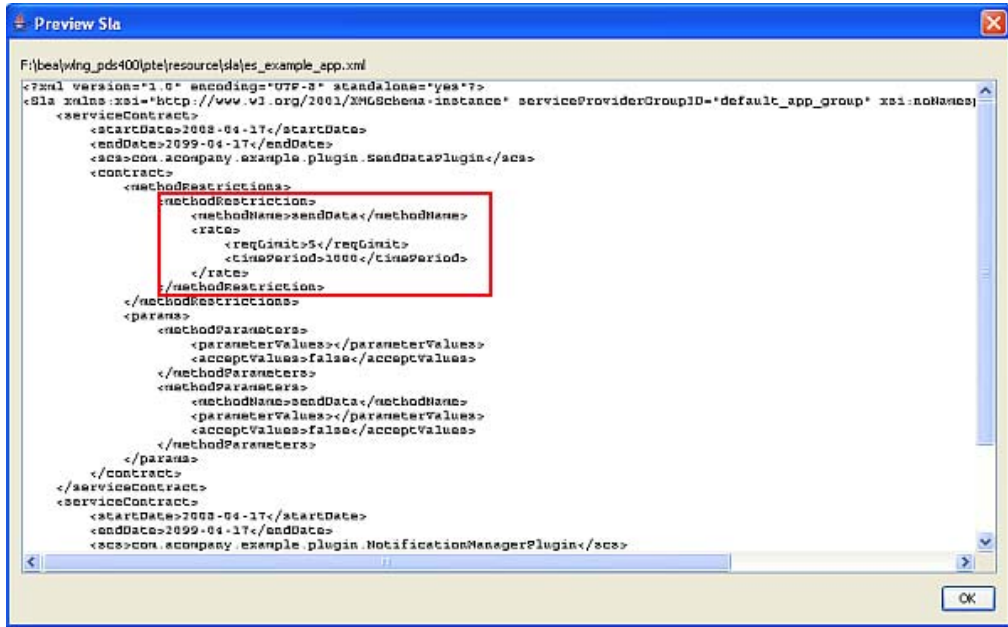
There are two basic kinds of editing you can do - the main **Service Contracts** and any **Overrides** you have specified. Each has its own window. To edit a Service Contract, select the item you are interested in and click the pencil icon in the upper right corner. This opens the **Contract Editor**. See [Figure 10-23](#).

Figure 10-23 The Contract Editor

The tags that can be edited appear as tabs at the top of the window. For more information on these tags, see the [“Defining Service Provider Group and Application Group SLAs”](#) chapter in *Managing Accounts and SLAs* a separate document in this set.

In the figure, a rate limit method restriction is being added to the `sendData` operation of the sample communication service. When you have made your edits, click the **OK** button and the window closes. Click **OK** once again (on [Figure 10-22](#)) the window closes. To preview the edits you have made in XML format, click the Eye icon at the top left of the SLA Browser ([Figure 10-21](#)). The **Preview SLA** window opens. See [Figure 10-24](#).

Figure 10-24 The Preview SLA window



The Method Restriction rate limit that was added in [Figure 10-23](#) is shown outlined in red.

When you have completed your edits, simply click the **Close** button on the **SLA Browser** window.

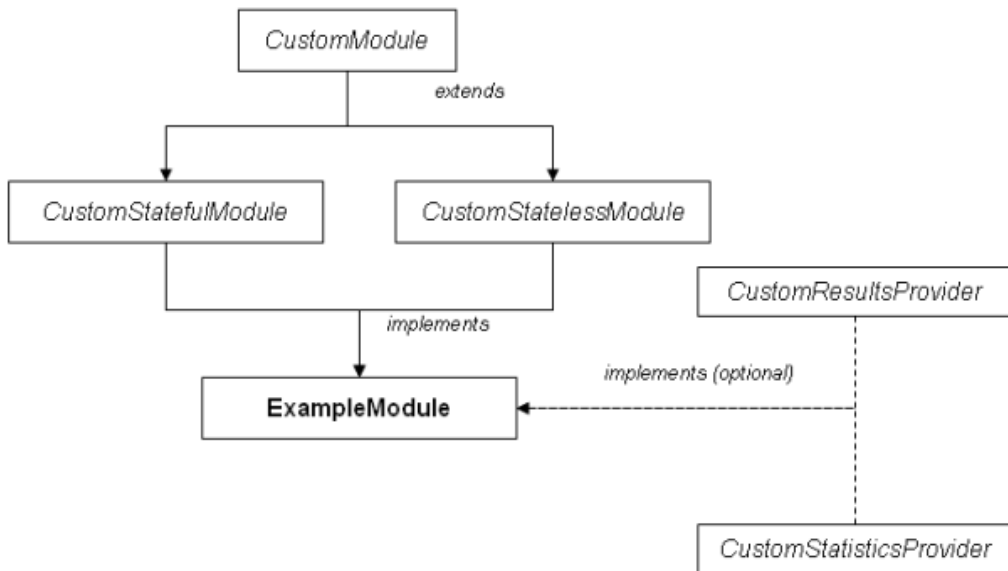
Extending the Platform Test Environment

One of the most common uses for the Platform Testing Environment is to test extension communication services. Depending on how those extensions are implemented, you may need to create one or more new modules so that the PTE can interact successfully with your new communication service. You can implement new client modules, and even new clients containing multiple modules, if support for the application-facing interface that you want your communication service to use is not already available in the PTE. You can also implement new simulators, if the network node type that you want your communication service to interact with is not available. From the point of view of the PTE, a module is a module.

The only relevant distinction in the PTE is between modules for operations that simply execute and return and those for operations that start a process which runs until it is turned off. These are called, respectively, stateless and stateful modules. See [Figure 10-12](#) for more information.

Stateless modules must implement the `CustomStatelessModule` SPI and stateful modules must implement the `CustomStatefulModule` SPI. There are two additional, optional interfaces that can be implemented if you would like your module to display results (for example, a notification, a message from the network delivered to a client Web Service) or provide statistics in the GUI. The custom module SPI hierarchy is as follows:

Figure 10-25 The Custom SPI Hierarchy

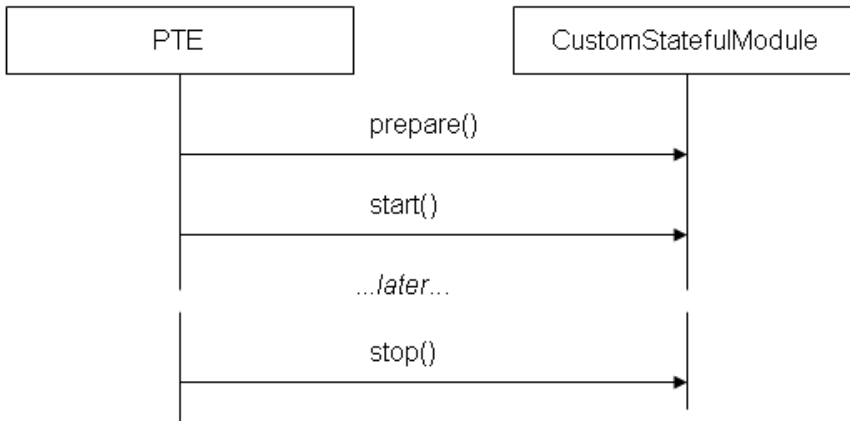


Any module that is created must be packaged as a .jar file which must be located in the `PTE_HOME/lib/modules/` directory. The root of the .jar file must include a descriptor file called `module.xml`. All custom modules automatically load when the PTE starts up.

Note: The modules created for use with the example communication service are located in `<bea_home>/wlng_pds400/example/pte_module`.

The Stateful SPI

Figure 10-26 shows the execution sequence for a stateful module:

Figure 10-26 The Execution Sequence for a Stateful Module

The following listing is the SPI that must be implemented by stateful PTE modules.

Listing 10-1 CustomStatefulModule SPI

```

package com.bea.wlcp.wlng.et.spi;

/**
 * This interface must be implemented by a custom stateful module.
 * A stateful module has a start() and a stop() method and will be
 * represented in the UI by the Start/Stop button.
 * Note: a stateful module is not used in duration tests.
 *
 * @author Copyright (c) 2008 by BEA Systems, Inc. All Rights Reserved.
 */
public interface CustomStatefullModule extends CustomModule {

    /**
     * Starts the module.
     *
     * @param context The custom module context
     */
  
```

```

    * @return true if the module successfully started
    * @throws Exception Any exception preventing the module to start
    */
    public boolean start(CustomModuleContext context) throws Exception;
}

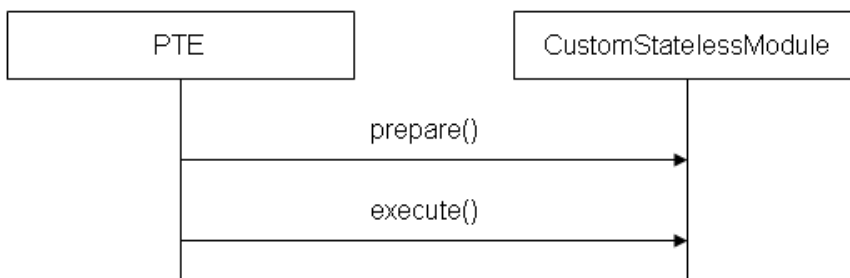
    * Stops the module.
    * @param context The custom module context
    * @return true if the module successfully stopped
    * @throws Exception Any exception preventing the module to stop
    */
    public boolean stop(CustomModuleContext context) throws Exception;
}

```

The Stateless SPI

Figure 10-27 shows the execution sequence for a stateless module:

Figure 10-27 The Execution Sequence for a Stateless Module



The following listing is the SPI that must be implemented by stateless PTE modules.

Listing 10-2 CustomStatelessModule SPI

```
package com.bea.wlcp.wlng.et.spi;

/**
 * This interface must be implemented by custom stateless module.
 * A stateless module has only an execute() method and will be
 * represented in the UI by the Send button.
 *
 * @author Copyright (c) 2008 by BEA Systems, Inc. All Rights Reserved.
 */

public interface CustomStatelessModule extends CustomModule {

    /**
     * Asks the module to execute its job and return the result.
     *
     * @param context The custom module context
     * @return The result of the execution
     * @throws Exception Any exception that occurred during the execution
     */
    public Object execute(CustomModuleContext context) throws Exception;
}
```

The Custom Base SPI

This following is the base SPI for custom PTE modules. It should *not* be implemented directly. See the first comment.

Listing 10-3 The Custom Base SPI

```
package com.bea.wlcp.wlng.et.spi;
```

```

import com.bea.wlcp.wlng.et.api.CustomModuleContext;

/**
 * This interface defines the general API a custom module must implement.
 * Note: a custom module should NOT implement this interface directly but
 * one of the subinterface like CustomStatefulModule or
CustomStatelessModule.
 *
 * @author Copyright (c) 2008 by BEA Systems, Inc. All Rights Reserved.
 */
public interface CustomModule {

    /**
     * Prepares the module with the given context. This method is invoked
before
     * the module is executed: it can be used by the module to prepare
     * any internal states needed.
     * Note: when a duration test is performed on the Platform Test
Environment,
     * prepare() is invoked only once at the beginning of the duration test.
     *
     * @param context The context of the custom module
     * @throws Exception Any exception that occurred during the module
preparation
     */
    public void prepare(CustomModuleContext context) throws Exception;

```

The Custom Results Provider SPI

The following listing is the SPI that must be implemented by modules that wish to display some sort of results in the GUI.

Listing 10-4 The CustomResultsProvider SPI

```
package com.bea.wlcp.wlng.et.spi;

/**
 * A custom module can implement this interface if it wants to provide
 * a list of results in the UI. The PTE will automatically display a list
 * and handle the user interaction with it.
 *
 * @author Copyright (c) 2008 by BEA Systems, Inc. All Rights Reserved.
 */

public interface CustomResultsProvider {

    /**
     * Clears the results.
     */

    public void clearResults();

    /**
     * Returns an array of string that will be used to create
     * the name of each column of the results table.
     * @return An array of string to create the column headers
     */

    public String[] getResultsHeaders();

    /**
     * Returns the results. Each result is composed of a map whose keys are
     * the same as the strings returned by getResultsHeaders().
     */
}
```



```

*

* Note: It is up to the custom module to accumulate the results until
* this method is invoked by the PTE.
*

* @return A list of results
*/

public List<Map<String,String>> getResults();
}

```

The Custom Statistics Provider SPI

The following listing is the SPI that must be implemented by modules that wish to display statistics in the GUI.

Listing 10-5 The CustomStatisticsProvider SPI

```

package com.bea.wlcp.wlng.et.spi;

/**
 * A custom module can implement this interface if it wants to provide
 * some statistics in the UI. The PTE will automatically display a list
 * and handle the user interaction with it.
 *
 * @author Copyright (c) 2008 by BEA Systems, Inc. All Rights Reserved.
 */

public interface CustomStatisticsProvider {

    /**
     * Clears the statistics.
     */

    public void clearStatistics();
}

```

```
/**
 * Returns a map of statistics. Each key represent a particular statistic
 * and the value the value of the statistic.
 * @return The map of statistics
 */
public Map<String,String> getStatistics();
}
```

The Context API

The following listing is the API that allows modules to acquire context.

