BEA WebLogic SIP Server™

Developing Applications with WebLogic SIP Server

Version 3.1
Revised: July 16, 2007

Copyright

Copyright © 1995-2007 BEA Systems, Inc. All Rights Reserved.

Restricted Rights Legend

This software is protected by copyright, and may be protected by patent laws. No copying or other use of this software is permitted unless you have entered into a license agreement with BEA authorizing such use. This document is protected by copyright and may not be copied photocopied, reproduced, translated, or reduced to any electronic medium or machine readable form, in whole or in part, without prior consent, in writing, from BEA Systems, Inc.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE DOCUMENTATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA SYSTEMS DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE DOCUMENT IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

Trademarks and Service Marks

Copyright © 1995-2006 BEA Systems, Inc. All Rights Reserved. BEA, BEA JRockit, BEA WebLogic Portal, BEA WebLogic Server, BEA WebLogic Workshop, Built on BEA, Jolt, JoltBeans, SteelThread, Top End, Tuxedo, and WebLogic are registered trademarks of BEA Systems, Inc. BEA AquaLogic, BEA AquaLogic Data Services Platform, BEA AquaLogic Enterprise Security, BEA AquaLogic Interaction, BEA AquaLogic Interaction Analytics, BEA AquaLogic Interaction Collaboration, BEA AquaLogic Interaction Content Services, BEA AquaLogic Interaction Data Services, BEA AquaLogic Interaction Integration Services, BEA AquaLogic Interaction Process, BEA AquaLogic Interaction Publisher, BEA AquaLogic Interaction Studio, BEA AquaLogic Service Bus, BEA AquaLogic Service Registry, BEA Builder, BEA Campaign Manager for WebLogic, BEA eLink, BEA Kodo, BEA Liquid Data for WebLogic, BEA Manager, BEA MessageQ, BEA SALT, BEA Service Architecture Leveraging Tuxedo, BEA WebLogic Commerce Server, BEA WebLogic Communications Platform, BEA WebLogic Enterprise, BEA WebLogic Enterprise Platform, BEA WebLogic Enterprise Security, BEA WebLogic Express, BEA WebLogic Integration, BEA WebLogic Java Adapter for Mainframe, BEA WebLogic JDriver, BEA WebLogic Log Central, BEA WebLogic Mobility Server, BEA WebLogic Network Gatekeeper, BEA WebLogic Personalization Server, BEA WebLogic Personal Messaging API, BEA WebLogic Platform, BEA WebLogic Portlets for Groupware Integration, BEA WebLogic Real Time, BEA WebLogic RFID Compliance Express, BEA WebLogic RFID Edge Server, BEA WebLogic RFID Enterprise Server, BEA WebLogic Server Process Edition, BEA WebLogic SIP Server, BEA WebLogic WorkGroup Edition, BEA Workshop for WebLogic Platform, BEA Workshop JSP, BEA Workshop JSP Editor, BEA Workshop Struts, BEA Workshop Studio, Dev2Dev, Liquid Computing, and Think Liquid are trademarks of BEA Systems, Inc. Accelerated Knowledge Transfer, AKT, BEA Mission Critical Support, BEA Mission Critical Support Continuum, and BEA SOA Self Assessment are service marks of BEA Systems, Inc.

All other names and marks are property of their respective owners.

Contents

1. Overview of SIP Servlets

What is a SIP Servlet? 1-1
Differences from HTTP Servlets 1-3
Multiple Responses 1-3
Receiving Responses 1-4
Proxy Functions 1-6
Message Body 1-6
ServletRequest 1-7
ServletResponse 1-7
SipServletMessage 1-8
Roles of a Servlet Container 1-8
Application Management 1-8
SIP Messaging 1-10
Utility Functions 1-13

2. Requirements and Best Practices for WebLogic SIP Server Applications

Overview of Developing and Porting Applications for WebLogic SIP Server 2-1
Applications Must Not Create Threads 2-2
Servlets Must Be Non-Blocking 2-3
Store all Application Data in the Session 2-3
All Session Data Must Be Serializable 2-3

. . . . Use setAttribute() to Modify Session Data in "No-Call" Scope2-3
send() Calls Are Buffered 2-6
Mark SIP Servlets as Distributable 2-6
Observe Best Practices for J2EE Applications 2-6
. . . . 2-6

3. Using Compact and Long Header Formats for SIP Messages

Overview of Header Format APIs and Configuration 3-1
Summary of Compact Headers 3-2
Assigning Header Formats with WlssSipServletMessage 3-2
Summary of API and Configuration Behavior 3-3

4. Composing SIP Applications

Overview of SIP Application Composition 4-1
Application Composition Model 4-2
Sample Composer Application 4-3
Troubleshooting Application Composition 4-5
Deployment Order Requirement 4-6

5. Developing Converged Applications

Overview of Converged Applications 5-1
Assembling and Packaging a Converged Application 5-2
Working with SIP and HTTP Sessions 5-2
Modifying the SipApplicationSession 5-4
Using the Converged Application Example 5-5

6. Using the Diameter Base Protocol API

Overview of Diameter Protocol Support 6-1
Overview of the Diameter API 6-2

Working with Diameter Nodes 6-4
Implementing a Diameter Application 6-5
Working with Diameter Sessions 6-6
Working with Diameter Messages. 6-7
Sending Request Messages. 6-8
Sending Answer Messages 6-8
Creating New Command Codes 6-8
Working with AVPs 6-9
Creating New Attributes. 6-9
Creating Converged Diameter and SIP Applications 6-10

7. Using the Diameter Sh Interface Application

Overview of Profile Service API and Sh Interface Support 7-1
Enabling the Sh Interface Provider 7-2
Overview of the Profile Service API. 7-2
Creating a Document Key for Application-Managed Profile Data 7-3
Using a Constructed Document Key to Manage Profile Data. 7-5
Monitoring Profile Data with ProfileListener 7-6
Prerequisites for Listener Implementations 7-7
Implementing ProfileListener. 7-7

8. Using the Diameter Rf Interface Application for Offline Charging

Overview of Rf Interface Support 8-1
Understanding Offline Charging Events 8-2
Event-Based Charging ¶. 8-2
Session-Based Charging ¶ 8-3
Configuring the Rf Application. 8-4

Using the Offline Charging API 8-5
Accessing the Rf Application 8-6
Implementing Session-Based Charging 8-6
Sending Asynchronous Requests 8-8
Specifying the Session Expiration 8-9
Implementing Event-Based Charging 8-9
Using the Accounting Session State ¶ 8-10

9. Using the Diameter Ro Interface Application for Online Charging

Overview of Ro Interface Support 9-1
Understanding Credit Authorization Models 9-2
Credit Authorization with Unit Determination 9-2
Credit Authorization with Direct Debiting 9-3
Determining Units and Rating 9-3
Configuring the Ro Application 9-3
Overview of the Online Charging API 9-4
Accessing the Ro Application 9-5
Implementing Session-Based Charging 9-6
Handling Re-Auth-Request Messages 9-6
Sending Credit-Control-Request Messages 9-8
Handling Failures 9-9

10. Developing Custom Profile Providers

Overview of the Profile Service API. 10-1
Implementing Profile API Methods 10-3
Configuring and Packaging Profile Providers. 10-3
Mapping Profile Requests to Profile Providers 10-5

Configuring Profile Providers Using the Administration Console 10-6

11.Using Content Indirection in SIP Servlets

Overview of Content Indirection 11-1

Using the Content Indirection API 11-2

Additional Information 11-2

12.Securing SIP Servlet Resources

Overview of SIP Servlet Security 12-1

WebLogic SIP Server Role Mapping Features 12-2

Using Implicit Role Assignment 12-3

Assigning Roles Using security-role-assignment 12-4

Important Requirement for WebLogic SIP Server 3.1 12-4

Assigning Roles at Deployment Time 12-6

Dynamically Assigning Roles Using the Administration Console 12-6

Assigning run-as Roles 12-7

Role Assignment Precedence for SIP Servlet Roles 12-8

Debugging Security Features. 12-9

weblogic.xml Deployment Descriptor Reference 12-9

13.Developing SIP Servlets Using Eclipse

Overview 13-1

SIP Servlet Organization 13-2

Setting Up the Development Environment 13-2

Creating a WebLogic SIP Server Domain 13-3

Configure the Default Eclipse JVM 13-3

Creating a New Eclipse Project 13-3

Creating an Ant Build File 13-4

Building and Deploying the Project 13-6

Debugging SIP Servlets 13-6

14. Enabling Message Logging

Overview 14-1
Enabling Message Logging 14-2
Specifying a Predefined Logging Level. 14-2
Customizing Log Records 14-3
Specifying Content Types for Unencrypted Logging 14-5
Example Message Log Configuration and Output 14-6
Configuring Log File Rotation 14-8

15. Generating SNMP Traps from Application Code

Overview 15-1
Requirement for Accessing SipServletSnmpTrapRuntimeMBean 15-2
Obtaining a Reference to SipServletSnmpTrapRuntimeMBean. 15-3
Generating a SNMP Trap 15-5

CHAPTER 1

Overview of SIP Servlets

What is a SIP Servlet?

The SIP Servlet API is a part of JAIN APIs and being standardized as JSR116 of JCP (Java Community Process). The SIP Servlet API version 1.0 was published in February, 2003.

Note: In this document, the term "SIP Servlet" is used to represent the API, and "SIP servlet" is used to represent an application created with the API.

J2EE provides Java Servlet that is a main technology of building Web applications. Although Java Servlet is used only to develop HTTP protocol-based applications on a Web application server, it basically has functions as a generic API for server applications. SIP Servlet is defined as the generic servlet API with SIP-specific functions added.

Figure 1-1 Servlet API and SIP Servlet API

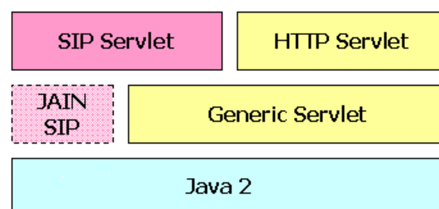

SIP Servlets are very similar to HTTP Servlets, and HTTP servlet developers will quickly adapt to the programming model. The service level defined by both HTTP and SIP Servlets is very similar, and you can easily design applications that support both HTTP and SIP. Listing 1 shows an example of a simple SIP servlet.

Listing 1-1 List 1: SimpleSIPServlet.java

```
package com.bea.example.simple;
import java.io.IOException;
import javax.servlet.*;
import javax.servlet.sip.*;

public class SimpleSIPServlet extends SipServlet {

    protected void doMessage(SipServletRequest req)
       throws ServletException, IOException
    {
       SipServletResponse res = req.createResponse(200);
       res.send();
    }
}
```

The above example shows a simple SIP servlet that sends back a 200 OK response to the SIP MESSAGE request. As you can see from the list, SIP Servlet and HTTP Servlet have many things in common:

1. Servlets must inherit the base class provided by the API. HTTP servlets must inherit HttpServlet, and SIP servlets must inherit SipServlet.
2. Methods doXxx must be overridden and implemented. HTTP servlets have doGet/doPost methods corresponding to GET/POST methods. Similarly, SIP servlets have doXxx methods corresponding to the method name (in the above example, the MESSAGE method). Application developers override and implement necessary methods.
3. The lifecycle and management method (init, destroy) of SIP Servlet are exactly the same as HTTP Servlet. Manipulation of sessions and attributes is also the same.
4. Although not appeared in the API, there is a deployment descriptor called sip.xml for a SIP servlet, which corresponds to web.xml. Application developers and service managers can edit this file to configure applications using multiple SIP servlets.

However, there are several differences between SIP and HTTP servlets. A major difference comes from protocols. The next section describes these differences as well as features of SIP servlets.

Differences from HTTP Servlets

Multiple Responses

You might notice from the List 1 that the doMessage method has only one argument. In HTTP, a transaction consists of a pair of request and response, so arguments of a doXxx method specify a request (HttpServletRequest) and its response (HttpServletResponse). An application takes information such as parameters from the request to execute it, and returns its result in the body of the response.

```
protected void doGet(HttpServletRequest req, HttpServletResponse res)
   throws ServletException, IOException
```

For SIP, more than one response may be returned to a single request.

Figure 1-2 Example of Request and Response in SIP

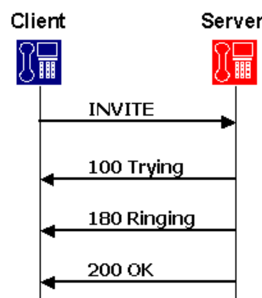

The above figure shows an example of a response to the INVITE request. In this example, the server sends back three responses 100, 180, and 200 to the single INVITE request. To implement such sequence, in SIP Servlet, only a request is specified in a doXxx method, and an application generates and returns necessary responses in an overridden method.