Listing 10-6 The Context API

```
package com.bea.wlcp.wlng.et.api;

/**
 * This interface defines the context available to a custom module.
 *
 * @author Copyright (c) 2008 by BEA Systems, Inc. All Rights Reserved.
 */
public interface CustomModuleContext {

    /**
     * Returns the custom module data object as described in the module.xml
     * @return The custom module data object
     */
    public Object getData();
}
```

```

/**
 * Returns the module of the specified type that this module depends on.
 * If there are many modules of the same type, the one chosen by the user
 * in the UI will be chosen.
 *
 * @param type The type of module
 * @return The module instance of the specified type
 */
public CustomModule getDependency(String type);

/**
 * Prepares the stub that the module will use to send a request. The PTE
 * will perform various changes to the stub depending on the UI settings,
 * like TCP Monitor or Override Endpoint.
 *
 * @param stub The stub to prepare
 * @param path The path to the parameter declared in module.xml that
corresponds
 * to the stub url. Use null if it doesn't have any corresponding parameter.
 */
public void prepareStub(Stub stub, String path);

/**
 * Deploy (or undeploy) a service using a specific WSDD file.
 *
 * @param wsddFile The WSDD file that the axis server will execute
 * @throws Exception Any exception when executing the command

```

```
*/  
  
public void axisDeploy(String wsddFile) throws Exception;
```

The Module.xml Descriptor File

Every module is packaged in a .jar file with a descriptor file, `module.xml`, in its root. What is in the file depends on the nature of the module.

The following is the listing for a client module and the simulator module supplied with the example communication service:

Listing 10-7 The example module.xml

```
<module-factory xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
  
xsi:noNamespaceSchemaLocation="http://www.bea.com/ns/wlng/40/et">  
  
  <module name="example_application_initiated"  
    type="client"  
    class="com.bea.wlcp.wlng.et.example.SendDataModule"  
    version="1.0"  
    depends="session"  
    uiPanel="client"  
    uiTabs="Other,Example,Application-Initiated"  
  >  
  
    <data>  
      <parameter name="Parameters"  
        class="com.bea.wlcp.wlng.et.example.SendDataData"  
        occurs="1">
```

```

    <parameter name="url"
        class="java.lang.String"
        occurs="1"
        default="http://${at.host}:${at.port}/example/SendData"
        monitor="true"/>

    <parameter name="data"

class="com.acompany.schema.example.data.send.local.SendData"
        occurs="1">

        <parameter name="address"
            class="java.net.URI"
            occurs="1"
            default="tel:1234"/>

        <parameter name="data"
            class="java.lang.String"
            occurs="1"
            default="Hello, world"/>
    </parameter>
</parameter>

</data>
</module>

...

```

Using the Platform Test Environment

```
<module name="example_simulator"
    type="netex"
    class="com.bea.wlcp.wlng.et.example.SimulatorModule"
    version="1.0"
    uiPanel="simulator"
    uiTabs="Netex"
    >
<data>
    <parameter name="Parameters"
        class="com.bea.wlcp.wlng.et.example.SimulatorData"
        occurs="1">

        <parameter name="port"
            class="int"
            occurs="1"
            default="5001"/>

    </parameter>
</data>
</module>

</module-factory>
```

Below is the entire .xsd file for module.xml:

Listing 10-8 The module.xsd File

```
<?xml version="1.0"?>

<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
            elementFormDefault="qualified">

    <!-- Main element that describes one or more modules -->
    <xs:element name="module-factory">
        <xs:complexType>
            <xs:sequence>
                <xs:element ref="module" minOccurs="0" maxOccurs="unbounded"/>
            </xs:sequence>
        </xs:complexType>
    </xs:element>

    <!-- Defines a single module -->
    <xs:element name="module">
        <xs:complexType>
            <xs:sequence>
                <!-- Optional data of the module -->
                <xs:element ref="data" minOccurs="0" maxOccurs="1"/>
            </xs:sequence>
            <!-- Name of the module. It will be used also for the display -->
            <xs:attribute name="name" type="xs:string" use="required"/>
            <!-- Type of the module -->
            <xs:attribute name="type" type="xs:string" use="required"/>
            <!-- Class of the module (fully qualified) -->
            <xs:attribute name="class" type="xs:string" use="required"/>
        </xs:complexType>
    </xs:element>
</xs:schema>
```

```
<!-- Version of the module -->
<xs:attribute name="version" type="xs:string" use="required"/>
<!-- Name of the module this module depends on.
Predefined types are:
- session : session module
- axis : axis server module
The PTE will make sure that before this module is started, the
dependent module is running. -->
<xs:attribute name="depends" type="xs:string" use="optional"/>
<!-- UI panel where the module will be located (see ui-panels) -->
<xs:attribute name="uiPanel" type="ui-panels" use="required"/>
<!-- Location of the module in the panel tabs.
The location is a list of UI tab names separated by comma. For example:
"Other,Example,SendData"
means that the module will be in a tab named "SendData"
located in the tab "Example" located in tab "Other".
The name of each tab is available in the UI.
If a tab doesn't exist for a particular name, it will be created.-->
<xs:attribute name="uiTabs" type="xs:string" use="required"/>
</xs:complexType>
</xs:element>

<!-- Available UI panels -->
<xs:simpleType name="ui-panels">
  <xs:restriction base="xs:string">
    <!-- Client panel -->
    <xs:enumeration value="client"/>
  </xs:restriction>
</xs:simpleType>
```



```

    <!-- Simulator panel -->
    <xs:enumeration value="simulator"/>
  </xs:restriction>
</xs:simpleType>

<!-- Data of the module -->
<xs:element name="data">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="parameter" minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>

<!-- A single parameter -->
<xs:element name="parameter">
  <xs:complexType>
    <xs:sequence>
      <!-- Can contain other parameters too -->
      <xs:element ref="parameter" minOccurs="0" maxOccurs="unbounded"/>
      <!-- Values restriction of the parameter (see restricted) -->
      <xs:element ref="restricted" minOccurs="0" maxOccurs="1"/>
      <!-- Internal use only -->
      <xs:element ref="instance" minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
    <!-- Name of the parameter. It will be used to access the member
    of the parameter (MBean-style) -->

```

```
<xs:attribute name="name" type="xs:string" use="required"/>
<!-- Fully qualified name of the parameter -->
<xs:attribute name="class" type="xs:string" use="required"/>
<!-- Occurrences of the parameter (see parameter-occurs) -->
<xs:attribute name="occurs" type="parameter-occurs" use="required"/>
<!-- Default value of the parameter -->
<xs:attribute name="default" type="xs:string" use="optional"/>
<!-- Set to true if this parameter represent an URL to a stub. If true,
it can be monitored by TCP monitor and have other properties -->
<xs:attribute name="stub" type="xs:boolean" use="optional"/>
<!-- Set to true if this parameter must be instanciated at creation
time.
This is only useful if the parameter is optional. -->
<xs:attribute name="instanciate" type="xs:boolean" use="optional"/>
<!-- Internal use only -->
<xs:attribute name="preview" type="xs:boolean" use="optional"/>
<!-- Internal use only -->
<xs:attribute name="help" type="xs:boolean" use="optional"/>
<!-- Internal use only -->
<xs:attribute name="multiline" type="xs:integer" use="optional"/>
<!-- Internal use only -->
<xs:attribute name="timebase" type="parameter-timebase"
use="optional"/>
<!-- Optional display string to use instead of the name in the UI -->
<xs:attribute name="display" type="xs:string" use="optional"/>
</xs:complexType>
</xs:element>
```

```

<!-- The value the parameter is restricted to -->
<xs:element name="restricted">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="value" minOccurs="1" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>

```

```

<!-- A value has a content only -->
<xs:element name="value">
  <xs:complexType>
    <xs:attribute name="content" type="xs:string" use="required"/>
  </xs:complexType>
</xs:element>

```

```

<!-- The occurrences of a parameter -->
<xs:simpleType name="parameter-occurs">
  <xs:restriction base="xs:string">
    <!-- required (one and only one) -->
    <xs:enumeration value="1"/>
    <!-- optional -->
    <xs:enumeration value="?" />
    <!-- one or more -->
    <xs:enumeration value="+" />
    <!-- zero or more -->
    <xs:enumeration value="*" />
  </xs:restriction>
</xs:simpleType>

```

```
<!-- tree of parameter-->

<xs:enumeration value="t"/>

</xs:restriction>

</xs:simpleType>

<!-- Internal use only -->

<xs:element name="instance">

  <xs:complexType>

    <xs:attribute name="v1" type="xs:string" use="required"/>

    <xs:attribute name="v2" type="xs:string" use="optional"/>

    <xs:attribute name="v3" type="xs:string" use="optional"/>

  </xs:complexType>

</xs:element>

<!-- Internal use only -->

<xs:simpleType name="parameter-timebase">

  <xs:restriction base="xs:string">

    <xs:enumeration value="ms"/>

    <xs:enumeration value="s"/>

    <xs:enumeration value="min"/>

    <xs:enumeration value="h"/>

  </xs:restriction>

</xs:simpleType>

</xs:schema>
```

Using the Unit Test Framework (UTFW) with the Platform Test Environment

Unit tests are a core part of any testing cycle. Data are input into the system and the results are retrieved from the system and compared to expected values, all programmatically.

The Unit Test Framework allows you to create unit tests for the PTE easily. You implement your test class based on the abstract class `WlngBaseTestCase` and it manages the mechanics of using JMX and JMS to connect to the PTE for you. A `test.properties` file located in the same directory can be used to define commonly changed properties of the test.

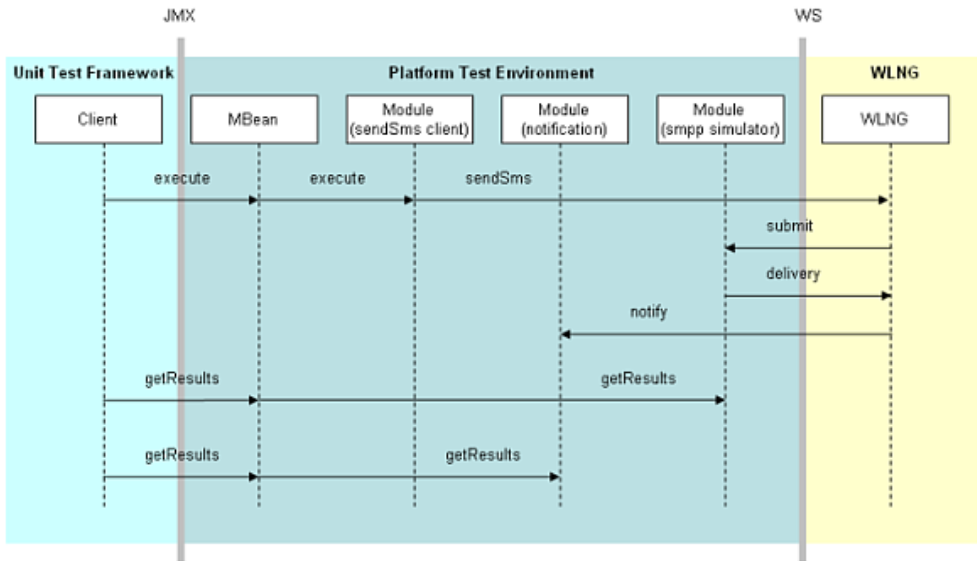
Note: The `WlngBaseTestCase` class is located in
`<bea_home>/wlng_pds400/lib/wlng/pte_api.jar`.

There are five basic steps to creating a unit test for the PTE:

1. Create any necessary client or simulator modules for the PTE using the required SPI and XML configuration files
2. Implement a test class based on the abstract class `WlngBaseTestCase`
3. Provision Network Gatekeeper
4. Start the Platform Test Environment and make sure the modules are correctly loaded
5. Run the test class

Note: The PTE should be running in Console (non-GUI) mode when you run your test. See [Installing and Running the Platform Test Environment](#) for more information on starting in Console mode.

Figure 10-28 An SMS Unit Test Sequence



The test sequence flow is as follows:

1. The test client calls `execute` on the PTE's Module Management MBean. The mechanics of the JMX call are taken care of by the base class.
2. The MBean calls `execute` on the specified Module, in this case `sendSMS`. This request includes a request for delivery receipts.
3. The `sendSMS` Module sends the request to Network Gatekeeper.
4. Network Gatekeeper processes it and submits it to the network simulator module, in this case the SMPP module
5. The simulator module returns a Delivery Receipt to Network Gatekeeper
6. Network Gatekeeper sends the receipt on to the Notification module (which represents the client Web Service implementation)
7. The test client retrieves the result of Network Gatekeeper's `submit` from the SMPP simulator
8. The test client retrieves the Delivery Receipt from the Notification module

To help you understand more clearly how all this works, there is an example unit test, which tests the example communication service, using the example clients and simulator. In standard installations, it is located in

<bea_home>/wlng_pds400/example/unit_test/src/com/bea/wlcp/wlng/pds/example
 . See [Listing 10-9](#) below.

Listing 10-9 A Unit Test for the Example Communication Service

```
package com.bea.wlcp.wlng.et.example;

import com.bea.wlcp.wlng.et.api.WlngBaseTestCase;

import java.util.List;
import java.util.Map;

/**
 * This class illustrates how to use the Unit Test Framework to
 * test the Communication Service Example. A few things are assumed before
 * running this class:
 * - the WLNG should be running and configured properly
 * - the CS example should be deployed and ready
 *
 * Note: this example uses also the wlngJmx to be able to access the WLNG
 * MBeans to ask the CS example plugin to connect to the Netex simulator.
 *
 */
public class TestSendData extends WlngBaseTestCase {
```

```
private static final String SEND_DATA_MBEAN =
"com.bea.wlcp.wlng.ptc:group=traffic,name=SendData";

private static final String NETWORK_TRIGGERED_MBEAN =
"com.bea.wlcp.wlng.ptc:group=traffic,name=NetworkTriggered";

private static final String NOTIF_MANAGER_MBEAN =
"com.bea.wlcp.wlng.ptc:group=client,name=NotificationManager";

private static final String NOTIF_MBEAN =
"com.bea.wlcp.wlng.ptc:group=client,name=Notification";

private static final String NETEX_SIMULATOR_MBEAN =
"com.bea.wlcp.wlng.ptc:group=netex,name=Simulator";

private static final String EXAMPLE_PLUGIN_MBEAN =
    "com.bea.wlcp.wlng:AppName=es_example_nt#4.0," +
    "InstanceName=example_netex_plugin," +

"Type=com.acompany.plugin.example.netex.management.ExampleMBean";

public TestSendData() throws Exception {
}

@Override
protected void setUp() throws Exception {
    super.setUp();
    wlngJmx.open("localhost", 8001, "weblogic", "weblogic");
    start(NETEX_SIMULATOR_MBEAN);
}

@Override
protected void tearDown() throws Exception {
```



```

        wlngJmx.close();

        stop(NETEX_SIMULATOR_MBEAN);

        super.tearDown();
    }

    public void testSendData() throws Exception {
        assertTrue(isRunning(NETEX_SIMULATOR_MBEAN));
        resetStatistics(NETEX_SIMULATOR_MBEAN);

        wlngJmx.invokeOperation(EXAMPLE_PLUGIN_MBEAN, "connect");

        String data = "Hello at " + System.currentTimeMillis();
        String to = "tel:1234";

        putParameter(SEND_DATA_MBEAN, "url",
            "http://localhost:8001/example/SendData");
        putParameter(SEND_DATA_MBEAN, "data.data", data);
        putParameter(SEND_DATA_MBEAN, "data.address", to);

        start(SESSION_MBEAN);
        assertTrue(isRunning(SESSION_MBEAN));

        execute(SEND_DATA_MBEAN);

        Thread.sleep(2000);

        stop(SESSION_MBEAN);
    }

```

Using the Platform Test Environment

```
Map<String,String> stats = listAllStatistics(NETEX_SIMULATOR_MBEAN);
System.out.println("Simulator statistics: "+stats);
assertEquals("MessageReceived", "1", stats.get("MessageReceived"));
assertEquals("MessageSent", "0", stats.get("MessageSent"));
}

public void testSendNetworkTriggeredData() throws Exception {
    String data = "Hello at " + System.currentTimeMillis();
    String from = "tel:1234";
    String to = "tel:7878";
    String correlator = "1234567890";

    assertTrue(isRunning(NETEX_SIMULATOR_MBEAN));
    resetStatistics(NETEX_SIMULATOR_MBEAN);

    wlngJmx.invokeOperation(EXAMPLE_PLUGIN_MBEAN, "connect");

    start(SESSION_MBEAN);
    assertTrue(isRunning(SESSION_MBEAN));

    putParameter(NOTIF_MANAGER_MBEAN, "url",
"http://localhost:8001/example/NotificationManager");
    putParameter(NOTIF_MANAGER_MBEAN, "start.address", "tel:7878");
    putParameter(NOTIF_MANAGER_MBEAN, "start.correlator", correlator);
    putParameter(NOTIF_MANAGER_MBEAN, "start.endpoint",
"http://localhost:13444/axis/services/Notification");
    putParameter(NOTIF_MANAGER_MBEAN, "stop.correlator", correlator);
    start(NOTIF_MANAGER_MBEAN);
}
```

```
start(NOTIF_MBEAN);

putParameter(NETWORK_TRIGGERED_MBEAN, "data", data);
putParameter(NETWORK_TRIGGERED_MBEAN, "fromAddress", from);
putParameter(NETWORK_TRIGGERED_MBEAN, "toAddress", to);

clearResults(NOTIF_MBEAN);

execute(NETWORK_TRIGGERED_MBEAN);

Thread.sleep(2000);

stop(NOTIF_MBEAN);
stop(NOTIF_MANAGER_MBEAN);
stop(SESSION_MBEAN);

Map<String,String> stats = listAllStatistics(NETEX_SIMULATOR_MBEAN);
System.out.println("Simulator statistics: "+stats);
assertEquals("MessageReceived", "0", stats.get("MessageReceived"));
assertEquals("MessageSent", "1", stats.get("MessageSent"));

List<Map<String,String>> results = listAllResults(NOTIF_MBEAN);
System.out.println("Notification results: "+results);
assertEquals("Correlator", correlator,
results.get(0).get("Correlator"));

assertEquals("From Address", from, results.get(0).get("From Address"));
assertEquals("Data", data, results.get(0).get("Data"));
```

Using the Platform Test Environment

}

}

Service Interceptors

The following sections give a high-level overview of service interceptors and describe both the out-of-the-box interceptors that ship with Network Gatekeeper and how to develop your own custom interceptors:

- [Overview](#)
- [Interceptor Decisions and Request Flow](#)
 - [Decisions](#)
 - [Flow Control](#)
 - [Changing the invocation order](#)
- [Standard Interceptors](#)
 - [Retry functionality for plug-ins](#)
- [Custom Interceptors](#)
 - [Developing Custom Interceptors](#)
 - [Deploying Custom Interceptors](#)

Overview

Interceptors are used to:

- Provide a mechanism to intercept and manipulate a request flowing through any arbitrary Communication Service in Network Gatekeeper

- Supply an easy way to modify the request flow
- Simplify the routing mechanism for plug-ins
- Centralize policy and SLA enforcement

Some typical use cases for interceptors are to:

- Deny a request if the user does not subscribe to a particular service in the application layer.
- Deny a request if a PIN is not valid
- Verify that a request's parameters are valid
- Perform argument manipulation like aliasing

A set of standard interceptors are provided out-of-the-box. Some are required, while others provide extra functionality. In addition, custom interceptors can be developed.

Interceptor Decisions and Request Flow

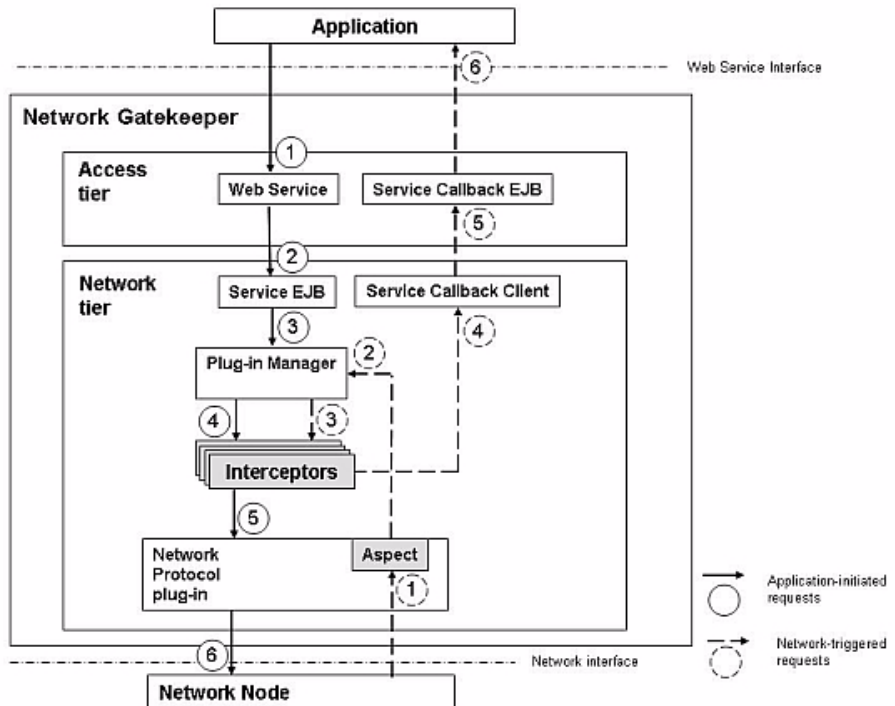
An interceptor makes a decision whether to permit, deny or stay neutral to a particular request: see [Decisions](#). The Plug-in Manager is responsible for calling the first interceptor in the chain of interceptors as defined in the interceptor configuration file: see [Flow Control](#). When changing the chain of interceptors, the interceptor module normally needs to be redeployed: see [Changing the invocation order](#).

Decisions

For application-initiated requests, the Plug-in Manager is called automatically by the service EJB for the application-facing interface. For network-triggered requests, the Plug-in Manager is called by an aspect that is woven prior to calling the service callback EJB for the application-facing interface.

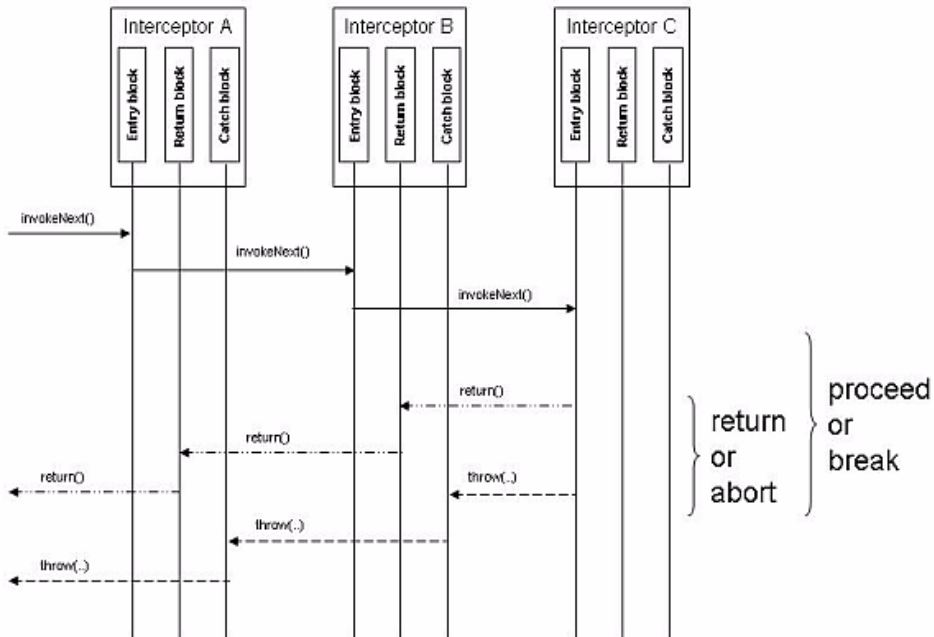
[Figure 11-1](#) illustrates where interceptors are triggered, both for application-initiated requests and for network-triggered.

Figure 11-1 Interceptors and the request flow



The interceptor chain is invoked at the point-cut that is a Java representation of the application-facing interface. Note that some application-initiated requests are not necessarily propagated to the network, and some network-triggered requests are not necessarily forwarded to the service callback client.

Each interceptor is responsible for deciding whether to continue to *proceed* down the chain of interceptors or to break it. The interceptor has two ways to break the chain, either to *return* or to *abort*.

Figure 11-2 Proceeding or breaking the interceptor chain

When the decision is to:

- *Proceed*, the request is passed on to the next interceptor in the chain and ultimately to the network protocol plug-in or to the application. When the request is returned from either one of these, the return path traverses the interceptors that were used in the calling path, making it possible to manipulate the request in the return path and ultimately return to the originator of the request, the application or the network node.
- *Return*, the request is rolled back through the previous interceptors using a regular return statement, making it possible for the previous interceptors to manipulate the request in the rollback path and ultimately return to the originator of the request, the application or the network node.
- *Abort* the request, is rolled back through each interceptors' exception catch-block rather than returning in a regular mode.

- For application-initiated requests the exception is reported back to the application. It is possible to reuse the exception catalogue to map the exception thrown by the interceptor to an exception defined by the application-facing interface.
`com.bea.wlcp.wlng.api.plugin.DenyPluginException` should be used by the interceptors for this scenario.
- For network-triggered requests, it is the responsibility of the plug-in to act on the thrown exception.

The interceptors have access to context data for the request. The actual data that is available depends on the context of the request. In general, the data available is the data that is defined by the application-facing interface, and includes the following items:

- The `RequestContext` for the request, including:
 - Service provider account ID.
 - Application account ID.
 - Application User ID.
 - Transaction ID.
 - Session ID.
 - A Java Map containing arbitrary request-specific data.
- The type of plug-in targeted by the request for (application-initiated requests).
- The type of object targeted by the request (network-triggered requests).
- The method targeted by the request.
- The arguments that will be used in the method targeted by the request.
- The set of `RequestInfo` available by the request, including:
 - method name.
 - arguments to the method.
 - plug-in type.
- A list of plug-ins that matches the specified `RequestInfo`.
- The interception point: is the request is network-triggered or is it application-initiated.

The following data can be set by the interceptor:

- In the `RequestContext`:

- Session ID.
- Transaction ID.
- Java Map.
- A list of plug-ins that matches the specified RequestInfo.
- Arguments to the method.

Flow Control

The invocation order of interceptors is defined in an XML-based configuration file that contains the interceptors: see [Standard Interceptors](#).

Each interceptor is identified by the class name of the entry point of the interceptor, that is, the class that implements the Service Provider Interface (SPI) `Interceptor`.

The configuration file, which is expressed in XML, contains the tags described in [Table 11-1](#).

Table 11-1 Description of interceptor configuration file

Tag	Description
<interceptor-config>	Main tag. Contains zero or more <position> tags.
<position>	<p>Contains one or more <interceptor> tags.</p> <p>Has an attribute name which is either:</p> <ul style="list-style-type: none">• MT_NORTH, which indicates that all <interceptor> tags encapsulated by this tag are valid for application-initiated (mobile terminated) requests.• MO_NORTH, which indicates that all <interceptor> tags encapsulated by this tag are valid for network-triggered (mobile originated) requests. <p>An interceptor may be present in both.</p>
<interceptor>	<p>Has the following attributes:</p> <ul style="list-style-type: none">• class, which identifies the class for the interceptor implementation, see above.• index, which indicates the invocation order relative to other interceptors within the same <position> tag. The order is ascending. Must be unique.