Currently, SIP Servlet defines the following doXxx methods:

```
protected void doInvite(SipServletRequest req);
```

```
protected void doAck(SipServletRequest req);
protected void doOptions(SipServletRequest req);
protected void doBye(SipServletRequest req);
protected void doCancel(SipServletRequest req);
protected void doSubscribe(SipServletRequest req);
protected void doNotify(SipServletRequest req);
protected void doMessage(SipServletRequest req);
protected void doInfo(SipServletRequest req);
protected void doPrack(SipServletRequest req);
```

Receiving Responses

One of the major features of SIP is that roles of a client and server are not fixed. In HTTP, Web browsers always send HTTP requests and receive HTTP responses: They never receive HTTP requests and send HTTP responses. In SIP, however, each terminal needs to have functions of both a client and server.

For example, both of two SIP phones must call to the other and disconnect the call.

Figure 1-3 Relationship between Client and Server in SIP

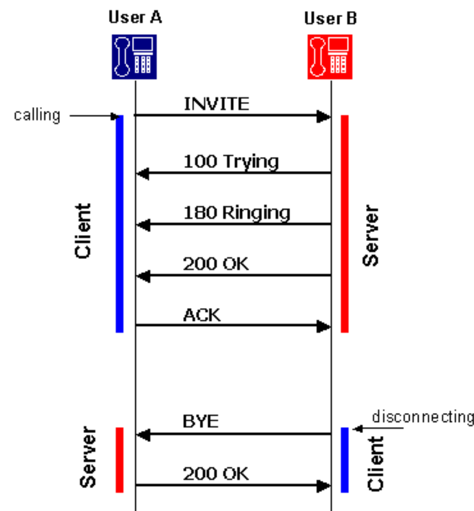

The above example indicates that a calling or disconnecting terminal acts as a client. In SIP, roles of a client and server can be changed in one dialog. This client function is called UAC (User Agent Client) and server function is called UAS (User Agent Server), and the terminal is called UA (User Agent). SIP Servlet defines methods to receive responses as well as requests.

```
protected void doProvisionalResponse(SipServletResponse res);

protected void doSuccessResponse(SipServletResponse res);

protected void doRedirectResponse(SipServletResponse res);

protected void doErrorResponse(SipServletResponse res);
```

These doXxx response methods are not the method name of the request. They are named by the type of the response as follows:

- doProvisionalResponse—A method invoked on the receipt of a provisional response (or 1xx response).
- doSuccessResponse—A method invoked on the receipt of a success response.
- doRedirectResponse—A method invoked on the receipt of a redirect response.

- doErrorResponse—A method invoked on the receipt of an error response (or 4xx, 5xx, 6xx responses).

Existence of methods to receive responses indicates that in SIP Servlet requests and responses are independently transmitted an application in different threads. Applications must explicitly manage association of SIP messages. An independent request and response makes the process slightly complicated, but enables you to write more flexible processes.

Also, SIP Servlet allows applications to explicitly create requests. Using these functions, SIP servlets can not only wait for requests as a server (UAS), but also send requests as a client (UAC).

Proxy Functions

Another function that is different from the HTTP protocol is "forking." Forking is a process of proxying one request to multiple servers simultaneously (or sequentially) and used when multiple terminals (operators) are associated with one telephone number (such as in a call center).

Figure 1-4 Proxy Forking

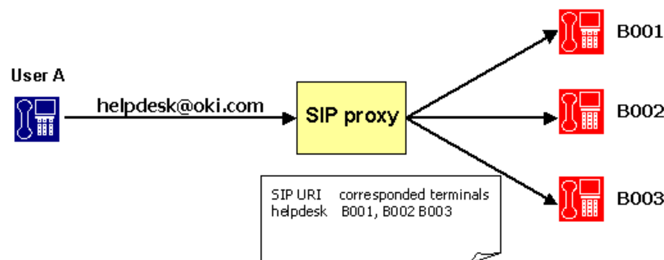

SIP Servlet provides a utility to proxy SIP requests for applications that have proxy functions.

Message Body

As the figure below, the structure of SIP messages is the same as HTTP.

Figure 1-5 SIP Message Example

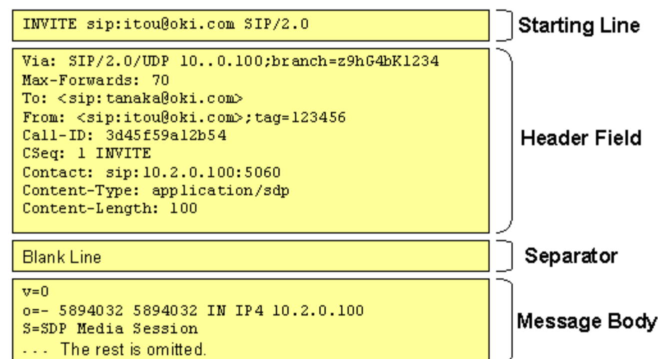

HTTP is basically a protocol to transfer HTML files and images. Contents to be transferred are stored in the message body. HTTP Servlet defines stream manipulation-based API to enable sending and receiving massive contents.

ServletRequest

```
ServletInputStream getInputStream()

BufferedReader     getReader()
```

ServletResponse

```
ServletOutputStream getOutputStream()

PrintWriter         getWriter()

int  getBufferSize()

void setBufferSize(int size)

void resetBuffer()

void flushBuffer()
```

In SIP, however, only low-volume contents are stored in the message body since SIP is intended for real-time communication. Therefore, above methods are provided only for compatibility, and their functions are disabled.

In SIP, contents stored in the body include:

- SDP (Session Description Protocol)—A protocol to define multimedia sessions used between terminals. This protocol is defined in RFC2373.
- Presence Information—A message that describes presence information defined in CPIM.
- IM Messages—IM (instant message) body. User-input messages are stored in the message body.

Since the message body is in a small size, processing it in a streaming way increases overhead. SIP Servlet re-defines API to manipulate the message body on memory as follows:

SipServletMessage

```
void   setContent(Object content, String contentType)
```

```
Object getContent()
```

```
byte[] getRawContent()
```

Roles of a Servlet Container

The following sections describes major functions provided by WebLogic SIP Server as a SIP servlet container:

- Application Management—Describes functions such as application management by servlet context, lifecycle management of servlets, application initialization by deployment descriptors.
- SIP Messaging—Describes functions of parsing incoming SIP messages and delivering appropriate SIP servlets, sending messages created by SIP servlets to appropriate UAS, and automatically setting SIP header fields.
- Utility Functions—Describes functions such as sessions, factories, and proxying that are available in SIP servlets.

Application Management

Like HTTP servlet containers, SIP servlet containers manage applications by servlet context (see Figure 6). Servlet contexts (applications) are normally archived in a WAR format and deployed in each application server.

Note: The method of deploying in application servers varies depending on your product. Refer to the documentation of your application server.

Figure 1-6 Servlet Container and Servlet Context

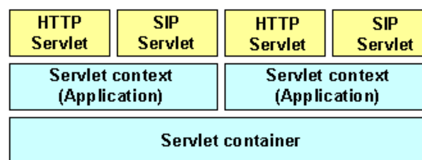

A servlet context for a converged SIP and Web application can include multiple SIP servlets, HTTP servlets, and JSPs.

WebLogic SIP Server can deploy applications using the same method as the application server you use as the platform. However, if you deploy applications including SIP servlets, you need a SIP specific deployment descriptor (sip.xml) defined by SIP servlets. The table below shows the file structure of a general converged SIP and Web application.

Table 1-1 File Structure Example of Application

File	Description
WEB-INF/	Place your configuration and executable files of your converged SIP and Web application in the directory. You cannot directly refer to files in this directory on Web (servlets can do this).
WEB-INF/web.xml	The J2EE standard configuration file for the Web application.
WEB-INF/sip.xml	The SIP Servlet-defined configuration files for the SIP application.
WEB-INF/classes/	Store compiled class files in the directory. You can store both HTTP and SIP servlets in this directory.
WEB-INF/lib/	Store class files archived as Jar files in the directory. You can store both HTTP and SIP servlets in this directory.
*.jsp, *.jpg	Files comprising the Web application (e.g. JSP) can be deployed in the same way as J2EE.

Information specified in the sip.xml file is similar to that in the web.xml except <servlet-mapping> setting that is different from HTTP servlets. In HTTP you specify a servlet associated with the file name portion of URL. But SIP has no concept of the file name. You set

filter conditions using URI or the header field of a SIP request. The following example shows that a SIP servlet called “register” is assigned all REGISTER methods.

Listing 1-2 List 1: Filter Condition Example of sip.xml

```
<servlet-mapping>
  <servlet-name>registrar</servlet-name>
  <pattern>
    <equal>
      <var>request.method</var>
      <value>REGISTER</value>
    </equal>
  </pattern>
</servlet-mapping>
```

Once deployed, lifecycle of the servlet context is maintained by the servlet container. Although the servlet context is normally started and shutdown when the server is started and shutdown, the system administrator can explicitly start, stop, and reload the servlet context.

SIP Messaging

SIP messaging functions provided by a SIP servlet container are classified under the following types:

- Parsing received SIP messages.
- Delivering parsed messages to the appropriate SIP servlet.
- Sending SIP servlet-generated messages to the appropriate UA
- Automatically generating a response (such as “100 Trying”).
- Automatically managing the SIP header field.

All SIP messages that a SIP servlet handles are represented as a SipServletRequest or SipServletResponse object. A received message is first parsed by the parser and then translated to one of these objects and sent to the SIP servlet container.

A SIP servlet container receives the following three types of SIP messages, for each of which you determine a target servlet.

- First SIP Request—When the SIP servlet container received a request that does not belong to any SIP session, it uses filter conditions in the sip.xml file (described in the previous section) to determine the target SIP servlet. Since the container creates a new SIP session when the initial request is delivered, any SIP requests received after that point are considered as subsequent requests.

  Note: Filtering should be done carefully. In WebLogic SIP Server, when the received SIP message matches multiple SIP servlets, it is delivered only to any one SIP servlet.

- Subsequent SIP Request—When the SIP Servlet container receives a request that belongs to any SIP session, it delivers the request to a SIP Servlet associated with that session. Whether the request belongs to a session or not is determined using dialog ID.

  Each time a SIP Servlet processes messages, a lock is established by the container on the call ID. If a SIP Servlet is currently processing earlier requests for the same call ID when subsequent requests are received, the SIP Servlet container queues the subsequent requests. The queued messages are processed only after the Servlet has finished processing the initial message and has returned control to the SIP Servlet container.

  This concurrency control is guaranteed both in a single containers and in clustered environments. Application developers can code applications with the understanding that only one message for any particular call ID will be processed at a given time.

- SIP Response—When the received response is to a request that a SIP servlet proxied, the response is automatically delivered to the same servlet since its SIP session had been determined. When a SIP servlet sends its own request, you must first specify a servlet that receives a response in the SIP session. For example, if the SIP servlet sending a request also receives the response, the following handler setting must be specified in the SIP session.

```
SipServletRequest req = getSipFactory().createRequest(appSession, ...);

req.getSession().setHandler(getServletName());
```

Normally, in SIP a “session” means a real-time session by RTP/RTSP. On the other hand, in HTTP Servlet a “session” refers to a way of relating multiple HTTP transactions. In this document, session-related terms are defined as follows:

Table 1-2 Session-Related Terminology

Realtime Session	A realtime session established by RTP/RTSP.
HTTP Session	A session defined by HTTP Servlet. A means of relating multiple HTTP transactions.
SIP Session	A means of implementing the same concept as in HTTP session in SIP. SIP (RFC3261) has a similar concept of “dialog,” but in this document this is treated as a different term since its lifecycle and generation conditions are different.
Application Session	A means for applications using multiple protocols and dialogs to associate multiple HTTP sessions and SIP sessions. Also called “AP session.”

WebLogic SIP Server automatically execute the following response and retransmission processes:

- Sending “100 Trying”—When WebLogic SIP Server receives an INVITE request, it automatically creates and sends “100 Trying.”
- Response to CANCEL—When WebLogic SIP Server receives a CANCEL request, it executes the following processes if the request is valid.
  a. Sends a 200 response to the CANCEL request.
  b. Sends a 487 response to the INVITE request to be cancelled.
  c. Invokes a doCancel method on the SIP servlet. This allows the application to abort the process within the doCancel method, eliminating the need for explicitly sending back a response.
- Sends ACK to an error response to INVITE—When a 4xx, 5xx, or 6xx response is returned for INVITE that were sent by a SIP servlet, WebLogic SIP Server automatically creates and sends ACK. This is because ACK is required only for a SIP sequence, and the SIP servlet does not require it.

  When the SIP servlet sends a 4xx, 5xx, or 6xx response to INVITE, it never receives ACK for the response.

- Retransmission process when using UDP—SIP defines that sent messages are retransmitted when low-trust transport including UDP is used. WebLogic SIP Server automatically do the retransmission process according to the specification.

Mostly, applications do not need to explicitly set and see header fields In HTTP Servlet since HTTP servlet containers automatically manage these fields such as Content-Length and Content-Type. SIP Servlet also has the same header management function.