Listing 11-1 Example of an interceptor configuration file

```

<?xml version="1.0" encoding="UTF-8"?>
<interceptor-config xmlns="http://www.bea.com/ns/wlmg/30"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.bea.com/ns/wlmg/30 config.xsd">
  <position name="MT_NORTH">
    <interceptor class="com.bea.wlcp.wlmg.interceptor.EnforceApplicationState"
      index="100"/>
    <interceptor class="com.bea.wlcp.wlmg.interceptor.EnforceSpAppBudget"
      index="200"/>
    <interceptor
      class="com.bea.wlcp.wlmg.interceptor.ValidateRequestUsingRequestFactory"
      index="300"/>
    <interceptor class="com.bea.wlcp.wlmg.interceptor.CreatePluginList"
      index="400"/>
    <interceptor class="com.bea.wlcp.wlmg.interceptor.RemoveInactivePlugin"
      index="500"/>
    <interceptor
      class="com.bea.wlcp.wlmg.interceptor.RemoveInvalidAddressPlugin" index="600"/>
    <interceptor
      class="com.bea.wlcp.wlmg.interceptor.FilterPluginListUsingCustomMatch"
      index="700"/>
    <interceptor class="com.bea.wlcp.wlmg.interceptor.RoundRobinPluginList"
      index="800"/>
    <interceptor class="com.bea.wlcp.wlmg.interceptor.EnforceNodeBudget"
      index="900"/>
    <interceptor
      class="com.bea.wlcp.wlmg.interceptor.InvokeServiceCorrelation" index="1000"/>
    <interceptor
      class="com.bea.wlcp.wlmg.interceptor.FindAndValidateSLAContract"
      index="1100"/>
    <interceptor class="com.bea.wlcp.wlmg.interceptor.CreatePolicyData"
      index="1200"/>
  </position>
</interceptor-config>

```

Service Interceptors

```
<interceptor
class="com.bea.wlcp.wlng.interceptor.CheckMethodParametersFromSLA"
index="1300"/>

<interceptor
class="com.bea.wlcp.wlng.interceptor.EnforceBlacklistedMethodFromSLA"
index="1400"/>

<interceptor
class="com.bea.wlcp.wlng.interceptor.InjectValuesInRequestContextFromSLA"
index="1500"/>

<interceptor class="com.bea.wlcp.wlng.interceptor.EvaluateILOGPolicy"
index="1600"/>

<interceptor class="com.bea.wlcp.wlng.interceptor.InvokePlugin"
index="1700"/>

</position>

<position name="MO_NORTH">

<interceptor class="com.bea.wlcp.wlng.interceptor.EnforceApplicationState"
index="100"/>

<interceptor
class="com.bea.wlcp.wlng.interceptor.InvokeServiceCorrelation" index="200"/>

<interceptor class="com.bea.wlcp.wlng.interceptor.CreatePolicyData"
index="300"/>

<interceptor class="com.bea.wlcp.wlng.interceptor.EvaluateILOGPolicy"
index="400"/>

<interceptor class="com.bea.wlcp.wlng.interceptor.InvokeApplication"
index="500"/>

</position>
</interceptor-config>
```

Each interceptor is responsible for calling the next interceptor in the chain, as opposed to being invoked by a delegator. This means that:

- For application-initiated request, the interceptors can change and add request-specific data. This data is then propagated to the next interceptor and ultimately to the network protocol plug-in. When the request returns from the plug-in, the data can be changed as the request is returning through the invocation chain.

- For network-triggered request, the interceptors can change and add request-specific data. This data is then propagated to the next interceptor and ultimately to the application. When the request returns from the application, the data can be changed as the request is returning through the invocation chain.

This is useful for aliasing of data, where the interceptor anonymizes request data such as telephone numbers so that an application is not aware of the true subscriber telephone number.

For application-initiated requests, the last interceptor in the chain is responsible for calling the plug-in. The out-of-the-box interceptor [InvokePlugin](#) does this.

For network-triggered requests, the last interceptor in the chain is responsible for calling the callback service EJB, which calls the application. The out-of-the-box interceptor [InvokeApplication](#) does this.

In either scenario, the first interceptor is called by the Plug-in Manager. The Plug-in Manager is, for application-triggered requests, invoked by the service EJB. For network triggered requests, the Plug-in Manager is invoked by an aspect applied to the north interface of the plug-in.

Changing the invocation order

As described in [Flow Control](#), the invocation order of the interceptors is defined in the interceptor configuration file, see [Table 11-3](#).

To rearrange the invocation chain, explode the ear file, edit the config.xml file and change the attribute `index` in the tag `<interceptor>`. Repackage the ear file and deploy it.

To exclude an interceptor chain, explode the ear file and delete or comment out the `<interceptor>` tag for it. Repackage the ear file and deploy it.

Always use the `interceptors.ear` deployed on the Administration server as the master and use standard WebLogic procedures to redeploy the application `interceptor.ear` to all servers in the network tier cluster from the Administration server.

Standard Interceptors

Below is a description of the interceptors that are available out of the box as a part of Network Gatekeeper. The name of the interceptor in the configuration file is the fully qualified class name. That is, it is prefixed with `com.bea.wlcp.wlmg.interceptor`.

Table 11-2 Out-of-the-box interceptors

Interceptor	Description
EnforceApplicationState	
	Enforces the application state. Verifies that the application with which the request is related has established a session with Network Gatekeeper.
EnforceSpAppBudget	
	Enforces the budget defined in the service provider group SLA and application group SLA. Is related to the SLA tag <rate> in <methodRestrictions>: see Defining Service Provider Level and Application Level Service Agreements .
ValidateRequestUsingRequestFactory	
	Validates the request using the RequestFactory corresponding to the type of plug-in the request is intended for. See description of the class RequestFactory .
CreatePluginList	
	Creates a list of plug-ins that are capable of handling the given request.
RemoveInactivePlugin	
	Removes any plug-in that is not active from the current plug-in list. CreatePluginList must have been invoked prior to this.
RemoveInvalidAddressPlugin	
	Matches configured plug-in routes with plug-ins. Removes any plug-in which does not support the address provided in the request from the current plug-in list. CreatePluginList must have been invoked prior to this.
FilterPluginListUsingCustomMatch	
	Invokes the custom match method of each plug-in in the current plug-in list. The custom match method either removes the plug-in from the current plug-in list or marks it as required. CreatePluginList must have been invoked prior to this.

Table 11-2 Out-of-the-box interceptors

Interceptor	Description
RoundRobinPluginList	<p>Performs a round-robin of the list of available plug-ins. This is not a strict round-robin, but a function of the number of plug-ins that match the request and the number of destination or target addresses in the request. If these parameters are consistent, a true round-robin is performed.</p> <p>CreatePluginList must have been invoked prior to this.</p>
EnforceNodeBudget	<p>Enforces all settings in the service provider node SLA and global node SLA, including validity of the dates and the budgets. See Writing Node SLAs.</p> <p>EnforceSpAppBudget must have been invoked prior to this.</p>
InvokeServiceCorrelation	<p>Invokes the service correlation feature, see Service Correlation.</p>
FindAndValidateSLAContract	<p>Enforces the existence of application level and service provider level SLAs for the given request. It also verifies that the dates given in the SLA are current. See Defining Service Provider Level and Application Level Service Agreements.</p>
CreatePolicyRequestData	<p>Creates the policy request data object needed by other interceptors.</p>
CheckMethodParametersFromSLA	<p>Checks and enforces that the request parameters are allowed as specified in the service provider group and application group SLAs.</p> <p>Is related to the SLA tags <parameterName> and <parameterValue> in <methodParameters>, see Defining Service Provider Level and Application Level Service Agreements.</p> <p>FindAndValidateSLAContract and CreatePolicyRequest must have been invoked prior to this.</p>
EnforceBlacklistedMethodFromSLA	

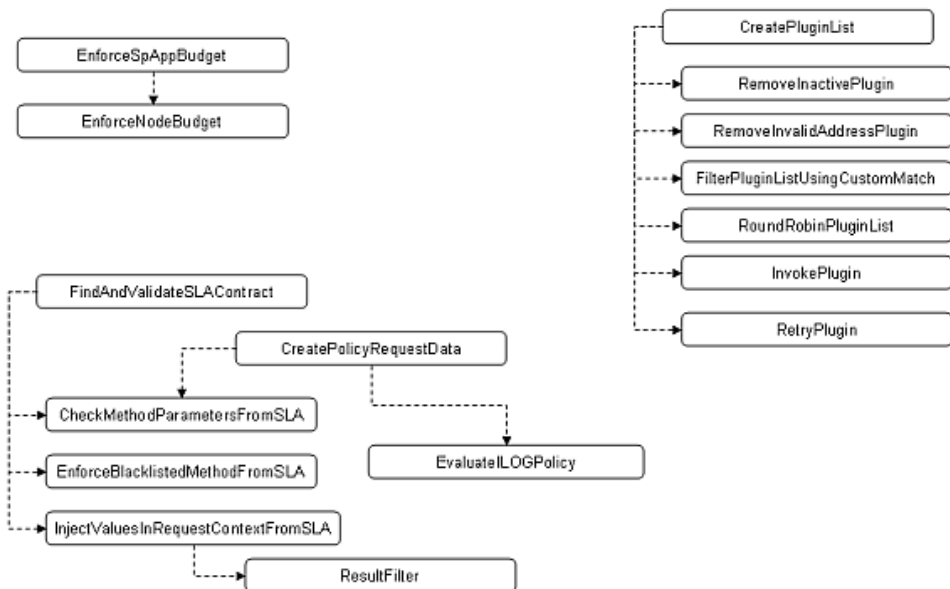
Table 11-2 Out-of-the-box interceptors

Interceptor	Description
	<p>Enforces the method blacklist as specified in the service provider group and application group SLAs.</p> <p>Is related to the SLA tag <blacklistedMethod> in <methodAccess>. See Defining Service Provider Level and Application Level Service Agreements.</p> <p>FindAndValidateSLAContract must have been invoked prior to this.</p>
InjectValuesInRequestContextFromSLA	<p>Adds any optional request context attribute as specified in the service provider group and application group SLAs.</p> <p>Is related to the SLA tags <attributeName>, <attributeValue>, and <contextAttribute> in <requestContext>. See Defining Service Provider Level and Application Level Service Agreements.</p> <p>FindAndValidateSLAContract must have been invoked prior to this.</p>
EvaluateILOGPolicy	<p>Evaluates any custom ILOG policy rules.</p> <p>CreatePolicyRequestData must have been invoked prior to this.</p>
InvokePlugin	<p>Invokes the plug-in(s). This should be the last interceptor for an application-initiated (mobile terminated) request.</p> <p>CreatePluginList must have been invoked prior to this.</p>
InvokeApplication	<p>Invokes the Application via the service callback EJB. This should be the last interceptor for an network-triggered (mobile originated) request.</p>
RetryPlugin	<p>Performs retries of request. See Retry functionality for plug-ins.</p> <p>CreatePluginList must have been invoked prior to this.</p>

Table 11-2 Out-of-the-box interceptors

Interceptor	Description
ResultFilter	<p>Applies result filters as specified in the service provider group and application group SLAs.</p> <p>Relates to the SLA tag <resultRestriction>. See Defining Service Provider Level and Application Level Service Agreements.</p> <p>InjectValuesInRequestContextFromSLA must have been invoked prior to this.</p>

Note: Some interceptors must be invoked before others can be invoked. A quick overview of the necessary sequences is seen in [Figure 11-3](#)

Figure 11-3 Required Interceptor Sequences

All out-of-the-box interceptors are classes packaged in `$DOMAIN_HOME/interceptors.ear`.

Below is a description of the contents of this ear:

Table 11-3 Contents of interceptor.ear

Path	Content
/	
	dummy.war Empty war file. Present in order to deploy the interceptors. Do not remove or change.
/APP-INF/classes/	
	config.xml Interceptor configuration file. See Flow Control .
	config.xsd Schema for config.xml
/APP-INF/classes/com/bea/wlcp/wlmg/interceptor/	
	Classes for the out-of-the-box interceptors, see Table 11-2 . Do not change the content of this directory.
/APP-INF/classes/com/bea/wlcp/wlmg/interceptor/deploy/	
	Infrastructure for the interceptor functionality. Do not change the content of this directory.
META-INF/	
	MANIFEST.MF Manifest file for the interceptor infrastructure.
	application.xml Deployment descriptor. Do not edit or remove.
	weblogic-application.xml WebLogic extensions to application.xml. Do not edit or remove.
WEB-INF/	
	No content.

Retry functionality for plug-ins

The RetryPlugin interceptor handles retry functionality for plug-ins. The retry is attempted among the plug-ins that were chosen based on the data provided in the request. Retries are only performed among the plug-ins in the same Network Gatekeeper instance.

The RetryPlugin is triggered when a plug-in throws a RetryPluginException. This exception is captured by the RetryPlugin interceptor, which removes the plug-in that threw the exception from the list of chosen plug-ins and calls the next interceptor in the chain.

The different decision scenarios are described below.

If the RequestInfo objects in the RequestContext are associated with:	The RetryPlugin interceptor:
PluginHolder objects that are marked as optional	removes the failed RequestInfo from the RequestContext and the next interceptor in the chain is invoked.
PluginHolder objects that are marked as required	treats the request itself as failed. No retry is performed, and an exception is thrown.
some PluginHolder objects that are marked as optional, and some that are marked as required	removes the RequestInfo objects that are associated with the PluginHolder objects that are marked as optional from the RequestContext and the next interceptor in the chain is invoked.

The following out-of-the-box plug-in throws the RetryPluginException:

- Subscriber Profile/LDAPv3.

Custom plug-ins can use the infrastructure for retries as provided by the RetryPlugin interceptor. This exception should be thrown if the communication with the underlying network node fails, or if an unexpected error is reported back from the plug-in.

Custom Interceptors

Developing Custom Interceptors

An interceptor implements the interface

```
com.bea.wlcp.wlmg.api.interceptor.Interceptor.
```

This interface defines the method:

```
Object invoke(com.bea.wlcp.wlmg.api.interceptor.Context context) throws  
Exception;
```

The interceptor is responsible for invoking the next interceptor in the invocation chain using the method:

```
Object com.bea.wlcp.wlmg.api.interceptor.Context.invokeNext(Interceptor  
current) throws Exception;
```

Since the interceptors call each other, the normal case would be just to return the object that was returned by the called interceptor. But in some cases, the returned object may be changed in order to do, for example, aliasing.

The decisions within the interceptor are expressed in these ways:

- To *proceed*, continue down the invocation chain by calling the next interceptor.
- To break the chain due to a violation: for example a parameter in the request is out-of-bounds, or that usage policies are violated. This *aborts* the request throwing a `PluginDenyException`.
- To break the chain because the request has been fulfilled (for example because there is no need to call the plug-in or the application in order to fulfill the needs of the request), simply *return* the request.

See [Interceptor Decisions and Request Flow](#).

Note: The interceptor must be thread safe.

[Listing 11-2](#) illustrates a very basic interceptor.

Listing 11-2 Example interceptor

```
import com.bea.wlcp.wlmg.api.interceptor.Interceptor;
```

```

public class SampleInterceptor implements Interceptor {

    private final int ABORT = 0;

    private final int RETURN = 1;

    public Object invoke(Context ctx) throws Exception {

        int decision = // Logic that evaluates the request and makes a decision.
        if (decision == ABORT) {

            throw new Exception();

        } else if (decision == RETURN) {

            Object returnValue = // Define a returnValue here if desired.

            return returnValue;

        } else {

            Object returnValue = ctx.invokeNext(this);

            // Define a new returnValue here if desired, for example for aliasing.

            return returnValue;

        }

    }

}

```

All necessary classes are available in the package:

com.bea.wlcp.wlmg.api.interceptor located in
 \$BEA_HOME/wlmg_pds400/lib/api/wlmg.jar

As an alternative to embedding the interceptor in `interceptors.ear` and defining the invocation order in `/APP-INF/classes/config.xml` it is possible to put the new interceptor in a separate ear file. Using this alternative, the interceptor must register the interceptor using the `InterceptorManager`, which is retrieved using the `InterceptorManagerFactory`.

When registering the interceptor manually, data corresponding to the data set in `/APP-INF/classes/config.xml` in `interceptors.ear` is supplied as parameters to the method:

```
void register(Interceptor interceptor, InterceptionPoint ip, int index);
```

in the `InterceptorManager` interface.

The attribute `name` in the tag `<position>` corresponds to the argument `ip`, the attribute `index` in the tag `<interceptor>` corresponds to the argument `index`.

[Listing 11-3](#) shows an example of how to register an interceptor manually.

Listing 11-3 Manually registering an interceptor

```
InterceptorManager im = InterceptorManagerFactory.getInstance(); // Get manager
im.register(myInterceptor, InterceptorManager.MT_NORTH, myIndex); // Register
im.update(); // Changes do not take effect until update() is called
```

Deploying Custom Interceptors

To deploy the interceptor in the common interceptor ear file, explode the `interceptors.ear` file and put the class files for the interceptor in `/APP-INF/classes`. Add a new `<interceptor>` tag with the attribute `class` referring to the entry point of the interceptor and a numeric value in the attribute `index` that corresponds to the location in the interceptor invocation chain.

For example:

If the interceptor main class is `com.acompany.interceptor.DoStuff`, the class `DoStuff` should be inserted into `interceptors.ear` in

`/APP-INF/classes/com/acompany/interceptor`, and the corresponding entry in `/APP-INF/classes/config.xml` shall be

```
<interceptor class="com.acompany.interceptor.DoStuff" index="1150"/>
```

See [Flow Control](#) for more information about `/APP-INF/classes/config.xml`. See [Standard Interceptors](#) to get information about where in the invocation chain to insert the new interceptor.

If deploying the interceptor in a separate ear, always deploy it using the Administration server and use standard WebLogic procedures to deploy the application to all servers in the cluster from the Administration server.

Service Interceptors

Subscriber-centric Policy

Making subscriber personalization easy and offering superior subscriber data protection is key to growing and maintaining a loyal subscriber base. The Platform Development Studio offers a straightforward way to extend the power of Network Gatekeeper's flexible policy-based control to the operator's subscriber base. The mechanism can be divided into three parts:

- [Service Classes and the Subscriber SLA](#)
- [The Profile Provider SPI and Subscriber Contracts](#)
- [Subscriber Policy Enforcement](#)

Note: There is an example Profile Provider in `<bea_home>/wlng_pds400/example`

Service Classes and the Subscriber SLA

The first step in adding subscriber-centric policy to Network Gatekeeper is to create a Subscriber SLA. This is an XML file based on the `sub_sla_file.xsd` schema.

Note: The schema file can be found in the `wlng.jar` file located in the `<bea_home>/wlng_pds400/lib/wlng` directory.

The SLA is used to define classes of service in the context of existing Service Provider and Application Groups. (For more information on Service Provider and Application Groups, see [“Managing Application Service Providers”](#) in *Concepts and Architectural Overview*, a separate document in this set.) These *service classes* can then be associated with subscribers, based on their preferences and permissions, defining individualized relationships between subscribers and Service Provider and Application Group functionality.

The <reference> tag

The <reference> tag specifies the operator's already-established Application and Service Provider Groups that are to be associated with this service class. There are two reference types that define the groups: the `ApplicationGroupReference` and the `ServiceProviderGroupReference`. In addition there are two additional reference types, the `ServiceReference` and the `MethodReference` that indicate specific service interfaces and methods, respectively, covered by those groups. In the [Listing 12-1](#) snippet, the service class `news_subscription` is defined. Evaluation of matches in the class occurs using the following rules:

- If no reference type is specified, everything of that type is a match
- Two or more entries of the same reference type creates an OR relationship
- The default relationship is AND

So, in the case of [Listing 12-1](#), the class covers any request that matches:

- Any of the service interfaces of the `silver_app_group`
(No `ServiceReference` type is specified, so everything is a match)
- **OR** the `gold_app_group`
(Two `ApplicationGroupReference` entries creates an **OR**)
 - **AND** the `SendSMS` service interface of the `gold_app_group`
(The default relationship)
 - **AND** the `content_sp_group`
(The default relationship)
 - **AND** the `SendSMS` service interface of the `content_sp_group`
(The default relationship)
 - **AND** either the `sendSms` **OR** the `getSmsDeliveryStatus` methods
(Two `MethodReference` entries creates an **OR**)

Listing 12-1 The <reference> element

```

<ServiceClass name="news_subscription">
    <references>
        <ApplicationGroupReference id="silver_app_group"/>
        <ApplicationGroupReference id="gold_app_group">
            <ServiceReference
serviceInterface="com.bea.wlcp.wlng.px21.plugin.SendSmsPlugin"/>
        </ApplicationGroupReference>
        <ServiceProviderGroupReference id="content_sp_group">
            <ServiceReference
serviceInterface="com.bea.wlcp.wlng.px21.plugin.SendSmsPlugin">
                <MethodReference methodName="sendSms" />
                <MethodReference methodName="getSmsDeliveryStatus" />
            </ServiceReference>
        </ServiceProviderGroupReference>
    </references>

```

Use of the empty tag, <references/>, matches everything.

The <restriction> tag

In addition to the <reference> tag, service classes may have a <restriction> tag. This tag is used to attach default rates and quotas that are used to create budgets for the classes. These rates and quotas can be replaced in specific contracts.

Note: The XSD requires you either to specify a rate/quota restriction or to use the <restrictAllType/> tag.

Listing 12-2 The <restriction> tag

```

<restriction>

```

```
<rate>

    <reqLimit>5</reqLimit>

    <timePeriod>1000</timePeriod>

</rate>

<quota>

    <qtaLimit>600</qtaLimit>

    <days>3</days>

    <limitExceedOK>true</limitExceedOK>

</quota>

</restriction>
```

These tags function exactly as they do in the other SLAs in Network Gatekeeper. For more information on these tags, see the **Contract structure** section of the [“Defining Service Provider Group and Application Group SLAs”](#) chapter of *Managing Accounts and SLAs*, a separate document in this set. If the `<limitExceedOK>` tag is set to true, the request is allowed even when if quota has been exceeded, but an alarm (Alarm id 200000) is fired

There is also a `<restrictAllType/>` tag. This tag, as its name implies, denies access to all requests.

Managing the Subscriber SLA

There are three management methods in the Service Level Agreement MBean for managing a Subscriber SLA. They are covered in detail in the [“Managing a Subscriber SLA”](#) chapter of *Managing Accounts and SLAs*, a separate document in this set. The methods allow you to load a Subscriber SLA as a string, to load a Subscriber SLA from a URL, and to retrieve a loaded Subscriber SLA.

The Profile Provider SPI and Subscriber Contracts

Once the Subscriber SLA is established, the various service classes it defines must be associated with individual subscribers. The combination of a subscriber (identified by URI) and a service class is called a *subscriber contract*. A subscriber (a URI) can have multiple subscriber contracts associated with it.

The subscriber contract object contains a URI designating the subscriber and the service class type with which it is associated. It also contains an expiration time, represented as a `java.util.Date`.

Note: The subscriber contract constructor will throw an exception if the URI, service class type, and expiration time are not specified.

The subscriber contract may also replace the default rate and/or quota settings in the service class, or set this subscriber to `RestrictAll`, that is, to deny access for all requests.

The operator or integrator is responsible for creating the mechanism, a Profile Provider, that supplies these subscriber contracts.

Note: All class files related to creating Profile Providers are in the `com.bea.wlcp.wlng.spi.subscriberdata` package, and can be found in the `wlng.jar` file in the `<bea_home>/wlng_pds400/lib/wlng` directory. The JavaDoc for the files can be found in the `<bea_home>/wlng_pds400/doc/javadoc` directory. An example implementation can be found in the `<bea_home>/wlng_pds400/example/profile_providers` directory. This sample implementation assumes the use of a properties file to assign subscriber URIs to particular service classes. An example properties file, `exampleSubscriberContractMappingFile.properties`, can be found in the same directory.

The Profile Provider must implement the Profile Provider SPI. The SPI defines three methods;

- `init`: Network Gatekeeper initializes the Profile Provider by passing in a list of the service classes that are defined in the Subscriber SLA and a list of any previously defined subscriber contracts. The Provider returns a list of updated subscriber contracts.
- `contractExpired`: Network Gatekeeper sends the Provider a list of service classes and a list of expired contracts. The Provider returns an updated list of contracts for those that have expired. The Provider can remove or add contracts to the returned list.
- `serviceClassesUpdated`: Whenever the Subscriber SLA is updated, and the service classes are thus modified, Network Gatekeeper sends the Provider a list of the updated service classes and a list of all current contracts. The Provider returns an updated list of contracts. The Provider can make any necessary updates to the subscriber contracts.

The Profile Provider implementation must have a public constructor with no parameters or a static method which returns `ProfileProvider`.

Note: There is a sample Profile Provider in `<bea_home>/wlng_pds400/example`.

Deploying the Custom Profile Provider

Once the `ProfileProviderImpl` has been created, the `.jar` file containing it must be added to the `app-inf/lib` directory of the `profile_providers.ear` file, which can be found in the `<bea_home>/wlmg_pds400/integration/profile_provider` directory. You must also modify the `app-inf/classes/ProfileProviders.prop` file, adding a line containing the package and implementation file name of each of your providers (multiple providers are possible). For example:

```
com.mycompany.mypackage.MyProfileProviderImpl
```

Once the EAR is modified, it can be deployed in the normal manner. For more information on deploying EAR files in Network Gatekeeper, see the [“Deployment model for Communication Services and Container Services”](#) chapter in the *System Administrator’s Guide*, a separate document in this set.

Subscriber Policy Enforcement

Once the `providers.ear` is deployed, the singleton `SubscriberProfileService` initializes the Profile Provider(s) and receives the relevant subscriber contracts. It uses the Budget Service to create budgets for the contracts, based on the specified rates and quotas, and also creates and schedules a timer based on the expiration times in the contracts. Both the Subscriber SLA and the subscriber contracts are persisted using the Storage Service.

Note: For more information on budgets in Network Gatekeeper, see the [“Managing and Configuring Budgets”](#) chapter in the *System Administrator’s Guide*, a separate document in this set.

When a request from an application arrives at Network Gatekeeper, it passes through the Interceptor Stack for policy evaluation. The `EnforceSubscriberBudget` interceptor manages policy enforcement for subscriber contracts. The process within the interceptor has two phases:

- [Do Relevant Subscriber Contracts Exist](#)
- [Is There Adequate Budget for the Contracts](#)

Do Relevant Subscriber Contracts Exist

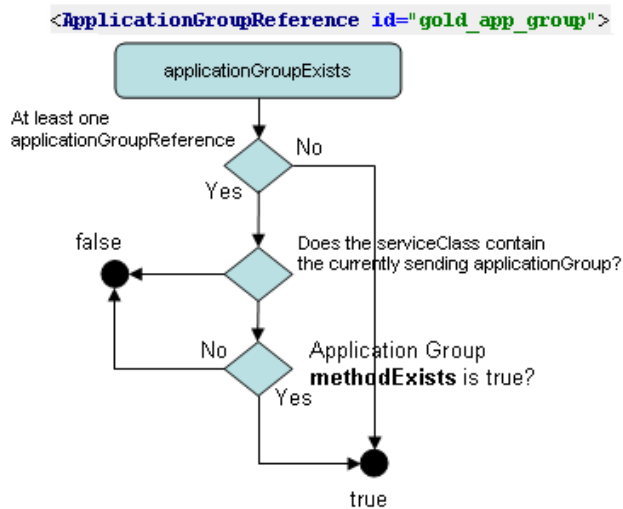
The first thing the interceptor must determine is whether one or more contracts exist that are relevant to the particular request that is being evaluated. The interceptor iterates through all the target URIs in the application request, and evaluates whether or not there are contracts in effect that it should enforce.

- If there are no contracts at all associated with a particular URI, the request is simply passed on to the next interceptor in the sequence.
- If there are contracts associated with a particular URI, a set of evaluations must be carried out. The figures below show the decision flow for the evaluations. All three sections must evaluate to `true` for there to be an enforceable contract.