In SIP, however, since important information about message delivery exists in some fields, these headers are not allowed to change by applications. Headers that can not be changed by SIP servlets are called “system headers.” The table below lists system headers:

Table 1-3 System Headers

Header Name	Description
Call-ID	Contains ID information to associate multiple SIP messages as Call.
From, To	Contains Information on the sender and receiver of the SIP request (SIP, URI, etc.). tag parameters are given by the servlet container.
CSeq	Contains sequence numbers and method names.
Via	Contains a list of servers the SIP message passed through. This is used when you want to keep track of the pass to send a response to the request.
Record-Route, Route	Used when the proxy server mediates subsequent requests.
Contact	Contains network information (such as IP address and port number) that is used for direct communication between terminals. For a REGISTER message, 3xx, or 485 response, this is not considered as the system header and SIP servlets can directly edit the information.

Utility Functions

SIP Servlet defines the following utilities that are available to SIP servlets:

1. SIP Session, Application Session
2. SIP Factory
3. Proxy

SIP Session, Application Session

As stated before, SIP Servlet provides a "SIP session" whose concept is the same as a HTTP session. In HTTP, multiple transactions are associated using information like Cookie. In SIP, this association is done with header information (Call-ID and tag parameters in From and To). Servlet containers maintain and manage SIP sessions. Messages within the same dialog can refer to the same SIP session. Also, For a method that does not create a dialog (such as MESSAGE), messages can be managed as a session if they have the same header information.

SIP Servlet has a concept of an "application session," which does not exist in HTTP Servlet. An application session is an object to associate and manage multiple SIP sessions and HTTP sessions. It is suitable for applications such as B2BUA.

SIP Factory

A SIP factory (SipFactory) is a factory class to create SIP Servlet-specific objects necessary for application execution. You can generate the following objects:

Table 1-4 Objects Generated with SipFactory

Class Name	Description
URI, SipURI, Address	Can generate address information including SIP URI from String.
SipApplicationSession	Creates a new application session. It is invoked when a SIP servlet starts a new SIP signal process.
SipServletRequest	Used when a SIP servlet acts as UAC to create a request. Such requests can not be sent with Proxy.proxyTo. They must be sent with SipServletRequest.send.

SipFactory is located in the servlet context attribute under the default name. You can take this with the following code.

```
ServletContext context = getServletContext();

SipFactory factory =

    (SipFactory) context.getAttribute("javax.servlet.sip.SipFactory");
```

Proxy

Proxy is a utility used by a SIP servlet to proxy a request. In SIP, proxying has its own sequences including forking. You can specify the following settings in proxying with Proxy:

- Recursive routing (recurse)—When the destination of proxying returns a 3xx response, the request is proxied to the specified target.
- Record-Route setting—Sets a <code>Record-Route</code> header in the specified request.
- Parallel/Sequential (parallel)—Determines whether forking is executed in parallel or sequentially.
- stateful—Determines whether proxying is transaction stateful.
- Supervising mode—In the event of the state change of proxying (response receipts), an application reports this.

CHAPTER 2

Requirements and Best Practices for WebLogic SIP Server Applications

The following sections describe requirements and best practices for developing applications for deployment to WebLogic SIP Server:

- "Overview of Developing and Porting Applications for WebLogic SIP Server" on page 2-1
- "Applications Must Not Create Threads" on page 2-2
- "Servlets Must Be Non-Blocking" on page 2-3
- "Store all Application Data in the Session" on page 2-3
- "All Session Data Must Be Serializable" on page 2-3
- "Use setAttribute() to Modify Session Data in "No-Call" Scope" on page 2-3
- "send() Calls Are Buffered" on page 2-6
- "Mark SIP Servlets as Distributable" on page 2-6
- "Observe Best Practices for J2EE Applications" on page 2-6

Overview of Developing and Porting Applications for WebLogic SIP Server

In a typical production environment, SIP applications are deployed to a cluster of WebLogic SIP Server instances that form the engine tier cluster. A separate cluster of servers in the data tier provides a replicated, in-memory database of the call states for active calls. In order for

applications to function reliably in this environment, you must observe the programming practices and conventions described in the sections that follow to ensure that multiple deployed copies of your application perform as expected in the clustered environment.

If you are porting an application from a previous version of WebLogic SIP Server, many of the conventions and restrictions described below may be new to you, because previous WebLogic SIP Server implementations did not support a clustering. As always, thoroughly test and profile your ported applications to discover problems and ensure adequate performance in the new environment.

Applications Must Not Create Threads

WebLogic SIP Server is a multi-threaded application server that carefully manages resource allocation, concurrency, and thread synchronization for the modules it hosts. To obtain the greatest advantage from the WebLogic SIP Server architecture, construct your application modules according to the SIP Servlet and J2EE API specifications.

Avoid application designs that require creating new threads in server-side modules such as SIP Servlets:

- The SIP Servlet container automatically locks the associated call state when invoking the `do`*`xxx`* method of a SIP Servlet. If the `do`*`xxx`* method spawns additional threads or accesses a different call state before returning control, *deadlock scenarios and lost updates to session data can occur.*
- Applications that create their own threads do not scale well. Threads in the JVM are a limited resource that must be allocated thoughtfully. Your applications may break or cause poor WebLogic SIP Server performance when the server load increases. Problems such as deadlocks and thread starvation may not appear until the application is under a heavy load.
- Multithreaded modules are complex and difficult to debug. Interactions between application-generated threads and WebLogic Server threads are especially difficult to anticipate and analyze.

WARNING: If your application must spawn threads, you must guard against deadlocks and carefully manage concurrent access to session data. At a minimum, never spawn threads inside the service method of a SIP Servlet. Instead, maintain a separate thread pool outside of the service method, and be careful to synchronize access to all session data.

Servlets Must Be Non-Blocking

SIP and HTTP Servlets must not block threads in the body of a SIP method because the call state remains locked while the method is invoked. For example, no Servlet method should actively wait for data to be retrieved or written before returning control to the SIP Servlet container.

Store all Application Data in the Session

If you deploy your application to more than one engine tier server (in a replicated WebLogic SIP Server configuration) you must store all application data in the session as session attributes. In a replicated configuration, engine tier servers maintain no cached information; all application data must be de-serialized from the session attribute available in data tier servers.

All Session Data Must Be Serializable

To support in-memory replication of SIP application call states, you must ensure that all objects stored in the SIP Servlet session are serializable. Every field in an object must be serializable or transient in order for the object to be considered serializable. If the Servlet uses a combination of serializable and non-serializable objects, WebLogic SIP Server cannot replicate the session state of the non-serializable objects.

Use setAttribute() to Modify Session Data in "No-Call" Scope

The SIP Servlet container automatically locks the associated call state when invoking the `do`*xxx* method of a SIP Servlet. However, applications may also attempt to modify session data in "no-call" scope. No-call scope refers to the context where call state data is modified outside the scope of a normal `do`*xxx* method. For example, data is modified in no-call scope when an HTTP Servlet attempts to modify SIP session data, or when a SIP Servlet attempts to modify a call state other than the one that the container locked before invoking the Servlet.

Applications must always use the SIP Session's `setAttribute` method to change attributes in no-call scope. Likewise, use `removeAttribute` to remove an attribute from a session object. Each time `setAttribute/removeAttribute` is used to update session data, the SIP Servlet container obtains and releases a lock on the associated call state. (The methods enqueue the object for updating, and return control immediately.) This ensures that only one application modifies the data at a time, and also ensures that your changes are replicated across data tier nodes in a cluster.

If you use other set methods to change objects within a session, WebLogic SIP Server cannot replicate those changes.

Note that the WebLogic SIP Server container does not persist changes to a call state attribute that are made *after* calling `setAttribute`. For example, in the following code sample the `setAttribute` call immediately modifies the call state, but the subsequent call to `modifyState()` does not:

```
Foo foo = new Foo(..);
appSession.setAttribute("name", foo); // This persists the call state.
foo.modifyState(); // This change is not persisted.
```

Instead, ensure that your Servlet code modifies the call state attribute value *before* calling `setAttribute`, as in:

```
Foo foo = new Foo(..);
foo.modifyState();
appSession.setAttribute("name", foo);
```

Also, keep in mind that the SIP Servlet container obtains a lock to the call state for *each* individual `setAttribute` call. For example, when executing the following code in an HTTP Servlet, the SIP Servlet container obtains and releases a lock on the call state lock twice:

```
appSess.setAttribute("foo1", "bar2");
appSess.setAttribute("foo2", "bar2");
```

This locking behavior ensures that only one thread modifies a call state at any given time. However, another process could potentially modify the call state between sequential updates. The following code is not considered thread safe when done no-call state:

```
Integer oldValue = appSession.getAttribute("counter");
Integer newValue = incrementCounter(oldValue);
appSession.setAttribute("counter", newValue);
```

To make the above code thread safe, you must enclose it using the `wlssAppSession.doAction` method, which ensures that all modifications made to the call state are performed within a single transaction lock, as in:

```
wlssAppSession.doAction(new WlssAction() {
       public Object run() throws Exception {
```

```
        Integer oldValue = appSession.getAttribute("counter");
        Integer newValue = incrementCounter(oldValue);
        appSession.setAttribute("counter", newValue);
        return null;
      }
    });
```

Finally, be careful to avoid deadlock situations when locking call states in a “`do`*`SipMethod`*” call, such as `doInvite()`. Keep in mind that the WebLogic SIP Server container has already locked the call state when the instructions of a `do`*`SipMethod`* are executed. If your application code attempts to access the current call state from within such a method (for example, by accessing a session that is stored within a data structure or attribute), the lock ordering results in a deadlock.

Listing 2-1 shows an example that can result in a deadlock. If the code is executed by the container for a call associated with `callAppSession`, the locking order is reversed and the attempt to obtain the session with `getApplicationSession(callId)` causes a deadlock.

Listing 2-1 Session Access Resulting in a Deadlock

```
WlssSipApplicationSession confAppSession = (WlssSipApplicationSession)
appSession;
confAppSession.doAction(new WlssAction() {
  // confAppSession is locked
  public Object run() throws Exception {
    String callIds = confAppSession.getAttribute("callIds");
    for (each callId in callIds) {
      callAppSess = Session.getApplicationSession(callId);
      // callAppSession is locked
      attributeStr += callAppSess.getAttribute("someattrib");
    }
    confAppSession.setAttribute("attrib", attributeStr);
  }
```

```
}
```

See “Modifying the SipApplicationSession” on page 5-4 for more information about using the `com.bea.wcp.sip.WlssAction` interface.

send() Calls Are Buffered

If your SIP Servlet calls the `send()` method within a SIP request method such as `doInvite()`, `doAck()`, `doNotify()`, and so forth, keep in mind that the WebLogic SIP Server container buffers all `send()` calls and transmits them in order *after* the SIP method returns. Applications cannot rely on `send()` calls to be transmitted immediately as they are called.

WARNING: Applications must not wait or sleep after a call to `send()`, because the request or response is not transmitted until control returns to the SIP Servlet container.

Mark SIP Servlets as Distributable

If you have designed and programmed your SIP Servlet to be deployed to a cluster environment, you must include the `distributable` marker element in the Servlet’s deployment descriptor when deploying the application to a cluster of engine tier servers. If you omit the `distributable` element, WebLogic SIP Server will not deploy the Servlet to a cluster of engine tier servers.

The `distributable` element is not required, and is ignored if you deploy to a single, combined-tier (non-replicated) WebLogic SIP Server instance.

Observe Best Practices for J2EE Applications

If you are deploying applications that use other J2EE APIs, observe the basic clustering guidelines associated with those APIs. For example, if you are deploying EJBs you should design all methods to be idempotent and make EJB homes clusterable in the deployment descriptor. See Clustering Best Practices in the WebLogic Server 9.2 Documentation for more information.

CHAPTER 3

Using Compact and Long Header Formats for SIP Messages

The following sections describe how to use the WebLogic SIP Server `WlssSipServletMessage` interface and configuration parameters to control SIP message header formats:

- "Overview of Header Format APIs and Configuration" on page 3-1
- "Summary of Compact Headers" on page 3-2
- "Assigning Header Formats with WlssSipServletMessage" on page 3-2
- "Summary of API and Configuration Behavior" on page 3-3

Overview of Header Format APIs and Configuration

Applications that operate on wireless networks may want to limit the size of SIP headers to reduce the size of messages and conserve bandwidth. JSR 116 provides the `SipServletMessage.setHeader` method, which enables application developers to set a given header to a specific string value, such as the compact (1-letter) format.

WebLogic SIP Server extends the `SipServletMessage` interface with `WlssSipServletMessage`. One feature of the `WlssSipServletMessage` API is the ability to set long or compact header formats for the entire SIP message using the `setUseHeaderForm` method.

In addition to `WlssSipServletMessage`, WebLogic SIP Server provides a container-wide configuration parameter that can control SIP header formats for all system-generated headers. This system-wide parameter can be used along with

`WlssSipServletMessage.setUseHeaderForm` and `SipServletMessage.setHeader` to further customize header formats.