Note: The XML snippets correspond to the relevant sections of [Listing 12-1](#):

- Is there an `ApplicationGroupReference` and is it relevant? See [Figure 12-1](#)

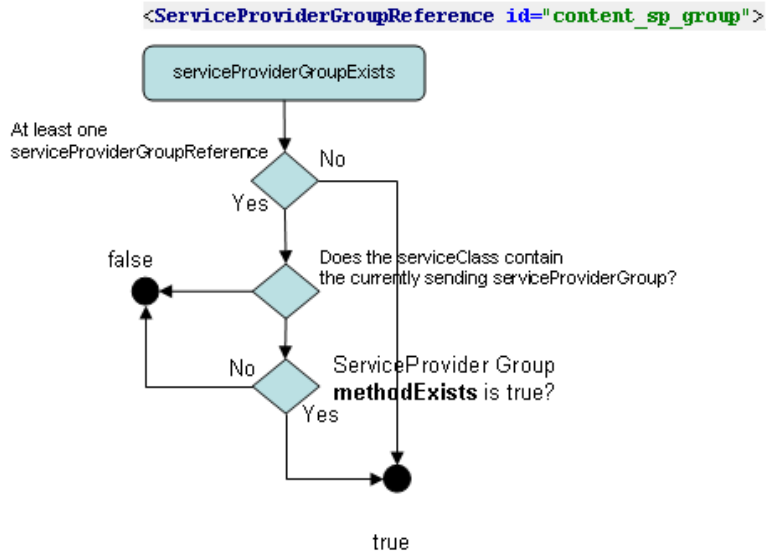
Figure 12-1 Application Group Reference Evaluation



Note: The evaluation for **methodExists** is covered in [Figure 12-3](#)

- Is there a `ServiceProviderGroupReference` and is it relevant? See [Figure 12-2](#).

Figure 12-2 Service Provider Group Reference Evaluation



Note: The evaluation for **methodExists** is covered in [Figure 12-3](#)

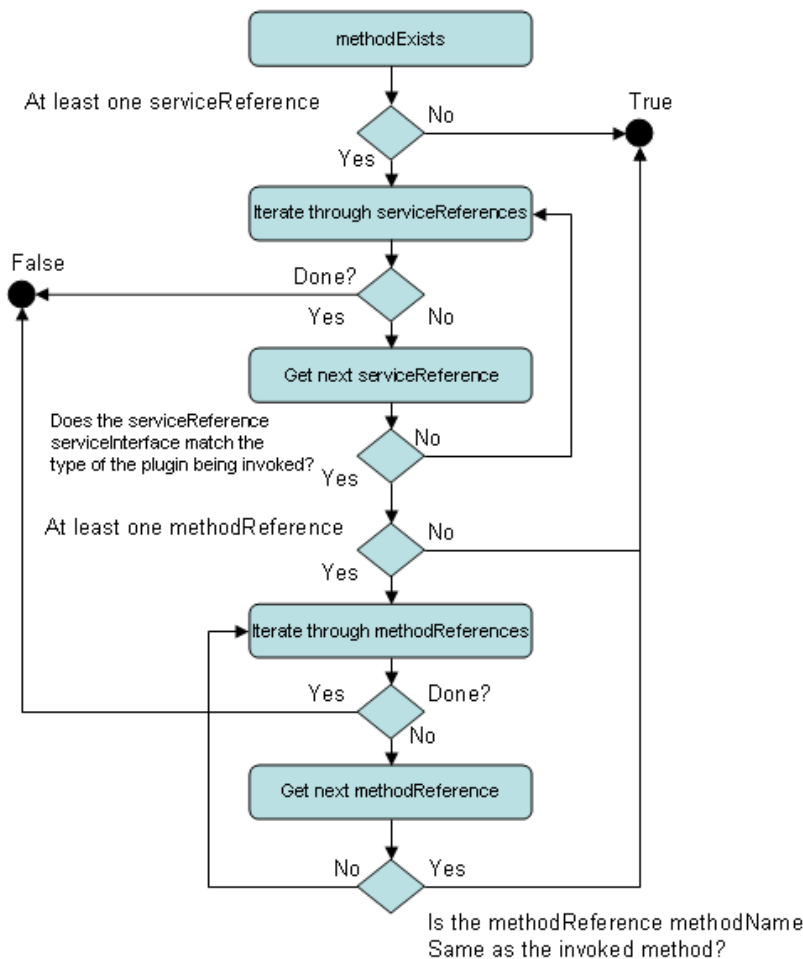
- Is there a Service Reference (and possibly a MethodReference) and are they relevant? See [Figure 12-3](#)

Figure 12-3 Service and Method Reference Evaluation

```

<ServiceReference serviceInterface="com.bea.wlcp.wlng.px21.plugin.SendSmsPlugin">
  <MethodReference methodName="sendSms" />
  <MethodReference methodName="getSmsDeliveryStatus" />
</ServiceReference>

```



Is There Adequate Budget for the Contracts

Once the interceptor determines that an enforceable contract exists, it first determines whether the contract includes a `<restriction>` tag set to `<restrictAll/>`. If so, the request is immediately denied, and processing on the request ceases.

If the `<restriction>` tag is not set to `<restrictAll/>`, the decision flow here is identical to the other budget evaluations that take place in Network Gatekeeper.

If there are no relevant contracts, or there are relevant contracts and there is adequate budget to cover them, budgets are adjusted as necessary and the request passes on to the next interceptor. If there are relevant contracts and there is not adequate budget to cover them, the request is denied.

Subscriber-centric Policy

Creating an EDR Listener and Generating SNMP MIBs

The following section describes how to create an external EDR listener.

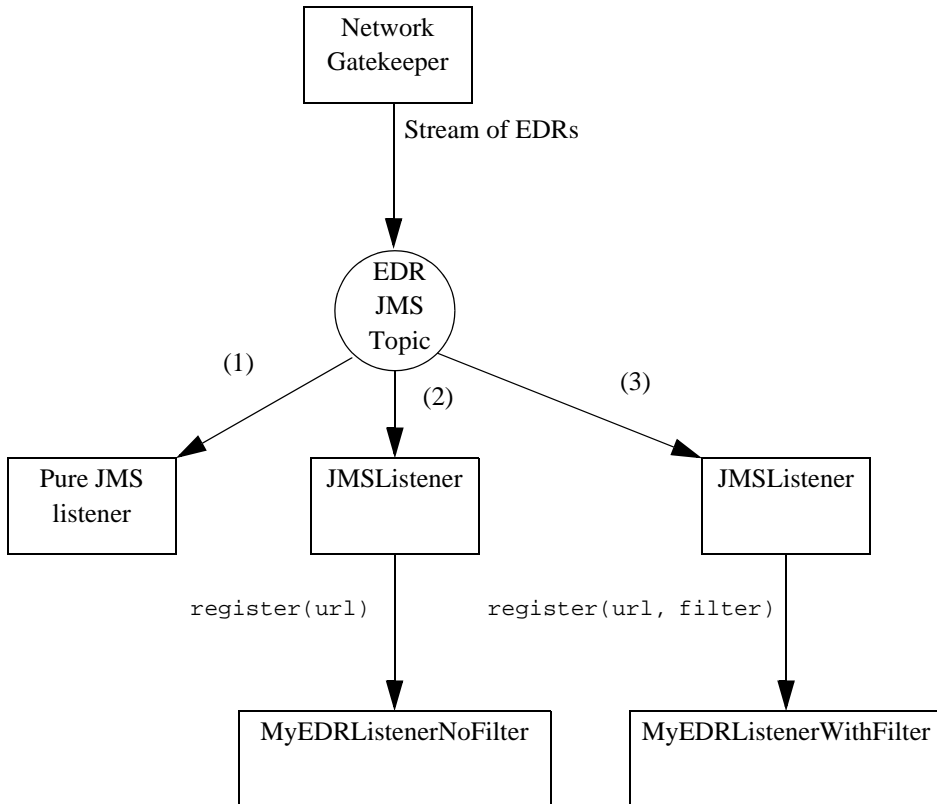
- Overview of External EDR listeners
 - Example using a pure JMS listener
 - Example using JMSListener utility with no filter
 - Using JMSListener utility with a filter
- Description of EDR listener utility
 - Class JMSListener
 - Class EdrFilterFactory
 - Class EdrData
 - Class ConfigDescriptor
 - Class EdrConfigDescriptor
 - Class AlarmConfigDescriptor
 - Class CdrConfigDescriptor
- Generating SNMP MIBs

Overview of External EDR listeners

External EDR listeners are JMS topic subscribers.

The diagram below illustrates three different ways of listening for EDRs as a JMS listener.

Figure 13-1 Flow for external EDR, alarm, and CDR listeners



EDRs are published externally using a JMS topic. This makes it possible to implement language-independent listeners anywhere on the network in a standard way. It is possible to implement an EDR listener in several ways:

- Alternative 1: Using a pure JMS listener. Implement the `javax.jms.MessageListener` interface. It is up to the implementation class to implement any filtering mechanism needed.

- Alternative 2: Using a subclass of JMSListener with no filter specified. In that case, the JMSListener class will use a tag, if available in the EDR, to filter the EDR into a specific category: EDR, alarm or CDR.
- Alternative 3: Using a subclass of JMSListener with a specified filter. This filter is used to perform the filtering. If a default filter is used to perform the same filtering as the WLNG, note that all classes used in the xml configuration files must be present in the current class loader. Otherwise, some EDRs will not be correctly filtered.

Example using a pure JMS listener

Listing 13-1 Example using a pure JMS listener

```
public class ClientJMSListener implements MessageListener {
    public void onMessage(Message msg) {
        // Extract the EdrData object or array
        if(o instanceof EdrData[]) {
            for(EdrData edr : (EdrData[])o) {
                //do something with each EDR
            }
        }
    }
}
```

Example using JMSListener utility with no filter

Listing 13-2 Example using a subclass of JMSListener with no filter specified

```
public class SampleEdrJMSListener extends JMSListener {
    public SampleEdrJMSListener(String url) throws Exception {
        // Register in the JMS topic. No filter is specified so
```

```
// the "tag" filtering mechanism will be used.
register(url);
}

@Override
public void onEdr(EdrData edr, ConfigDescriptor descriptor) {
    // The "tag" mechanism will filter the stream of EDRs according
    // to the internal WLNG filtering. To know which type of EDR is
    // actually provided in this method, we have to determine the
    // instance of the ConfigDescriptor as follow:
    if(descriptor instanceof EdrConfigDescriptor) {
        // do something with this EDR
    } else if(descriptor instanceof AlarmConfigDescriptor) {
        // do something with this alarm
    } else if(descriptor instanceof CdrConfigDescriptor) {
        // do something with this CDR
    }
}
}
```

Using JMSListener utility with a filter

Listing 13-3 Using a subclass of JMSListener with a specified filter

```
public class SampleEdrJMSListener extends JMSListener {

    public SampleEdrJMSListener(String url) throws Exception {
        // Register in the JMS topic. Use the default alarm filter.
    }
}
```



```

// Note that in this case all classes needed by the alarm.xml file
// must be in the current class loader in order for the filtering
// to work correctly.

register(url, EdrFilterFactory.createDefaultFilterForAlarm());
}

@Override

public void onEdr(EdrData edr, ConfigDescriptor descriptor) {
    // Only AlarmConfigDescriptor should be received here.
    // Just check before casting.
    if(descriptor instanceof AlarmConfigDescriptor) {
        ... do something with this alarm
    }
}
}
}

```

Note: When using the JMSListener class, make sure that any modification to an EDR, CDR, or alarms descriptor in network Gatekeeper is also updated in the edrjmslistener.jar file.

Description of EDR listener utility

The EDR listener utility contains a set of classes to use when creating an external JMS listener using the JMSListener.

The helper classes are found in the domain home directory in Network Gatekeeper, in:

\$ET_Home/lib/edrjmslistener.jar

Class JMSListener

Table 13-1 JMSListener

Method	Description
<code>public void register(String url)</code>	Registers the JMS listener to the EDR topic using no filter. The filtering will be done using the tagging mechanism. The parameter url specifies the URL of a network tier WLNG server.
<code>public void register(String url, EdrFilter filter)</code>	Registers the JMS listener to the EDR topic using the specified filter.
<code>public void onEdr(EdrData edr, ConfigDescriptor descriptor)</code>	Method that the subclass can override to get notified each time an EDR is received. The descriptor will be a subclass of ConfigDescriptor that will identify the type of EDR: either EdrConfigDescriptor, AlarmConfigDescriptor or CdrConfigDescriptor.

Class EdrFilterFactory

Table 13-2 EdrFilterFactory

Method	Description
<code>public static EdrFilter createDefaultFilterForEdr()</code>	Creates the default filter using in the WLNG to filter the EDRs using the edr.xml file embedded in the edrjmslistener.jar file.
<code>public static EdrFilter createDefaultFilterForAlarm()</code>	Creates the default filter using in the WLNG to filter the alarms using the alarm.xml file embedded in the edrjmslistener.jar file.
<code>public static EdrFilter createDefaultFilterForCdr()</code>	Creates the default filter using in the WLNG to filter the CDRs using the cdr.xml file embedded in the edrjmslistener.jar file.

Class EdrData

This class contains all the values that an EDR (alarm and CDR) have.

Table 13-3 EdrData

Method	Description
public String getValue(String key)	Gets the value associated with the specified key.
public List<String> getValues(String key)	Gets the values associated with the specified key.

Class ConfigDescriptor

This class is the parent class of EdrConfigDescriptor, AlarmConfigDescriptor and CdrConfigDescriptor.

Class EdrConfigDescriptor

This class contains the data that is specified in the descriptors in the edr.xml configuration file: the identifier and the description.

Table 13-4 EdrConfigDescriptor

Method	Description
public long getIdentifier()	Returns the identifier of the EDR.
public String getDescription()	Returns the description of the EDR.

Class AlarmConfigDescriptor

This class contains the data that is specified in the descriptors in the alarm.xml configuration file: the identifier, the severity and the description.

Table 13-5 AlarmConfigDescriptor

Method	Description
public long getIdentifier()	Returns the identifier of the alarm.
public String getSeverity()	Returns the severity of the alarm.
public String getDescription()	Returns the description of the alarm.

Class CdrConfigDescriptor

This class identifies a CDR. This descriptor does not contain any additional data.

Updating EDR configuration files

If you are using external EDR listeners, and the alarm, CDR, or EDR descriptors have been updated in Network Gatekeeper, the corresponding files need to be updated in `edrjmslistener.jar`. Update the corresponding xml file with the updated entries in the `edr` directory in `edrjmslistener.jar`.

Generating SNMP MIBs

Alarms can be forwarded as SNMP traps, see [Managing and Configuring the SNMP service](#) in *System Administrator's Guide*.

The MIB file that corresponds to the alarms can be generated using the ant task `mibgenerator` defined in `com.bea.wlcp.wlmg.ant.MIBGeneratorTask`.

The ant task is packaged in `$PDS_HOME/wlmg/lib/ant-mib-generator.jar`

There is an example build file that uses the an task in `$PDS_HOME/integration`

When the alarms descriptor is changed, a new MIB should be generated and distributed to the SNMP clients. Copy the contents of the alarm descriptor and paste it into an xml file. Use this xml file when generating the MIB file.

Converting Traffic Paths and Plug-ins to Communication Services

Traffic paths and network protocol plug-ins developed as extensions to Network Gatekeeper 3.0 can be converted to communication services and deployed in this release using the procedure described in this section.

Plug-ins and traffic paths developed for Network Gatekeeper 2.2 and earlier should be re-engineered in order to take full advantage of the improvements of the platforms.

A pre-requisite is to have the source code for the traffic path and plug-in that is to be converted.

- [Converting Network Protocol Plug-ins](#)
- [Converting Traffic Paths](#)
- [Checklist](#)

Converting Network Protocol Plug-ins

The procedure for converting a plug-in for an existing communication service is:

1. Generate a new plug-in using the Platform Development Studio Eclipse Wizard, see [Using the Eclipse Wizard](#).
2. Copy the `src` directory of the plug-in to be converted to the `src` directory of the new plug-in.
3. If no MBean class is declared, remove the `javadoc2annotation` target from the `build.xml` file for the new plug-in.

Converting Traffic Paths

The procedure for converting a traffic path to a communication service is:

1. Generate a new common service with the same settings as the traffic path to be converted using the Platform Development Studio Eclipse Wizard, see [Using the Eclipse Wizard](#).
2. Copy any customized part from the `traffic_path` directory for the traffic path to be converted to `common` directory for new communication service.
3. Copy the `src` directory of the plug-in to be converted to the `src` directory of the new plug-in.
4. If no MBean class is declared, remove the `javadoc2annotation` target from the `build.xml` file for the new plug-in.

Checklist

Below are a few items that should be verified during the conversion process:

- Make sure that the version used in the deploy targets in the main `build.xml` matches the one specified in the `common.xml` file.
- Make sure that the class specified in the property `plugin.class` defined in the `build.xml` for the plug-in is correct.
- Remove all references to `com.incomit.policy.DenyException` since it is not supported anymore.

Policy

For most installations of WebLogic Network Gatekeeper, the ability rapidly and accurately to evaluate the status of requests in terms of Policy, or rules governing a variety of service characteristics, is one of the most important features that the system offers.

Note: Some evaluations, such as enforcement of SLAs, are performed by the Interceptor Stack. See [Chapter 11, “Service Interceptors”](#) for more details. The Policy system described in this chapter allows you to add additional types of evaluation to the request flow, including adding rules to be used for the Callable Policy Web Service.

If you extend the Network Gatekeeper, particularly if you add a new Communication Service, you may also need to make changes in the Policy system to cover new functionality that you have added. This chapter provides a very high level description of the process by which policy requests are processed and though which new rules can be added. It covers:

- [Overview](#)
- [Policy Request Data](#)
- [Adding a New Rule](#)
- [Using RequestContext Parameters Defined in Service Level Agreements](#)

Overview

When an application service request arrives at the service interceptor `CreatePolicyRequestData`, its parameters are put in a `PolicyRequest` object. The service interceptor `EvaluateILOGPolicy` evaluates these custom policy rules. The rules themselves are written in the ILOG IRL language.

Policy Request Data

A `PolicyRequest` object has a standard form. The values in the object must be mapped to the variables in the Policy Rule that will be used to evaluate them. The Policy Request object can contain subsets of this standard data:

- `applicationID`: The Application ID of the requesting party
- `serviceProviderID`: The Service Provider ID of the requesting party
- `nodeID`: Used internally by Network Gatekeeper - ignore
- `serviceName`: The name of the software module in which the policy request originates. Used in the rules to match to service contracts in the SLAs and to look-up any rules specific to the service.
- `methodName`: The name of the method which the request wishes to have executed. Access to this method is what is being evaluated.
- `serviceCode`: The service code provided by the application, which is written to CDRs for tracking purposes
- `requesterID`: An additional ID that may be provided by the application for tracking purposes. (dependent on the northbound interface being used)

All of the above are Strings.

- `transactionID`: Used internally by Network Gatekeeper - ignore
- `noOfActiveSessions`: Used internally by Network Gatekeeper - ignore
- `timeStamp`: The time the request was fed to the Policy Engine.
- `reqCounter`: The number of target addresses in the request. If only one target address is used in the request this value is set to 1. If using multiple target addresses in the request, it is the number of target addresses.

All of these are Longs.

In addition to these standard values, Policy Request objects contain all the parameters passed in from the application in its initial request, as `AdditionalParameters`, an array of `AdditionalDataValue`. An `AdditionalDataValue` consist of a name-value pair. The following data types can be defined in an `AdditionalDataValue` object.

- `intValue(int val)`: Integer values

- `longValue(long val)`: Long values
- `stringValue(String val)`: Strings.
- `stringArrayValue(String[] val)`: Arrays of String values.
- `booleanValue(boolean val)`: Boolean values.
- `shortValue(short val)`: Short values.
- `charValue(char val)`: Char values.
- `floatValue(float val)`: Float values.
- `doubleValue(double val)`: Double values.
- `intArrayValue(int[] val)`: Arrays of int values.

The name of the name-value pair is defined in the `dataName` member variable in the `AdditionalData` object. See [Listing 15-1](#)

Listing 15-1 Defining AdditionalData

```
AdditionalData adArray[] = new AdditionalData[1];
AdditionalDataValue targetAddressValue = new AdditionalDataValue();
AdditionalData adTargetAddressString = new AdditionalData();
targetAddressValue.stringValue(address);
adTargetAddressString.dataName = "targetAddress";
adTargetAddressString.dataValue = targetAddressValue;
adArray[0] = adTargetAddressString;
policyRequest.additionalParameters = adArray;
```

If any of the incoming parameters from the application are complex types, the objects are automatically examined and broken down into simple Java types. So, for example, the Parlay X 2.1 complex type `ChargingInformation` can contain a description, which is a string, a currency kind, which is also a string, an amount, which is a decimal number, and a code, which is a string. When the data is sent to the Policy Engine, it is broken down into a string value called

`parameters.charging.currency`, another string value called `parameters.charging.code`, and so forth.

Adding a New Rule

New rules can be added to the Policy Service. The rule must have a name and a priority.

High priority rules are evaluated before low priority rules. There are a set of pre-defined priority levels, which are mapped to a numerical value:

- minimum, where the value is -1×10^9
- low, where the value is -1×10^6
- high, where the value is 1×10^6
- maximum, where the value is 1×10^9

[Listing 15-2](#) shows the basic structure of a rule:

Listing 15-2 Skeleton of a rule

```
rule DenySubscriberNotExists
{
  priority = high;
  when
  {
    // fetch the policy request data and perform evaluations.
  }
  then
  {
    // Take action on
  }
};
```

Mapping PolicyRequest Data

In order to perform an evaluation, the data in the `PolicyRequest` object must be fetched by the rule in the Policy Engine and mapped to the equivalent variable name in the rule. The standard types of request data in the Policy Request are associated with variables of the same name in the rules. Below is an example of a rule assigning the `PolicyRequest` member variable `serviceName` to the rule variable `sname` via the Policy Request object. The rule object `pr` is assigned to the `PolicyRequest` object.

Listing 15-3 Policy Request data is fetched

```
?pr: event PolicyRequest(?sname: serviceName);
```

If the Policy Engine has evaluated the request and made the decision to deny it, the Policy Engine's representation of the `PolicyRequest` object (`pr`) must be *retracted*. Retracting the `PolicyRequest` object aborts further rule enforcement.

Listing 15-4 Retract a request

```
retract (?pr);
```

If the Policy Engine has evaluated the request and made the decision to allow it, the Policy Engine's representation of the request (`pr`) must still be retracted, but in the last rule of the execution flow. For example, this could be achieved by adding a general finalizing allow rule that retracts the request. This rule should have priority `minimum`.

Listing 15-5 General finalizing allow rule that retracts a request

```
rule AllowServiceRequest
{
    priority = minimum;
    when
```

```
{
    ?pr: event PolicyRequest();

}
then
{
    retract (?pr);
    ?pr.allow();

}
};
```

Data that is defined as `AdditionalValues` must be fetched as shown in [Listing 15-6](#). The Additional Value named `targetAddress` is stored in the variable `addDataValue`. The `PolicyRequest` object is `pr`.

Listing 15-6 Fetching `AdditionalValue` data

```
bind ?addDataValue = ?pr.getAdditionalDataStringValue("targetAddress");
```

The particular signature of the fetching method depends on the type of data:

- `getAdditionalDataIntValue(...)`, for int values
- `getAdditionalDataLongValue(...)`, for long value.
- `getAdditionalDataStringValue(...)`, for String values
- `getAdditionalDataStringArrayValue(...)`, for arrays of String values
- `getAdditionalDataBooleanValue(...)`, for boolean values
- `getAdditionalDataShortValue(...)`, for short values

- `getAdditionalDataCharValue(...)`, for char values
- `getAdditionalDataFloatValue(...)`, for float values
- `getAdditionalDataDoubleValue(...)`, for double values
- `getAdditionalDataIntArrayValue(...)` for arrays of int values.

If the data type is unknown, it can be determined by invoking the `discriminator` method on the `AdditionalDataValue` object.

Listing 15-7 Determine the type of an AdditionalDatavalue

```
bind ?type = ?pr.getAdditionalData.dataValue.discriminator().value();
```

Where `type` is one of the following:

- `AdditionalDataType._P_ADDITIONAL_INT`
- `AdditionalDataType._P_ADDITIONAL_LONG`
- `AdditionalDataType._P_ADDITIONAL_STRING`
- `AdditionalDataType._P_ADDITIONAL_STRING_ARRAY`
- `AdditionalDataType._P_ADDITIONAL_BOOLEAN`
- `AdditionalDataType._P_ADDITIONAL_SHORT`
- `AdditionalDataType._P_ADDITIONAL_CHAR`
- `AdditionalDataType._P_ADDITIONAL_FLOAT`
- `AdditionalDataType._P_ADDITIONAL_DOUBLE`
- `AdditionalDataType._P_ADDITIONAL_INT_ARRAY`

Creating a New Rule File by Extending an Existing File: an Example

The following shows an example of extending an existing rule file:

1. List the Current Services' Rule Files

2. Select the Service Whose Rule File You Wish to Extend
3. Add a New Extended Rule
4. Load the New Rule File.

Use the operations in the PolicyService to manage the rule files, see [Managing the PolicyService](#) in the *System Administrator's Guide*.

Using RequestContext Parameters Defined in Service Level Agreements

It is possible to use generic data specified in service provider and application-level SLAs in a plug-in. This is useful when the choice of the action or behavior a plug-in should make is based on which service provider or application originates the request originates. For example, this can be used for information about parameters that corresponds to a certain group of applications. For instance a certain group might get the priority on their SMS set to LOW because they pay less. The priority might be a parameter that is sent down to the network which handles this.

In an SLA, a `<contextAttribute>` is defined as a name/value pair, where the name is defined in the tag `<attributeName>` and the value is specified in `<attributeValue>`.

A plug-in can retrieve the value specified in `<attributeValue>` using the name specified in `<attributeName>`. The value is retrieved using the RequestContext for the request:

```
String attributeValue =
    (String)RequestContextManager.getCurrent().get("<attributeName>");
```

For example, the value associated with the contextAttribute with the attributeName `com.bea.wlcp.wlng.plugin.sms.testName1` is retrieved using:

```
String value1 =
    (String)RequestContextManager.getCurrent().get("com.bea.wlcp.wlng.plugin.s
ms.testName1");
```

Callable Policy Web Service

The following section describes how to use the callable policy interface exposed by Network Gatekeeper.

- [Introduction](#)
- [Callable Policy Web Service interface definition](#)
 - [Endpoints](#)
 - [Detailed service description](#)
 - [XML Schema data type definition](#)
 - [AdditionalDataValue](#) structure
 - [AdditionalDataValueType](#) enumeration
 - [Interface: Policy](#)
 - [Operation: evaluate](#)
 - [Interface: PolicyManagement](#)
 - [Operation: viewRuleFile](#)
 - [Operation: deleteRuleFile](#)
 - [Operation: loadRules](#)
 - [Operation: listRuleFiles](#)
- [Rule files](#)

Introduction

The callable policy service in WebLogic Network Gatekeeper exposes two Web Services interfaces related to callable policy:

- Policy evaluation
- Policy management

The callable policy service is intended to allow applications and network nodes that have no policy evaluation capabilities themselves to use the policy evaluation capabilities in Network Gatekeeper. The service is not designed to expose the service to external service providers. Rather it is to be used internally as a way of exposing generic policy capabilities to network nodes within the telecom network where Network Gatekeeper is deployed. Communication Services deployed in Network Gatekeeper do not use the interfaces exposed by the callable policy Web Service.