Summary of Compact Headers

Table 3-1 defines the compact header abbreviations described in the SIP specification (RFC3261). Specifications that introduce additional headers may also include compact header abbreviations.

Table 3-1 Compact Header Abbreviations

Header Name (Long Format)	Compact Format
Call-ID	i
Contact	m
Content-Encoding	e
Content-Length	l
Content-Type	c
From	f
Subject	s
Supported	k
To	t
Via	v

Assigning Header Formats with WlssSipServletMessage

All instances of `SipServletRequest`, `SipServletResponse`, and `WlssSipServletResponse` can be cast to `WlssSipServletMessage` in order to use the extended API. A pair of getter/setter methods, `setUseHeaderForm` and `getUseHeaderForm`, are used to assign or retrieve the header formats used in the message. These methods assign or return a `HeaderForm` object, which is a simple Enumeration that describes the header format:

- `COMPACT`—Forces all headers in the message to use compact format. This behavior is similar to the container-wide configuration value of "force compact," as described in `use-compact-form` in the *Configuration Reference*.

- `LONG`—Forces all headers in the message to use long format. This behavior is similar to the container-wide configuration value of "force long," as described in `use-compact-form` in the *Configuration Reference*.
- `DEFAULT`—Defers the header format to the container-wide configuration value set in `use-compact-form`.

`WlssSipServletResponse.setUseHeaderForm` can be used in combination with `SipServletMessage.setHeader` and the container-level configuration parameter, `use-compact-form`. "Summary of API and Configuration Behavior" on page 3-3.

Summary of API and Configuration Behavior

Header formats can be specified at the header, message, and SIP Servlet container levels. Table 3-2 shows the header format that results when adding a new header with `SipServletMessage.setHeader`, given different container configurations and message-level settings with `WlssSipServletResponse.setUseHeaderForm`.

Table 3-2 API Behavior when Adding Headers

SIP Servlet Container Header Configuration (use-compact-form Setting)	WlssSipServletMessage. setUseHeaderForm Setting	SipServletMessage. setHeader Value	Resulting Header
COMPACT	DEFAULT	"Content-Type"	"Content-Type"
COMPACT	DEFAULT	"c"	"c"
COMPACT	COMPACT	"Content-Type"	"c"
COMPACT	COMPACT	"c"	"c"
COMPACT	LONG	"Content-Type"	"Content-Type"
COMPACT	LONG	"c"	"Content-Type"
LONG	DEFAULT	"Content-Type"	"Content-Type"
LONG	DEFAULT	"c"	"c"
LONG	COMPACT	"Content-Type"	"c"
LONG	COMPACT	"c"	"c"
LONG	LONG	"Content-Type"	"Content-Type"
LONG	LONG	"c"	"Content-Type"

Table 3-2 API Behavior when Adding Headers

FORCE_COMPACT	DEFAULT	“Content-Type”	“c”
FORCE_COMPACT	DEFAULT	“c”	“c”
FORCE_COMPACT	COMPACT	“Content-Type”	“c”
FORCE_COMPACT	COMPACT	“c”	“c”
FORCE_COMPACT	LONG	“Content-Type”	“Content-Type”
FORCE_COMPACT	LONG	“c”	“Content-Type”
FORCE_LONG	DEFAULT	“Content-Type”	“Content-Type”
FORCE_LONG	DEFAULT	“c”	“Content-Type”
FORCE_LONG	COMPACT	“Content-Type”	“c”
FORCE_LONG	COMPACT	“c”	“c”
FORCE_LONG	LONG	“Content-Type”	“Content-Type”
FORCE_LONG	LONG	“c”	“Content-Type”

Table 3-3 shows the system header format that results when setting the header format with `WlssSipServletResponse.setUseHeaderForm` given different container configuration values.

Table 3-3 API Behavior for System Headers

SIP Servlet Container Header Configuration (use-compact-form Setting)	WlssSipServletMessage. setUseHeaderForm Setting	Resulting Contact Header
COMPACT	DEFAULT	“m”
COMPACT	COMPACT	“m”
COMPACT	LONG	“Contact”
LONG	DEFAULT	“Contact”
LONG	COMPACT	“m”
LONG	LONG	“Contact”

Table 3-3 API Behavior for System Headers

FORCE_COMPACT	DEFAULT	“m”
FORCE_COMPACT	COMPACT	“m”
FORCE_COMPACT	LONG	“Contact”
FORCE_LONG	DEFAULT	“Contact”
FORCE_LONG	COMPACT	“m”
FORCE_LONG	LONG	“Contact”

CHAPTER 4

Composing SIP Applications

The following sections describe how to use WebLogic SIP Server application composition features:

- "Overview of SIP Application Composition" on page 4-1
- "Application Composition Model" on page 4-2
- "Sample Composer Application" on page 4-3
- "Troubleshooting Application Composition" on page 4-5

Note: The application composition strategy described in this section is applicable only to WebLogic SIP Server 3.1. Application composition techniques will likely change in future versions of the SIP Servlet specification, rendering the current composition techniques obsolete.

Overview of SIP Application Composition

Application composition is the process of "chaining" multiple SIP applications, such as Proxies, User Agent Servers (UAS), User Agent Clients (UAC), redirect servers, and Back-to-Back User Agents (B2BUA), into a logical path that processes a given SIP request. The sections that follow describe an application composition programming model that can be deployed to WebLogic SIP Server. By using this programming model, you can define a set of applications responsible for processing a given initial SIP request, as well as the logic for how and when each application should modify the request. The WebLogic SIP Server container ensures that each application remains on the call path for subsequent message processing requests.

Application Composition Model

The basic WebLogic SIP Server application composition model involves creating a main "composer" application that examines initial SIP requests to determine which deployed applications should process the request. For example, a composer application may examine the user specified in the Request-URI header and select applications based on the user's subscription level.

The composer application should insert one or more Route headers into the request, with each Route header specifying the name and location of a deployed SIP application that should process the request. Application names are defined similar to user addresses, using the format:

application@address

where *application* is the deployment name of the SIP application and *address* is the address of the load balancer used to contact the WebLogic SIP Server installation, the cluster address, or the listen address of the server itself (for example, `proxyapp1@mycompany.com`). The order of the Route headers in the message should dictate the required order of application execution. The Request-URI header of the initial request should remain unchanged.

After inserting Route headers to chain the required applications, the composer application then proxies the message using the original Request-URI. The container examines the contents of the initial Route header in the request. If the user name portion of the route header matches a deployed application name and the address matches a configured server address, then the container delivers the request to the named application.

After each application processes the request, the top Route header is removed the message is then proxied to the next application as shown in Figure 4-1.

Figure 4-1 Composed Application Model

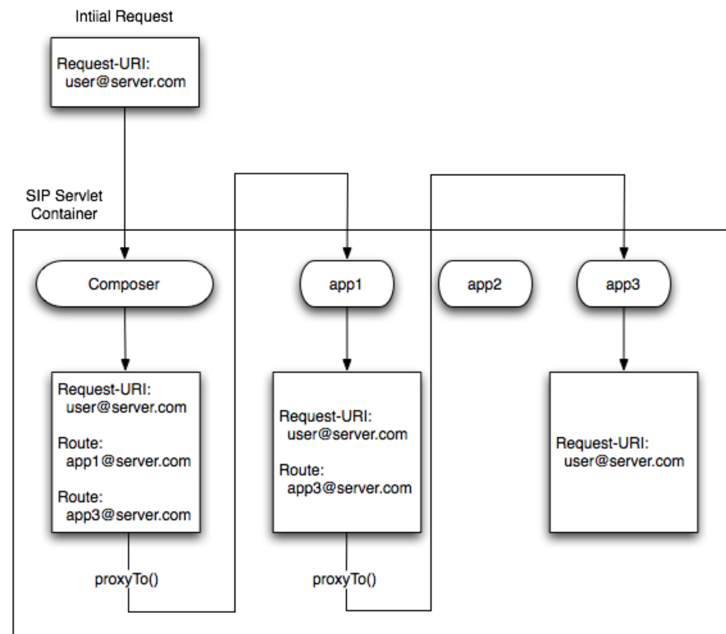

If a request is proxied to another server, the SIP Servlet container inserts the session IDs of chained applications into the Record-Route header of the message. The session IDs ensure that each server hosting a chained application remains in the call path for subsequent requests.

Sample Composer Application

Listing 4-1 shows the organization of a simple composer application.

Listing 4-1 Sample Composer Application

```
package example;


import javax.servlet.sip.SipFactory;

import javax.servlet.sip.SipServletRequest;
```

```
import javax.servlet.sip.SipURI;
import javax.servlet.sip.SipServlet;
import javax.servlet.ServletConfig;
import javax.servlet.ServletException;
import java.io.IOException;

public class Composer extends SipServlet {
  private SipFactory factory;

  private static final String CLUSTER_ADDRESS = "example.com";

  public void init(ServletConfig sc) throws ServletException {
    super.init(sc);
    factory = (SipFactory)
      getServletContext().getAttribute("javax.servlet.sip.SipFactory");
  }

  protected void doRequest(SipServletRequest req)
    throws ServletException, IOException {

    if (!req.isInitial()) {
      super.doRequest(req);
      return;
    }

    SipURI[] routeSet = getRouteSet(req);
    for (int i = 0; i < routeSet.length; i++) {
```

```
      req.pushRoute(routeSet[i]);
    }

    req.getProxy().proxyTo(req.getRequestURI());
  }

  /*
  * Returns application route set for specified request. Ideally, this route
  * set should be based on the requesting user's subscribed services. In
  * this example, it is fixed for all users.
  */
  private SipURI[] getRouteSet(SipServletRequest req) {
    return new SipURI[] { createRouteURI("app1"), createRouteURI("app2") };
  }

  private SipURI createRouteURI(String appName) {
    SipURI uri = factory.createSipURI(appName, CLUSTER_ADDRESS);
    uri.setLrParam(true);
    return uri;
  }
}
```

Troubleshooting Application Composition

WebLogic SIP Server examines the first Route header in a message to determine two things:

1. Does the username portion of the header match the name of a deployed SIP application?
2. Does the address portion of the header indicate that the application is intended for this WebLogic SIP Server instance?

Both of these conditions must be met in order for the SIP Servlet container to route a request to an application specified in the Route header. When both conditions are met, the SIP container checks the mapping rules to verify that the designated application is capable of handling request. If either condition is not met, Weblogic SIP Server uses the default Servlet mapping rules defined in the Servlet's deployment descriptor to process the request.

For example, if the username portion of the first Route header does not match a deployed application name, default Servlet mapping rules are used to process the request. Always ensure that the composer application embeds the correct application names into Route headers when chaining applications together.

Even if the username matches a deployed application, the address portion must also match one of the configured addresses for the WebLogic SIP Server instance:

- A load balancer URI configured in `sipserver.xml`
- The cluster address for the WebLogic SIP Server engine tier
- A listen address for the server instance itself (default listen address or the listen address of a network channel)

To ensure that the address of an application matches the server address, ensure that the composer application is embedding the proper address string in Route headers. Also ensure that the server instances are configured using the same address string. See loadbalancer and Configuring WebLogic SIP Server Network Resources in *Configuring Network Resources*.

Deployment Order Requirement

When using this application composition model, please keep in mind that the order of deployment for composed applications is significant and the composer application must be deployed before all other applications. Because redeployment changes the deployment order, use application upgrade mechanisms instead of redeployment to update the composer application. Alternately, you can configure the Servlet mapping rules of all other applications in such a way to insure that they are never invoked first (without coming through the composer Servlet).

CHAPTER 5

Developing Converged Applications

The following sections describe how to develop converged HTTP and SIP applications with WebLogic SIP Server:

- "Overview of Converged Applications" on page 5-1
- "Assembling and Packaging a Converged Application" on page 5-2
- "Working with SIP and HTTP Sessions" on page 5-2
- "Using the Converged Application Example" on page 5-5

Overview of Converged Applications

In a *converged application*, SIP protocol functionality is combined with other protocols (generally HTTP) to provide a unified communication service. For example, an online push-to-talk application might enable a customer to initiate a voice call to ask questions about products in their shopping cart. The SIP session initiated for the call is associated with the customer's HTTP session, which enables the employee answering the call to view customer's shopping cart contents or purchasing history.

You assemble converged applications using the basic SIP Servlet directory structure outlined in JSR 116. Converged applications require both a `sip.xml` and a `web.xml` deployment descriptor files.

The HTTP and SIP sessions used in a converged application can be accessed programmatically via a common application session object. WebLogic SIP Server provides an extended API to help you associate HTTP sessions with an application session.

Assembling and Packaging a Converged Application

JSR 116 fully describes the requirements and restrictions for assembling converged applications. The following statements summarize the information in the SIP Servlet specification:

- Use the standard SIP Servlet directory structure for converged applications.
- Store all SIP Servlet files under the `WEB-INF` subdirectory; this ensures that the files are not served up as static files by an HTTP Servlet.
- Include deployment descriptors for both the HTPP and SIP components of your application. This means that both `sip.xml` and `web.xml` descriptors are required. A `weblogic.xml` deployment descriptor may also be included to configure Servlet functionality in the WebLogic SIP Server container.
- Observe the following restrictions on deployment descriptor elements:
  - The `distributable` tag must be present in both `sip.xml` and `web.xml`, or it must be omitted entirely.
  - `context-param` elements are shared for a given converged application. If you define the same `context-param` element in `sip.xml` and in `web.xml`, the parameter must have the same value in each definition.
  - If either the `display-name` or `icons` element is required, the element must be defined in both `sip.xml` and `web.xml`, and it must be configured with the same value in each location.

Working with SIP and HTTP Sessions

As shown in Figure 5-1, each application deployed to the WebLogic SIP Server container has a single `SipApplicationSession`, which can contain one or more `SipSession` and `HttpSession` objects. The basic API provided by `javax.servlet.SipApplicationSession` enables you to iterate through all available sessions available in a given `SipApplicationSession`. However, the basic API specified by JSR 116 does not define methods to obtain a given `SipApplicationSession` or to create or associate HTTP sessions with a `SipApplicationSession`.

Figure 5-1 Sessions in a Converged Application

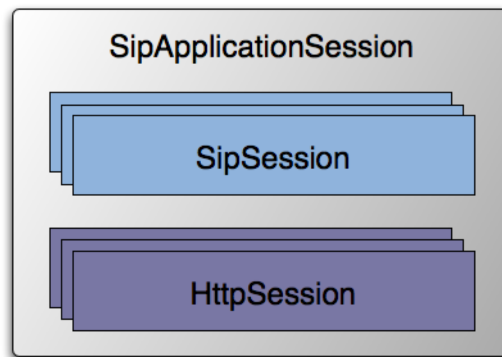

WebLogic SIP Server extends the basic API to provide methods for:

- Creating new HTTP sessions from a SIP Servlet
- Adding and removing HTTP sessions from `SipApplicationSession`
- Obtaining `SipApplicationSession` objects using either the call ID or session ID
- Encoding HTTP URLs with session IDs from within a SIP Servlet

These extended API methods are available in the utility class `com.bea.wcp.util.Sessions`. Table 5-1 provides a summary of each method. See the JavaDoc for more details on this utility class.

Table 5-1 Summary of com.bea.wcp.util.Sessions Methods

Method	Description
getApplicationSession	Obtain a `SipApplicationSession` object with a specified session ID.
getApplicationSessionsByCallId	Obtain an Iterator of `SipApplicationSession` objects associated with the specified call ID.
createHttpSession	Create an HTTP session from within a SIP Servlet. You can modify the HTTP session state and associate the new session with an existing `SipApplicationSession` for later use.

Table 5-1 Summary of com.bea.wcp.util.Sessions Methods

Method	Description
setApplicationSession	Associate an HTTP session with an existing `SipApplicationSession`.
removeApplicationSession	Removes an HTTP session from an existing `SipApplicationSession`.
getEncodeURL	Encodes an HTTP URL with the `jsessionid` of an existing HTTP session object.

Modifying the SipApplicationSession

When using a replicated domain, WebLogic SIP Server automatically provides concurrency control when a SIP Servlet modifies a `SipApplicationSession` object. In other words, when a SIP Servlet modifies the `SipApplicationSession` object, the SIP container automatically locks other applications from modifying the object at the same time.

Non-SIP applications, such as HTTP Servlets, must themselves ensure that the application call state is locked before modifying it in a replicated environment. This is also required if a single SIP Servlet needs to modify other call state objects, such as when a conferencing Servlet joins multiple calls.

To help application developers manage concurrent access to the application session object, WebLogic SIP Server extends the standard `SipApplicationSession` object with `com.bea.wcp.sip.WlssSipApplicationSession`, and adds a new interface, `com.bea.wcp.sip.WlssAction` to encapsulate changes to the session. When these APIs are used, the SIP container ensures that all business logic contained within the `WlssAction` object is executed on a locked copy of the associated `SipApplicationSession` instance.

Listing 5-1 Example Code using WlssSipApplicationSession and WlssAction API

```
SipApplicationSession appSession = ...;

WlssSipApplicationSession wlssAppSession = (WlssSipApplicationSession)
appSession;

wlssAppSession.doAction(new WlssAction() {

       public Object run() throws Exception {
```

```
        // Add all business logic here.
        appSession.setAttribute("counter", latestCounterValue);
        sipSession.setAttribute("currentState", latestAppState);
        // The SIP container ensures that the run method is invoked
        // while the application session is locked.
        return null;
      }
    });
```

Using the Converged Application Example

WebLogic SIP Server includes a sample converged application that uses the com.bea.`wcp.util.Sessions` API. All source code, deployment descriptors, and build files for the example are installed in `WLSS_HOME\samples\sipserver\examples\src\convergence`. See the `readme.html` file in the example directory for instructions about how to build and run the example.

CHAPTER 6

Using the Diameter Base Protocol API

The following sections provide an overview of using the WebLogic SIP Server Diameter Base protocol implementation to create your own Diameter applications:

- "Overview of Diameter Protocol Support" on page 6-1
- "Overview of the Diameter API" on page 6-2
- "Working with Diameter Nodes" on page 6-4
- "Implementing a Diameter Application" on page 6-5
- "Working with Diameter Sessions" on page 6-6
- "Working with Diameter Messages" on page 6-7
- "Working with AVPs" on page 6-9

Overview of Diameter Protocol Support

Diameter is a peer-to-peer protocol that involves delivering attribute-value pairs (AVPs). A Diameter message includes a header and one or more AVPs. The collection of AVPs in each message is determined by the type of Diameter application, and the Diameter protocol also allows for extension by adding new commands and AVPs. Diameter enables multiple peers to negotiate their capabilities with one another, and defines rules for session handling and accounting functions.

WebLogic SIP Server includes an implementation of the base Diameter protocol that supports the core functionality and accounting features described in RFC 3588. WebLogic SIP Server uses the

base Diameter functionality to implement multiple Diameter applications, including the Sh, Rf, and Ro applications described later in this document.

You can also use the base Diameter protocol to implement additional client and server-side Diameter applications. The base Diameter API provides a simple, Servlet-like programming model that enables you to combine Diameter functionality with SIP or HTTP functionality in a converged application.

The sections that follow provide an overview of the base Diameter protocol packages, classes, and programming model used for developing client and server-side Diameter applications. See also the following sections for information about using the provided Diameter protocol applications in your SIP Servlets:

- "Using the Diameter Sh Interface Application" on page 7-1 describes how to access and manage subscriber profile data using the Diameter Sh application.
- "Using the Diameter Rf Interface Application for Offline Charging" on page 8-1 describes how to issue offline charging requests using the Diameter Rf application.
- "Using the Diameter Ro Interface Application for Online Charging" on page 9-1 describes how to perform online charging using the Diameter Ro application.

Overview of the Diameter API

All classes in the Diameter base protocol API reside in the root `com.bea.wcp.diameter` package. Table 6-1 describes the key classes, interfaces, and exceptions in this package.

Table 6-1 Key Elements of the Diameter Base Protocol API

Category	Element	Description
Diameter Node	Node	A class that represents a Diameter node implementation. A diameter node can represent a client- or server-based Diameter application, as well as a Diameter relay agent.

Table 6-1 Key Elements of the Diameter Base Protocol API

Diameter Applications	Application, ClientApplication	A class that represents a basic Diameter application. ClientApplication extends Application for client-specific features such as specifying destination hosts and realms. All Diameter applications must extend one of these classes to return an application identifier. The classes can also be used directly to create new Diameter sessions.
	ApplicationId	A class that represents the Diameter application ID. This ID is used by the Diameter protocol for routing messages to the appropriate application. The `ApplicationId` corresponds to one of the Auth-Application-Id, Acct-Application-Id, or Vendor-Specific-Application-Id AVPs contained in a Diameter message.
	Session	A class that represents a Diameter session. Applications that perform session-based handling must extend this class to provide application-specific behavior for managing requests and answering messages.
Message Processing	Message, Request, Answer	The Message class is a base class used to represent request and answer message types. Request and Answer extend the base class.
	Command	A class that represents a Diameter command code.
	RAR, RAA	These classes extend the `Request` and `Answer` classes to represent re-authorization messages.
	ResultCode	A class that represents a Diameter result code, and provides constant values for the base Diameter protocol result codes.
AVP Handling	Attribute	A class that provides Diameter attribute information.
	Avp, AvpList	Classes that represent one or more attribute-value pairs in a message. AvpList is also used to represent AVPs contained in a grouped AVP.
	Type	A class that defines the supported AVP datatypes.
Error Handling	DiameterException	The base exception class for Diameter exceptions.
	MessageException	An exception that is raised when an invalid Diameter message is discovered.
	AvpException	An exception that is raised when an invalid AVP is discovered.

Table 6-1 Key Elements of the Diameter Base Protocol API

Supporting Interfaces	Enumerated	An enum value that implements this interface can be used as the value of an AVP of type INTEGER32, INTEGER64, or ENUMERATED.
	SessionListener	An interface that applications can implement to subscribe to messages delivered to a Diameter session.
	MessageFactory	An interface that allows applications to override the default message decoder for received messages, and create new types of `Request` and `Answer` objects. The default decoding process begins by decoding the message header from the message bytes using an instance of `MessageFactory`. This is done so that an early error message can be generated if the message header is invalid. The actual message AVPs are decoded in a separate step by calling `decodeAvps`. AVP values are fully decoded and validated by calling `validate`, which in turn calls `validateAvp` for each partially-decoded AVP in the message.

In addition to these base Diameter classes, accounting-related classes are stored in the `com.bea.wcp.diameter.accounting` package, and credit-control-related classes are stored in `com.bea.wcp.diameter.cc`. See "Using the Diameter Ro Interface Application for Online Charging" on page 9-1 and "Using the Diameter Rf Interface Application for Offline Charging" on page 8-1 for more information about classes in these packages.

Working with Diameter Nodes

A diameter node is represented by the `com.bea.wcp.diameter.Node` class. A Diameter node may host one or more Diameter applications, as configured in the `diameter.xml` file. In order to access a Diameter application, a deployed application (such as a SIP Servlet) must obtain the diameter Node instance and request the application. Listing 6-1 shows the sample code used to access the Rf application.

Listing 6-1 Accessing a Diameter Node and Application

```
ServletContext sc = getServletConfig().getServletContext();

Node node = sc.getAttribute("com.bea.wcp.diameter.Node");
```

```
RfApplication rfApp = (RfApplication)
node.getApplication(Charging.RF_APPLICATION_ID);
```

Diameter Nodes are generally configured and started as part of a WebLogic SIP Server instance. However, for development and testing purposes, you can also run a Diameter node as a standalone process. To do so:

1. Set the environment for your domain:

   ```
   cd ~/bea/user_projects/domains/diameter/bin
   . ./setDomainEnv.sh
   ```

2. Locate the `diameter.xml` configuration file for the Node you want to start:

   ```
   cd ../config/custom
   ```

3. Start the Diameter node, specifying the `diameter.xml` configuration file to use:

   ```
   java com.bea.wcp.diameter.Node diameter.xml
   ```

Implementing a Diameter Application

All Diameter applications must extend either the base `Application` class or, for client applications, the `ClientApplication` class. The model for creating a Diameter application is similar to that for implementing Servlets in the following ways:

- Diameter applications override the `init()` method for initialization tasks.
- Initialization parameters configured for the application in `diameter.xml` are made available to the application.
- A session factory is used to generate new application sessions.

Diameter applications must also implement the `getId()` method to return the proper application ID. This ID is used to deliver Diameter messages to the correct application.

Applications can optionally implement `rcvRequest()` or `rcvAnswer()` as needed. By default, `rcvRequest()` answers with UNABLE_TO_COMPLY, and `rcvRequest()` drops the Diameter message.

Listing 6-2 shows a simple Diameter client application that does not use sessions.

Listing 6-2 Simple Diameter Application

```
public class TestApplication extends ClientApplication {
```

```
    protected void init() {
      log("Test application initialized.");
    }
    public ApplicationId getId() {
      return ApplicationId.BASE_ACCOUNTING;
    }
    public void rcvRequest(Request req) throws IOException {
      log("Got request: " + req.getHopByHopId());
      req.createAnswer(ResultCode.SUCCESS).send();
    }
  }
```

Working with Diameter Sessions

The base Session class represents a Diameter session. If you extend the base Session class, you must implement either `rcvRequest()` or `rcvAnswer()`, and may implement both methods.

The base Application class is used to generate new Session objects. After a session is created, all session-related messages are delivered directly to the session object. The WebLogic SIP Server container automatically generates the session ID and encodes the ID in each message. Session attributes are supported much in the same fashion as attributes in `SipApplicationSession`.

Listing 6-3 shows a simple Diameter session implementation.