For example, a node in the network might need to enforce a set of rules for requests flowing through it, to allow or deny requests based on time of day and originator of the request. In this case, the node might determine the originator of the request and use the callable policy evaluation Web Service to evaluate that request. The rule that is being evaluated uses the data provided in the web services call and makes its decision based on them. Modifications to the rules can be done using the policy management Web Service.

A user of the policy evaluation and policy management Web Services interfaces is registered using the same service provider and application model that is used for users of the Communication Services. If the system requires sessions, the user must be logged in using the same session manager interface exposed to these service provider applications.

Note: If there is no specific rule file associated with a ServiceName loaded in the rule engine, it uses the default rule file in its evaluation. If you are using Callable Policy, you must make sure that an appropriate rule file is loaded into the rule engine. For more information, see the “[Managing the PolicyService](#)” chapter in the *System Administration Guide*.

It necessary to have service provider group and application group Service Level Agreements defined for the user of the callable policy service. To use the policy evaluation interface, the tag `<scs>` must contain the value `com.bea.wlcp.wlmg.px21.plugin.PolicyPlugin`.

To use the policy management interface, the tag `<scs>` must contain the value `com.bea.wlcp.wlmg.px21.plugin.PolicyManagementPlugin`.

Listing 16-1 Example of SLA that allows the use of both the policy evaluation and policy management interfaces

```

<serviceContract>
    <scs>com.bea.wlcp.wlng.px21.plugin.PolicyPlugin</scs>
</serviceContract>
<serviceContract>
    <scs>com.bea.wlcp.wlng.px21.plugin.PolicyManagementPlugin</scs>
</serviceContract>

```

Callable Policy Web Service interface definition

Endpoints

The endpoint for the Policy evaluation interface is:

`http://<host:port>/callable_policy/Policy`

The endpoint for the Policy management interface is:

`http://<host:port>/callable_policy/PolicyManagement`

Detailed service description

Policy Evaluation

The policy evaluation interface makes it possible for an external application to evaluate a request containing a set of parameters. The parameters in the request include authentication information, information on the type of service the request should be evaluated against, the method name of the method that should be evaluated, and arbitrary additional data provided as name-value pairs.

All request parameters are evaluated according to a policy rule.

When evaluated, a copy of the data provided in the evaluation process is returned together with information on the outcome of the requests, that is, if the request was allowed or denied. If the request was allowed, the application calling the Web Service must use the returned copy of the parameters for further processing, because the returned parameters in the request may have been changed by the policy rule processing.

Policy management

The policy management web service interface makes it possible to load and delete policy rules.

XML Schema data type definition

AdditionalDataValue structure

Defines the AdditionalDataValue structure.

Element Name	Element type	Optional	Description
name	xsd:string	N	Name part of the additional data name-value pair.
value	xsd:string	N	Value part of the additional data name-value pair.
type	callable_policy_local_xsd:AdditionalDataValueType	N	Identifies the data type. See AdditionalDataValueType enumeration .

AdditionalDataValueType enumeration

Describes a data type.

Enumeration value	Description
STRING_TYPE	Data type is String.
INTEGER_TYPE	Data type is Integer.
FLOAT_TYPE	Data type is float.
DOUBLE_TYPE	Data type is double.
CHAR_TYPE	Data type is Char.
BOOLEAN_TYPE	Data type is boolean.

Enumeration value	Description
INT_ARRAY_TYPE	Data type is int array.
STRINGARRAY_TYPE	Data type is String array.

Interface: Policy

Operations to evaluate a request.

Operation: evaluate

The policy evaluation interface makes it possible for an external application to evaluate a request containing a set of parameters. All of the request parameters are evaluated according to a Policy rule.

Input message: evaluateRequest

Part name	Part type	Optional	Description
type	xsd:string	N	Service type to be evaluated.
serviceName	xsd:string	N	ServiceName associated with the rule file.
methodName	xsd:string	N	Name of method to be evaluated.
requesterID	xsd:string	N	The application ID as given by the operator.
timeStamp	xsd:dateTime	N	Defines the date and time of the request.
additionalData	callable_policy_local_xsd:additionalDataValue	Y	Specifies any other data, specified as name-value pairs. See AdditionalDataValue structure .

Output message: evaluateResponse

Part name	Part type	Optional	Description
modifiedRequest	callable_policy_local_xsd:evaluateRequest	N	The response that Network Gatekeeper returns after being evaluated by policy rules. Same data structure as evaluateRequest, but data may have been changed by the policy evaluation.
returnValue	xsd:string	N	Return value the policy rules passed back.
thrownException	xsd:string	N	Name of the exception thrown during evaluation.
thrownPolicyException	xsd:string	N	Name of the policy rejection exception thrown during evaluation.
denyReasonDescription	xsd:string	N	Description of the reason of denying the request.
denyCode	xsd:string	N	Code identifying the reason of denying the request.

Referenced faults

ServiceException:

If there is an internal error during evaluation process, a ServiceException is thrown.

PolicyException:

If the policy evaluation request is rejected, a PolicyException is thrown.

Interface: PolicyManagement

Operations to manage policy rules.

Operation: viewRuleFile

Fetches a policy rule file of a given type and service from the rules engine.

Input message: viewRuleFile

Part name	Part type	Optional	Description
type	xsd:string	N	Type of SLA, either: <ul style="list-style-type: none">• Application• Serviceprovider
serviceName	xsd:String	N	ServiceName associated with the rule file.

Output message: viewRuleFileResponse

Part name	Part type	Optional	Description
return	xsd:String	N	The rule file.

Referenced faults

ServiceException:

If there is an internal error during evaluation process, a ServiceException is thrown.

PolicyException:

If the policy evaluation request is rejected, a PolicyException is thrown.

Operation: deleteRuleFile

Deletes a policy rule file of a given type and service from the rules engine.

Input message: deleteRuleFile

Part name	Part type	Optional	Description
type	xsd:string	N	Type of rule file, either: <ul style="list-style-type: none">• Application• Serviceprovider
serviceName	xsd:String	N	ServiceName associated with the rule file.

Output message: deleteRuleFileResponse

Part name	Part type	Optional	Description
-	-	-	-

Referenced faults

ServiceException:

If there is an internal error during evaluation process, a ServiceException is thrown.

PolicyException:

If the policy evaluation request is rejected, a PolicyException is thrown.

Operation: loadRules

Loads a policy rule file of a given type and service into the rules engine.

Input message: loadRules

Part name	Part type	Optional	Description
type	xsd:string	N	Type of rule file, either: <ul style="list-style-type: none">• Application• Serviceprovider
urlUrl	xsd:string	N	URL to rule file to be loaded.
serviceName	xsd:string	N	ServiceName associated with the rule file.

Output message: loadRulesResponse

Part name	Part type	Optional	Description
-	-	-	-

Referenced faults

ServiceException:

If there is an internal error during evaluation process, a ServiceException is thrown.

PolicyException:

If the policy evaluation request is rejected, a PolicyException is thrown.

Operation: listRuleFiles

Lists the rule files of a given type that are loaded into the rules engine.

Input message: listRuleFiles

Part name	Part type	Optional	Description
type	xsd:string	N	Type of rule file, either: <ul style="list-style-type: none">• Application• Serviceprovider

Output message: listRuleFilesResponse

Part name	Part type	Optional	Description
ruleFile	Array of xsd:string	Y	A list of rule files matching the given criteria.

Referenced faults

ServiceException:

If there is an internal error during evaluation process, a ServiceException is thrown.

PolicyException:

If the policy evaluation request is rejected, a PolicyException is thrown.

Rule files

The rule files are written in IRL, ILog Rule Language.

When writing rules in the context of Network Gatekeeper policy rules, the following apply:

The rule is associated with a service name when loaded into Network Gatekeeper policy service, [Input message: loadRules](#).

Which rule to be triggered by [Input message: evaluateRequest](#) is correlated with the parameter `serviceName` given in the Web Service request.

When the evaluate request triggers the rule, a set of general parameters can be accessed by the policy rule:

- String `applicationID`: Application ID associated with the request.
- String `serviceProviderID`: Service provider ID associated with the request.
- String `serviceName`: Service name from which the request originates or is destined for.
- String `methodName`: Method that triggered the request.
- String `serviceCode`.
- String `requesterID`.
- long `transactionID`.
- int `noOfActiveSessions`.
- long `timeStamp`: Time the request was sent to the rules engine for processing. Milliseconds from start of UNIX epoch.
- long `reqCounter`: Defines the increase rate for related counters.

A rule must have a name and a priority. High priority rules are evaluated before low priority rules. There are a set of pre-defined priority levels, which are mapped to a numerical value:

- minimum, where the value is -1×10^9
- low, where the value is -1×10^6
- high, where the value is 1×10^6
- maximum, where the value is 1×10^9

Listing 16-2 shows the basic structure of a rule:

Listing 16-2 Skeleton of a rule

```
rule DenySubscriberNotExists
{
    priority = high;
    when
    {
        // fetch the policy request data and perform evaluations.
```

```
    }  
    then  
    {  
        // Take action on  
    }  
};
```

In order to perform an evaluation, the data in the `PolicyRequest` object must be fetched by the rule and mapped to the equivalent variable names in the rule. The standard types of request data in the Policy Request are associated with variables of the same name in the rules. Below is an example of a rule assigning the `PolicyRequest` member variable `serviceName` to the rule variable `sname` via the Policy Request object. The rule object `pr` is assigned to the `PolicyRequest` object.

Listing 16-3 Policy Request data is fetched

```
?pr: event PolicyRequest(?sname: serviceName);
```

If the Policy Engine has evaluated the request and made the decision to deny it, the Policy Engine's representation of the `PolicyRequest` object (`pr`) must be retracted. Retracting the `PolicyRequest` object aborts further rule enforcement.

Listing 16-4 Retract a request

```
retract (?pr);
```

If the Policy Engine has evaluated the request and made the decision to allow it, the Policy Engine's representation of the request (`pr`) must still be retracted, but in the last rule of the execution flow. For example, this could be achieved by adding a general finalizing allow rule that retracts the request. This rule should have priority minimum.

Listing 16-5 General finalizing allow rule that retracts a request

```
rule AllowServiceRequest
{
    priority = minimum;
    when
    {
        ?pr: event PolicyRequest();
    }
    then
    {
        retract (?pr);
        ?pr.allow();
    }
}
```

Data that is defined as `AdditionalValues` must be fetched as shown below. The `Additional Value` named `targetAddress` is stored in the variable `addDataValue`. The `PolicyRequest` object is `pr`.

Listing 16-6 Fetching `AdditionalValue` data

```
bind ?addDataValue = ?pr.getAdditionalDataStringValue("targetAddress");
```

The particular signature of the fetching method depends on the type of data:

- `getAdditionalDataIntValue(...)`, for int values
- `getAdditionalDataLongValue(...)`, for long value.

- `getAdditionalDataStringValue(...)`, for String values
- `getAdditionalDataStringArrayValue(...)`, for arrays of String values
- `getAdditionalDataBooleanValue(...)`, for boolean values
- `getAdditionalDataShortValue(...)`, for short values
- `getAdditionalDataCharValue(...)`, for char values
- `getAdditionalDataFloatValue(...)`, for float values
- `getAdditionalDataDoubleValue(...)`, for double values
- `getAdditionalDataIntArrayValue(...)` for arrays of int values.

If the data type is unknown, it can be determined by invoking the discriminator method on the `AdditionalDataValue` object.

Listing 16-7 Determine the type of an AdditionalDatavalue

```
bind ?type = ?pr.getAdditionalData.dataValue.discriminator().value();
```

Where `type` is one of the following:

- `AdditionalDataType._P_ADDITIONAL_INT`
- `AdditionalDataType._P_ADDITIONAL_LONG`
- `AdditionalDataType._P_ADDITIONAL_STRING`
- `AdditionalDataType._P_ADDITIONAL_STRING_ARRAY`
- `AdditionalDataType._P_ADDITIONAL_BOOLEAN`
- `AdditionalDataType._P_ADDITIONAL_SHORT`
- `AdditionalDataType._P_ADDITIONAL_CHAR`
- `AdditionalDataType._P_ADDITIONAL_FLOAT`
- `AdditionalDataType._P_ADDITIONAL_DOUBLE`
- `AdditionalDataType._P_ADDITIONAL_INT_ARRAY`

Checklist

This section contains a short summary checklist to use when creating extensions to Network Gatekeeper:

- When creating the management interface, consider if the management operations and attributes should be cluster-wide or local.
- Make sure to follow the plug-in naming convention: `Plugin_<web service interface part>_<network protocol>`.
- Make sure to implement `customMatch` of the `PluginInstance` (or `ManagedPluginInstance`) to be sure that requests end up in the correct plug-in. This is important when there are multiple plug-ins for the same communication service.
- Create exception types that are very specific to various error scenarios. This will allow fine grain control of the alarms that are generated.
- Have a clean separation between the north and the south side of the plug-in.
- Make sure to return all north interfaces (callback included) and souths interfaces when implementing the `getNorthInterfaces()` and `getSouthInterfaces()` of `PluginInstance`.
- Make sure to implement the `resolveAppInstanceGroupId()` method for each `PluginSouth` instance (if applicable).
- Annotate each parameter in the south object methods that you need to have when aspect calls back the `resolveAppInstanceGroupId()` or the `prepareRequestContext()` methods.

Checklist

- Consider what additional EDR fields you need to add. Annotate all the methods you want to be woven using the `@Edr` annotation.
- Annotate the specific arguments you want to see in the EDR for each annotated methods. Use either `@ContextKey` or `@ContextTranslate` depending on the kind of argument.
- Add all the EDRs you are triggering to the EDR descriptor.