Listing 6-3 Simple Diameter Session

```
public class TestSession extends Session {
  public TestSession(TestApplication app) {
    super(app);
  }
  public void rcvRequest(Request req) throws IOException {
    getApplication().log("rcvReuest: " + req.getHopByHopId());
```

```
    req.createAnswer(ResultCode.SUCCESS).send();
  }
}
```

To use the sample session class, the `TestApplication` in Listing 6-2 would need to add a factory method:

```
public class TestApplication extends Application {
  ...
  public TestSession createSession() {
    return new TestSession(this);
  }
}
```

`TestSession` could then be used to create new requests as follows:

```
TestSession session = testApp.createSession();
Request req = session.creatRequest();
req.send();
```

The answer is delivered directly to the Session object.

Working with Diameter Messages

The base `Message` class is used for both Request and Answer message types. A Message always includes an application ID, and optionally includes a session ID. By default, messages are handled in the following manner:

1. The message bytes are parsed.
2. The application and session ID values are determined.
3. The message is delivered to a matching session or application using the following rules:
   a. If the Session-Id AVP is present, the associated Session is located and the session's `rcvMessage()` method is called.
   b. If there is no Session-Id AVP present, or if the session cannot be located, the Diameter application's `rcvMessage()` method is called

c. If the application cannot be located, an UNABLE_TO_DELIVER response is generated.

The message type is determined from the Diameter command code. Certain special message types, such as RAR, RAA, ACR, ACA, CCR, and CCA, have getter and setter methods in the `Message` object for convenience.

Sending Request Messages

Either a `Session` or `Application` can originate and receive request messages. Requests are generated using the `createRequest()` method. You must supply a command code for the new request message. For routing purposes, the destination host or destination realm AVPs are also generally set by the originating session or application.

Requests can be sent asynchronously using the `send()` method, or synchronously using the blocking `sendAndWait()` method. Answers for requests that were sent asynchronously are delivered to the originating session or application. You can specify a request timeout value when sending the message, or can use the global `request-timeout` configuration element in `diameter.xml`. An UNABLE_TO_DELIVER result code is generated if the timeout value is reached before an answer is delivered. `getResultCode()` on the resulting Answer returns the result code.

Sending Answer Messages

New answer messages are generated from the `Request` object, using `createAnswer()`. All generated answers should specify a ResultCode and an optional Error-Message AVP value. The `ResultCode` class contains pre-defined result codes that can be used.

Answers are delivered using the `send()` method, which is always asynchronous (non-blocking).

Received answers can be obtained using `Request.getAnswer()`. After receiving an answer, you can use `getSession()` to obtain the relevant session ID and `getResultCode()` to determine the result. You can also use `Answer.getRequest()` to obtain the original request message.

Creating New Command Codes

The `Command` class represents pre-defined commands codes for the Diameter base protocol, and can be used to create new command codes. Command codes share a common name space based on the code itself.

The `define()` method enables you to define codes, as in:

```
static final Command TCA = Command.define(1234, “Test-Request”, true, true);
```

The `define()` method registers a new Command, or returns a previous command definition if one was already defined. Commands can be compared using the reference equality operator (==).

Working with AVPs

The `Avp` class represents a Diameter attribute-value pair. You can create new AVPs with an attribute value in the following way:

```
Avp avp = new Avp(Attribute.ERROR_MESSAGE, "Bad request");
```

You can also specify the attribute name directly, as in:

```
Avp avp = new Avp("Error-Message", "Bad request");
```

The value that you specify must be valid for the specified attribute type.

To create a grouped AVP, use the `AvpList` class, as in:

```
AvpList avps = new AvpList();

avps.add(new Avp("Event-Timestamp", 1234));

avps.add(new Avp("Vendor-Id", 1111));
```

Creating New Attributes

The Attribute class represents an AVP attribute, and includes the AVP code, name, flags, optional vendor ID, and type of attribute. The class also maintains a registry of defined attributes. All attributes share a common namespace based on the attribute code and vendor ID.

The `define()` method enables you to define new attributes, as in:

```
static final Attribute TEST = Attribute.define(1234, "Test-Attribute", 0,
Attribute.FLAG_MANDATORY, Type.INTEGER32);
```

Table 6-2 lists the available attribute types and describes how they are mapped to Java types.

The `define()` method registers a new attribute, or returns a previous definition if one was already defined. Attributes can be compared using the reference equality operator (==).

Table 6-2 Attribute Types

Diameter Type	Type Constant	Java Type
Integer32	Type.INTEGER32	Integer
Integer64	Type.INTEGER64	Long

Table 6-2 Attribute Types

Diameter Type	Type Constant	Java Type
Float32	Type.FLOAT32	Float
OctetString	Type.BYTES	ByteBuffer (read-only)
UTF8String	Type.STRING	String
Address	Type.ADDRESS	InetAddress
Grouped	Type.GROUPED	AvpList

Creating Converged Diameter and SIP Applications

The Diameter API enables you to create converged applications that utilize both SIP and Diameter functionality. A SIP Servlet can access an available Diameter application via the Diameter Node, as shown in Listing 6-4.

Listing 6-4 Accessing the Rf Application from a SIP Servlet

```
ServletContext sc = getServletConfig().getServletContext();

Node node = (Node) sc.getAttribute("com.bea.wcp.diameter.Node");

RfApplication rfApp = (RfApplication)
node.getApplication(Charging.RF_APPLICATION_ID);
```

WebLogic SIP Server automatically links the Diameter session to the currently-active call state by encoding the Call-id into the Diameter session ID. When a Diameter message is received, the container automatically retrieves the associated call state and locates the Diameter session. A Diameter session is are serializable, so you can store the session as an attribute in a the `SipApplicationSession` object, or vice versa.

Converged applications can use the Diameter `SessionListener` interface to receive notification when a Diameter message is received by the session. The `SessionListener` interface defines a single method, `rcvMessage()`. `Listing 6-5` shows an example of how to implement the method.

Listing 6-5 Implementing SessionListener

```
Session session = app.createSession();
session.setListener(new SessionListener() {
  public void rcvMessage(Message msg) {
    if (msg.isRequest()) System.out.println("Got request!");
  }
});
```

Note: The `SessionListener` implementation must be serializable for distributed applications.

CHAPTER 7

Using the Diameter Sh Interface Application

The following sections describe how to use the Diameter Sh interface application, based on the WebLogic SIP Server Diameter protocol implementation, in your own applications:

- "Overview of Profile Service API and Sh Interface Support" on page 7-1
- "Enabling the Sh Interface Provider" on page 7-2
- "Overview of the Profile Service API" on page 7-2
- "Creating a Document Key for Application-Managed Profile Data" on page 7-3
- "Using a Constructed Document Key to Manage Profile Data" on page 7-5
- "Monitoring Profile Data with ProfileListener" on page 7-6

Overview of Profile Service API and Sh Interface Support

The IMS specification defines the Sh interface as the method of communication between the Application Server (AS) function and the Home Subscriber Server (HSS), or between multiple IMS Application Servers. The AS uses the Sh interface in two basic ways:

- To query or update a user's data stored on the HSS
- To subscribe to and receive notifications when a user's data changes on the HSS

The user data available to an AS may be defined by a service running on the AS (*repository data*), or it may be a subset of the user's IMS profile data hosted on the HSS. The Sh interface specification, 3GPP TS 29.328, defines the IMS profile data that can be queried and updated via

Sh. All user data accessible via the Sh interface is presented as an XML document with the schema defined in 3GPP TS 29.328.

The IMS Sh interface is implemented as a provider to the base Diameter protocol support in WebLogic SIP Server. The provider transparently generates and responds to the Diameter command codes defined in the Sh application specification. A higher-level Profile Service API enables SIP Servlets to manage user profile data as an XML document using XML Document Object Model (DOM). Subscriptions and notifications for changed profile data are managed by implementing a profile listener interface in a SIP Servlet.

Figure 7-1 Profile Service API and Sh Provider Implementation

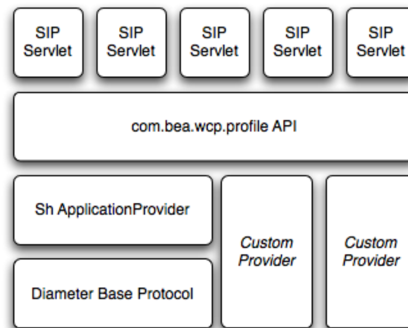

WebLogic SIP Server 3.1 includes a provider for the Diameter Sh interface. Future versions of WebLogic SIP Server may include new providers to support additional interfaces defined in the IMS specification. Applications using the profile service API will be able to use additional providers as they are made available.

Enabling the Sh Interface Provider

See Configuring Diameter Sh Client Nodes and Relay Agents in *Configuring Network Resources* for full instructions on setting up Diameter support.

Overview of the Profile Service API

WebLogic SIP Server provides a simple profile service API that SIP Servlets can use to query or modify subscriber profile data, or to manage subscriptions for receiving notifications about

changed profile data. Using the API, a SIP Servlet explicitly requests user profile documents via the Sh provider application. The provider returns an XML document, and the Servlet can then use standard DOM techniques to read or modify profile data in the local document. Updates to the local document are applied to the HSS after a "put" operation.

Creating a Document Key for Application-Managed Profile Data

Servlets that manage profile data can explicitly obtain an Sh XML document from a factory using a key, and then work with the document using DOM.

The document selector key identifies the XML document to be retrieved by a Diameter interface, and uses the format `protocol://uri/reference_type[/access_key]`.

The `protocol` portion of the selector identifies the Diameter interface provider to use for retrieving the document. In WebLogic SIP Server version 3.1, only the Sh interface is provided, so only Sh XML documents (beginning with the `sh://` protocol designation) can be retrieved.

With Sh document selectors, the next element, `uri`, generally corresponds to the User-Identity or Public-Identity of the user whose profile data is being retrieved. If you are requesting an Sh data reference of type LocationInformation or UserState, the URI value can be the User-Identity or MSISDN for the user.

Table 7-1 summarizes the possible URI values that can be supplied depending on the Sh data reference you are requesting. 3GPP TS 29.328 describes the possible data references and associated reference types in more detail.

Table 7-1 Possible URI Values for Sh Data References

Sh Data Reference Number	Data Reference Type	Possible URI Value in Document Selector
0	RepositoryData	User-Identity or Public-Identity
10	IMSPublicIdentity	
11	IMSUserState	
12	S-CSCFName	
13	InitialFilterCriteria	

Table 7-1 Possible URI Values for Sh Data References

Sh Data Reference Number	Data Reference Type	Possible URI Value in Document Selector
14	LocationInformation	User-Identity or MSISDN
15	UserState	
17	Charging information	User-Identity or Public-Identity
17	MSISDN	

The final element of the document selector, `reference_type`, specifies the data reference type being requested. For some data reference requests, only the `uri` and `reference_type` are required. Other Sh requests use an access key, which requires a third element in the document selector corresponding to the value of the Attribute-Value Pair (AVP) defined in the key.

Table 7-2 summarizes the required document selector elements for each type of Sh data reference request.

Table 7-2 Summary of Document Selector Elements for Sh Data Reference Requests

Data Reference Type	Required Document Selector Elements	Example Document Selector
RepositoryData	sh://uri/reference_type/Service-Indication	sh://sip:user@bea.com/RepositoryData/Call Screening/
IMSPublicIdentity	sh://uri/reference_type/[*Identity-Set*] where *Identity-Set* is one of: • All-Identities • Registered-Identities • Implicit-Identities	sh://sip:user@bea.com/IMSPublicIdentity/Registered-Identities
IMSUserState	sh://uri/reference_type	sh://sip:user@bea.com/IMSUserState/
S-CSCFName	sh://uri/reference_type	sh://sip:user@bea.com/S-CSCFName/
InitialFilterCriteria	sh://uri/reference_type/Server-Name	sh://sip:user@bea.com/InitialFilterCriteria/www.bea.com/

Table 7-2 Summary of Document Selector Elements for Sh Data Reference Requests

Data Reference Type	Required Document Selector Elements	Example Document Selector
LocationInformation	sh://uri/reference_type/(CS-Domain | PS-Domain)	sh://sip:user@bea.com/LocationInformation/CS-Domain/
UserState	sh://uri/reference_type/(CS-Domain | PS-Domain)	sh://sip:user@bea.com/UserState/PS-Domain/
Charging information	sh://uri/reference_type	sh://sip:user@bea.com/Charging information/
MSISDN	sh://uri/reference_type	sh://sip:user@bea.com/MSISDN/

Using a Constructed Document Key to Manage Profile Data

WebLogic SIP Server provides a helper class, `com.bea.wcp.profile.ProfileService`, to help you easily retrieve a profile data document. The `getDocument()` method takes a constructed document key, and returns a read-only `org.w3c.dom.Document` object. To modify the document, you make and edit a copy, then send the modified document and key as arguments to the `putDocument()` method.

Note: If Diameter Sh client node services are not available on the WebLogic SIP Server instance when `getDocument()` the profile service throws a "No registered provider for protocol" exception.

WebLogic SIP Server caches the documents returned from the profile service for the duration of the service method invocation (for example, when a `doRequest()` method is invoked). If the service method requests the same profile document multiple times, the subsequent requests are served from the cache rather than by re-querying the HSS.

Listing 7-1 shows a sample SIP Servlet that obtains and modifies profile data.

Listing 7-1 Sample Servlet Using ProfileService to Retrieve and Write User Profile Data

```
package demo;

import com.bea.wcp.profile.*;
```

```
import javax.servlet.sip.SipServletRequest;
import javax.servlet.sip.SipServlet;
import org.w3c.dom.Document;
import java.io.IOException;
public class MyServlet extends SipServlet {
      private ProfileService psvc;
      public void init() {
        psvc = (ProfileService)
getServletContext().getAttribute(ProfileService.PROFILE_SERVICE);
      }
      protected void doInvite(SipServletRequest req) throws IOException {
        String docSel = "sh://" + req.getTo() + "/IMSUserState/";
        // Obtain and change a profile document.
        Document doc = psvc.getDocument(docSel); // Document is read only.
        Document docCopy = (Document) doc.cloneNode(true);
        // Modify the copy using DOM.
        psvc.putDocument(docSel, docCopy); // Apply the changes.
      }
}
```

Monitoring Profile Data with ProfileListener

The IMS Sh interface enables applications to receive automatic notifications when a subscriber's profile data changes. WebLogic SIP Server provides an easy-to-use API for managing profile data subscriptions. A SIP Servlet registers to receive notifications by implementing the `com.bea.wcp.profile.ProfileListener` interface, which consists of a single `update` method that is automatically invoked when a change occurs to profile to which the Servlet is subscribed. Notifications are not sent if that same Servlet modifies the profile information (for example, if a user modifies their own profile data).

Note: In a replicated environment, Diameter relay nodes always attempt to push notifications directly to the engine tier server that subscribed for profile updates. If that engine tier

server is unavailable, another server in the engine tier cluster is chosen to receive the notification. This model succeeds because session information is stored in the data tier, rather than the engine tier.

Prerequisites for Listener Implementations

In order to receive a call back for subscribed profile data, a SIP Servlet must do the following:

- Implement `com.bea.wcp.profile.ProfileListener`.
- Create one or more subscriptions using the `subscribe` method in the `com.bea.wcp.profile.ProfileService` helper class.
- Register itself as a listener using the `listener` element in `sip.xml`.

"Implementing ProfileListener" on page 7-7 describes how to implement `ProfileListener` and use the `susbscribe` method. In addition to having a valid listener implementation, the Servlet must declare itself as a listener in the `sip.xml` deployment descriptor file. For example, it must add a `listener` element declaration similar to:

```
<listener>
   <lisener-class>com.mycompany.MyLisenerServlet</listener-class>
</listener>
```

Implementing ProfileListener

Actual subscriptions are managed using the `subscribe` method of the `com.bea.wcp.profile.ProfileService` helper class. The subscribe method requires that you supply the current `SipApplicationSession` and the key for the profile data document you want to monitor. See "Creating a Document Key for Application-Managed Profile Data" on page 7-3.

Applications can cancel subscriptions by calling `ProfileSubscription.cancel()`. Also, pending subscriptions for an application are automatically cancelled if the application session is terminated.

Listing 7-2 shows sample code for a Servlet that implements the `ProfileListener` interface.

Listing 7-2 Sample Servlet Implementing ProfileListener Interface

```
package demo;

    import com.bea.wcp.profile.*;
```

```
import javax.servlet.sip.SipServletRequest;
import javax.servlet.sip.SipServlet;
import org.w3c.dom.Document;
import java.io.IOException;
public class MyServlet extends SipServlet implements ProfileListener {
  private ProfileService psvc;
  public void init() {
    psvc = (ProfileService)
getServletContext().getAttribute(ProfileService.PROFILE_SERVICE);
  }
  protected void doInvite(SipServletRequest req) throws IOException {
    String docSel = "sh://" + req.getTo() + "/IMSUserState/";
    // Subscribe to profile data.
    psvc.subscribe(req.getApplicationSession(), docSel, null);
}
  public void update(ProfileSubscription ps, Document document) {
    System.out.println("IMSUserState updated: " +
ps.getDocumentSelector());
  }
}
```

CHAPTER 8

Using the Diameter Rf Interface Application for Offline Charging

The following sections describe how to use the Diameter Rf interface application, based on the WebLogic SIP Server Diameter protocol implementation, in your own applications:

- "Overview of Rf Interface Support" on page 8-1
- "Understanding Offline Charging Events" on page 8-2
- "Configuring the Rf Application" on page 8-4
- "Using the Offline Charging API" on page 8-5

Overview of Rf Interface Support

Offline charging is used for network services that are paid for periodically. For example, a user may have a subscription for voice calls that is paid monthly. The Rf protocol allows an IMS Charging Trigger Function (CTF) to issue offline charging events to a Charging Data Function (CDF). The charging events can either be one-time events or may be session-based.

WebLogic SIP Server provides a Diameter Offline Charging Application that can be used by deployment application to generate charging events based on the Rf protocol. The offline charging application uses the base Diameter protocol implementation, and allows any application deployed on WebLogic SIP Server to act as CTF to a configured CDF.

For basic information about offline charging, see RFC 3588: Diameter Base Protocol. For more information about the Rf protocol, see 3GPP TS 32.299.

Understanding Offline Charging Events

For both event and session based charging, the CTF implements the accounting state machine described in RFC 3588. The server (CDF) implements the accounting state machine “SERVER, STATELESS ACCOUNTING” as specified in RFC 3588.

The reporting of offline charging events to the CDF is managed through the Diameter Accounting Request (ACR) message. Rf supports the ACR event types described in Table 8-1.

Table 8-1 Rf ACR Event Types

Request	Description
START	Starts an accounting session.
INTERIM	Updates an accounting session.
STOP	Stops an accounting session
EVENT	Indicates a one-time accounting event.

The START, INTERIM, and STOP event types are used for session-based accounting. The EVENT type is used for event based accounting, or to indicate a failed attempt to establish a session.

Event-Based Charging ¶

Event-based charging events are reported through the ACR EVENT message. Listing 8-1 shows the basic message flow.

Listing 8-1 Message Flow for Event-Based Charging

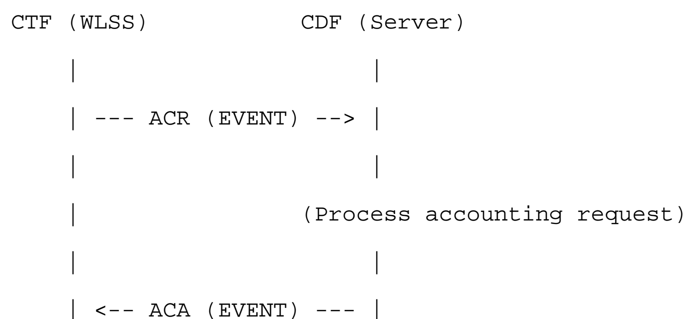

```
CTF (WLSS)            CDF (Server)
    |                      |
    | --- ACR (EVENT) -->  |
    |                      |
    |                 (Process accounting request)
    |                      |
    | <-- ACA (EVENT) ---  |
```

```
        |                         |
```

Session-Based Charging ¶

Session-based charging uses the ACR START, INTERIM, and STOP requests to report usage to the CDF. During a session, the CTF may report multiple ACR INTERIM requests depending on the session lifecycle. Listing 8-2 shows the basic message flow

Listing 8-2 Message Flow for Session-Based Charging

```
CTF (WLSS)             CDF (Server)
    |                        |
    | --- ACR (START) ----> |
    |                        |
    |                  (Open CDR)
    |                        |
    | <-- ACA (START) ----- |
    |                        |
   ...                      ...
    | --- ACR (INTERIM) --> |
    |                        |
    |                  (Update CDR)
    |                        |
    | <-- ACA (INTERIM) --- |
   ...                      ...
    | --- ACR (STOP) -----> |
    |                        |
    |                  (Close CDR)
    |                        |
```

```
| <-- ACA (STOP) ------ |
|                       |
```

Here, ACA START is sent a receipt of a service request by WebLogic SIP Server. ACA INTERIM is typically sent upon expiration of the AII timer. ACA STOP is issued upon request for service termination by WebLogic SIP Server.

Configuring the Rf Application

The `RfApplication` is packaged as a Diameter application similar to the Sh application used for managing profile data. The Rf Diameter application can be configured and enabled by editing the Diameter configuration file located in *DOMAIN_ROOT*`/config/custom/diameter.xml`, or by using the Diameter console extension. Additionally, configuration of both the CDF realm and host can be specified using the cdf.realm and cdf.host initialization parameters to the Diameter Rf application.

Listing 8-3 shows a sample excerpt from `diameter.xml` that enables Rf with a CDF realm of "bea.com" and host "cdf.bea.com:"

Listing 8-3 Sample Rf Application Configuration (diameter.xml)

```
<application>
  <application-id>3</application-id>
  <accounting>true</accounting>
  <class-name>com.bea.wcp.diameter.charging.RfApplication</class-name>
  <param>
    <name>cdf.realm</name>
    <value>bea.com</value>
  </param>
  <param>
    <name>cdf.host</name>
    <value>cdf.bea.com</value>
  </param>
```

```
</application>
```

Because the `RfApplication` uses the Diameter base accounting messages, its Diameter application id is 3 and there is no vendor ID.

Using the Offline Charging API

WebLogic SIP Server provides an offline charging API to enable any deployed application to act as a CTF and issue offline charging events. This API supports both event-based and session-based charging events.

The classes in package `com.bea.wcp.diameter.accounting` provide general support for Diameter accounting messages and sessions. Table 8-1 summarizes the classes.

Table 8-2 Diameter Accounting Classes

Class	Description
ACR	An Accounting-Request message.
ACA	An Accounting-Answer message.
ClientSession	A Client-based accounting session.
RecordType	Accounting record type constants.

In addition, classes in package `com.bea.wcp.diameter.charging` support the Rf application specifically. Table 8-1 summarizes the classes.

Table 8-3 Diameter Rf Application Support Classes

Charging	Common definitions for 3GPP charging functions
RfApplication	Offline charging application
RfSession	Offline charging session

The `RfApplication` class can be used to directly send ACR requests for event-based charging. The application also has the option of directly modifying the ACR request before it is sent out. This is necessary in order for an application to add any custom AVPs to the request.

In particular, an application should set the Service-Information AVP it carries the service-specific parameters for the CDF. The Service-Information AVP of the ACR request is used to send the application-specific charging service information from the CTF (WLSS) to the CDF (Charging

Server). This is a grouped AVP whose value will depend on the application and its charging function. The Offline Charging API allows the application to set this information on the request before it is sent out.

For session-based accounting, the `RfApplication` class can also be used to create new accounting sessions for generating session-based charging events. Each accounting session is represented by an instance of `RfSession`, which encapsulates the accounting state machine for the session.

Accessing the Rf Application

If the Rf application is deployed, then applications deployed on WebLogic SIP Server can obtain an instance of the application from the Diameter node (`com.bea.wcp.diameter.Node` class). Listing 8-4 shows the sample Servlet code used to obtain the Diameter `Node` and access the Rf application.

Listing 8-4 Accessing the Rf Application

```
ServletContext sc = getServletConfig().getServletContext();
Node node = sc.getAttribute("com.bea.wcp.diameter.Node");
RfApplication rfApp = (RfApplication)
node.getApplication(Charging.RF_APPLICATION_ID);
```

Applications can safely use a single instance of `RfApplication` to issue offline charging requests concurrently, in multiple threads. Each instance of `RfSession` actually holds the per-session state unique to each call.

Implementing Session-Based Charging

For session-based charging requests, an application first uses the `RfApplication` to create an instance of `RfSession`. The application can then use the session object to create one or more charging requests.

The first charging request must be an ACR START request, followed by zero or more ACR INTERIM requests. The session ends with an ACR STOP request. Upon receipt of the corresponding ACA STOP message, the `RfApplication` automatically terminates the `RfSession`.

Listing 8-5 shows the sample code used to start a new session-based accounting session.

Listing 8-5 Starting a Session-Based Account Session

```
RfSession session = rfApp.createSession();
sipRequest.getApplicationSession().setAttribute("RfSession", session);
ACR acr = session.createACR(RecordType.START);
acr.addAvp(Charging.SERVICE_INFORMATION, ...);
ACA aca = acr.sendAndWait(1000);
if (!aca.getResultCode().isSuccess()) {
  ... error ...
}
```

In Listing 8-5, the `RfSession` is stored as a SIP application session attribute so that it can be used to send additional accounting requests as the call progresses. Listing 8-6 shows how to send an INTERIM request.

Listing 8-6 Sending an INTERIM request

```
  RfSession session = (RfSession)
req.getApplicationSession().getAttribute("RfSession");
  ACR acr = session.createACR(RecordType.INTERIM);
  ACA aca = acr.sendAndWait(1000);
  if (!aca.getResultCode().isSuccess()) {
    ... error ...
  }
```

An application may want to send one or more ACR INTERIM requests while a call is in progress. The frequency of ACR INTERIM requests is usually based on the Acct-Interim-Interval AVP value in the ACA START message sent by the CDF. For this reason, an application timer should be used to send ACR INTERIM requests at the requested interval. See 3GPP TS 32.299 for more details about interim requests.

Sending Asynchronous Requests

Applications will generally use the synchronous `sendAndWait()` method. However, if latency is critical, an asynchronous API is provided wherein the application Servlet is asynchronously notified when an answer message is received from the CDF. To use the asynchronous API, an application first registers an instance of `SessionListener` in order to asynchronously receive messages delivered to the session, as shown in Listing 8-7

Listing 8-7 Registering a SessionListener

```
RfSession session = rfApp.createSession();
session.setAttribute("SAS", sipReq.getApplicationSession());
session.setListener(this);
```

Attributes can be stored in an `RfSession` instance similar to the way SIP application session attributes are stored. In the above example, the associated SIP application was stored as an `RfSession` so that it is available to the listener callback.

When a Diameter request or answer message is received from the CDF, the application Servlet is notified by calling the `rcvMessage(Message msg)` method. The associated SIP application session can then be retrieved from the `RfSession` if it was stored as a session attribute, as shown in Listing 8-8.

Listing 8-8 Retrieving the RfSession after a Notification

```
public void rcvMessage(Message msg) {
  if (msg.getCommand() != Command.ACA) {
    if (msg.isRequest()) {
      ((Request) msg).createAnswer(ResultCode.UNABLE_TO_COMPLY,
"Unexpected request").send();
    }
    return;
  }
  ACA aca = (ACA) msg;
```

```
    RfSession session = (RfSession) aca.getSession();
    SipApplicationSession appSession = (SipApplicationSession)
session.getAttribute("SAS");
    ...
  }
```

Specifying the Session Expiration

The Acct-Interim-Interval (AII) timer value is used to indicate the expiration time of an Rf accounting session. It is specified when ACR START is sent to the CDF to initiate the accounting session. The CDF responds with its own AII value, which must be used by the CTF to start a timer upon whose expiration an ACR INTERIM message must be sent. This INTERIM message informs the CDF that the session is still in use. Otherwise, the CDF terminates the session automatically.

It is the application's responsibility to send ACR INTERIM messages, because these are used to send updated Service-Information data to the CDF. BEA recommends creating a ServletTimer that is set to expire according to the AII value. When the timer expires, the application should send an ACR INTERIM message with the updated service information data.

Implementing Event-Based Charging

For an event-based charging request, the charging request is a one-time event and the session is automatically terminated upon receipt of the corresponding EVENT ACA message. The `sendAndWait(long timeout)` method can be used to synchronously send the EVENT request and block the thread until a response has been received from the CDF. Listing 8-9 shows an example that uses an `RfSession` for sending an event-based charging request.

Listing 8-9 Event-Based Charging Using RfSession

```
  RfSession session = rfApp.createSession();
  ACR acr = session.createACR(RecordType.EVENT);
  acr.addAvp(Charging.SERVICE_INFORMATION, ...);
  ACA aca = acr.sendAndWait(1000);
  if (!aca.getResultCode().isSuccess()) {
    ... send error response ...
```

```
}
```

For convenience, it is also possible send event-based charging requests using the `RfApplication` directly, as shown in Listing 8-10.

Listing 8-10 Event-Based Charging Using RfApplication

```
ACR acr = rfApp.createEventACR();
acr.addAvp(Charging.SERVICE_INFORMATION, ...);
ACA aca = acr.sendAndWait(1000);
```

Internally, the `RfApplication` creates an instance of `RfSession` associated with the ACR request, so this method is equivalent to creating the session explicitly.

For both session and event based accounting, the `RfSession` class automatically handles creating session IDs, as well as updating the Accounting-Record-Number AVP used to sequence messages within the same accounting session.

In the above cases the applications waits for up to 1000 ms to receive an answer from the CDF. If no answer is received within that time, the Diameter core delivers an UNABLE_TO_COMPLY error response to the application, and cancels the request. If no timeout is specified with `sendAndWait()`, then the default request timeout of 30 seconds is used. This default value can be configured using the Diameter console extension.

Using the Accounting Session State ¶

The accounting session state for offline charging is serializable, so it can be stored as a SIP application session attribute. Because the client APIs are synchronous, it is not necessary to maintain any state for the accounting session once the Servlet has finished handling the call.

For event-based charging events it is not necessary for the application to maintain any accounting session state because it is only used internally, and is disposed once the ACA response has been received.

CHAPTER 9

Using the Diameter Ro Interface Application for Online Charging

The following sections describe how to use the Diameter Rf interface application, based on the WebLogic SIP Server Diameter protocol implementation, in your own applications:

- "Overview of Ro Interface Support" on page 9-1
- "Understanding Credit Authorization Models" on page 9-2
- "Configuring the Ro Application" on page 9-3
- "Overview of the Online Charging API" on page 9-4
- "Accessing the Ro Application" on page 9-5
- "Implementing Session-Based Charging" on page 9-6
- "Sending Credit-Control-Request Messages" on page 9-8
- "Handling Failures" on page 9-9

Overview of Ro Interface Support

Online charging, also known as credit-based charging, is used to charge prepaid services. A typical example of a prepaid service is a calling card purchased for voice or video. The Ro protocol allows a Charging Trigger Function (CTF) to issue charging events to an Online Charging Function (OCF). The charging events can be immediate, event-based, or session-based.

WebLogic SIP Server provides a Diameter Online Charging Application that deployed applications can use to generate charging events based on the Ro protocol. This enables deployed

applications to act as CTF to a configured OCF. The Diameter Online Charging Application uses the base Diameter protocol that underpins both the Rf and Sh applications.

The Diameter Online Charging Application is based on IETF RFC 4006: Diameter Credit Control Application. However, the application supports only a subset of the RFC 4006 required for compliance with 3GPP TS 32.299: Telecommunication management; Charging management; Diameter charging applications. Specifically, the WebLogic SIP Server Diameter Online Charging Application provides no direct support for service-specific Attribute-Value Pairs (AVPs), but the API that is provided is flexible enough to allow applications to include custom service-specific AVPs in any credit control request.

Understanding Credit Authorization Models

RFC 4006 defines two basic types of credit authorization models:

- Credit authorization with unit reservation, and
- Credit authorization with direct debiting.

Credit authorization with unit reservation can be performed with either event-based or session-based charging events. Credit authorization with direct debiting uses immediate charging events. In both models, the CTF requests credit authorization from the OCF prior to delivering services to the end user. In both models

The sections that follow describe each model in more detail.

Credit Authorization with Unit Determination

RFC 4006 defines both Event Charging with Unit Reservation (ECUR) and Session Charging with Unit Reservation (SCUR). Both charging events are session-based, and require multiple transactions between the CTF and OCF. ECUR begins with an interrogation to reserve units before delivering services, followed by an additional interrogation to report the actual used units to the OCF upon service termination. With SCUR, it is also possible to include one or more intermediate interrogations for the CTF in order to report currently-used units, and to reserve additional units if required. In both cases, the session state is maintained in both the CTF and OCF.

For both ECUR and SCUR, the online charging client implements the “CLIENT, SESSION BASED” state machine described in RFC 4006.

Credit Authorization with Direct Debiting

For direct debiting, Immediate Event Charging (IEC) is used. With IEC, a single transaction is created where the OCF deducts a specific amount from the user's account immediately after completing the credit authorization. After receiving the authorization, the CTF delivers services. This form of credit authorization is a one-time event in which no session state is maintained.

With IEC, the online charging client implements the "CLIENT, EVENT BASED" state machine described in IETF RFC 4006.

Determining Units and Rating

Unit determination refers to calculating the number of non-monetary units (service units, time, events) that can be assigned prior to delivering services. Unit rating refers to determining a price based on the non-monetary units calculated by the unit determination function.

It is possible for either the OCF or the CTF to handle unit determination and unit rating. The decision lies with the client application, which controls the selection of AVPs in the credit control request sent to the OCF.

Configuring the Ro Application

The `RoApplication` is packaged as a Diameter application similar to the Sh application used for managing profile data. The Ro Diameter application can be configured and enabled by editing the Diameter configuration file located in *DOMAIN_ROOT*`/config/custom/diameter.xml`, or by using the Diameter console extension.

The application init parameter `ocs.host` specifies the host identity of the OCF. The OCF host must also be configured in the peer table as part of the global Diameter configuration. Alternately, the init parameter `ocs.realm` can be used to specify more than one OCF host using realm-based routing. The corresponding realm definition must also exist in the global Diameter configuration.

Listing 9-1 shows a sample excerpt from `diameter.xml` that enables Ro with an OCF host name of "myocs.bea.com."

Listing 9-1 Sample Ro Application Configuration (diameter.xml)

```
<application>
  <application-id>4</application-id>
```

```
    <class-name>com.bea.wcp.diameter.charging.RoApplication</class-name>
    <param>
      <name>ocs.host</name>
      <value>myocs.bea.com</value>
    </param>
  </application>
```

Because the `RoApplication` is based on the Diameter Credit Control Application, its Diameter application id is 4.

Overview of the Online Charging API

WebLogic SIP Server provides an online charging API to enable any deployed application to act as a CTF and issue online charging events to an OCS via the Ro protocol. All online charging requests use the Diameter Credit-Control-Request (CCR) message. The CC-Request-Type AVP is used to indicate the type of charging used. In the charging API, the CC-Request-Type is represented by the `RequestType` class in package `com.bea.wcp.diameter.cc`. Table 9-1 shows the request types associated with different credit authorization models.

Table 9-1 Credit Control Request Types

Type	Description	RequestType Field in com.bea.wcp.diameter.cc.RequestType
IEC	Immediate Event Charging	`EVENT_REQUEST`
ECUR	Event Charging with Unit Reservation	`INITIAL` or `TERMINATION_REQUEST`
SCUR	Session Charging with Unit Reservation	`INITIAL`, `UPDATE`, or `TERMINATION_REQUEST`

For ECUR and SCUR, units are reserved prior to service delivery and committed upon service completion. Units are reserved with `INITIAL_REQUEST` and committed with a `TERMINATION_REQUEST`. For SCUR, units can also be updated with `UPDATE_REQUEST`.

The base diameter package, `com.bea.wcp.diameter`, contains classes to support the re-authorization requests used in Ro. The `com.bea.wcp.diameter.cc` package contains classes to support credit-control applications, including Ro applications.

`com.bea.wcp.diameter.charging` directly supports the Ro credit-control application. Listing 9-2 summarizes the classes of interest to Ro credit-control.

Table 9-2 Summary of Ro Classes

Class	Description	Package
`Charging`	Constant definitions	`com.bea.wcp.diameter.charging`
`RoApplication`	Online charging application	`com.bea.wcp.diameter.charging`
`RoSession`	Online charging session	`com.bea.wcp.diameter.charging`
`CCR`	Credit Control Request	`com.bea.wcp.diameter.cc`
`CCA`	Credit Control Answer	`com.bea.wcp.diameter.cc`
`ClientSession`	Credit control client session	`com.bea.wcp.diameter.cc`
`RequestType`	Credit-control request type	`com.bea.wcp.diameter.cc`
`RAR`	Re-Auth-Request message	`com.bea.wcp.diameter`
`RAA`	Re-Auth-Answer message	`com.bea.wcp.diameter`

Accessing the Ro Application

If the Ro application is deployed, then applications deployed on WebLogic SIP Server can obtain an instance of the application from the Diameter node (`com.bea.wcp.diameter.Node` class). Listing 9-2 shows the sample Servlet code used to obtain the Diameter `Node` and access the Ro application.

Listing 9-2 Accessing the Ro Application

```
private RoApplication roApp;

void init(ServletConfig conf) {

    ServletContext ctx = conf.getServletContext();

    Node node = (Node) ctx.getParameter("com.bea.wcp.diameter.Node");

    roApp = node.getApplication(Charging.RO_APPLICATION_ID);

  }
```

This code example would make `RoApplication` available to the Servlet as an instance variable. The instance of `RoApplication` is safe for use by multiple concurrent threads.

Implementing Session-Based Charging

The `RoApplication` can be used to create new sessions for session-based credit authorization. The `RoSession` class implements the appropriate state machine depending on the credit control type, either ECUR (Event-Based Charging with Unit Reservation) or SCUR (Session-based Charging with Unit Reservation). The `RoSession` class is also serializable, so it can be stored as a SIP session attribute. This allows the session to be restored when necessary to terminate the session or update credit authorization.

The example in Listing 9-3 creates a new `RoSession` for event-based charging, and sends a CCR request to start the first interrogation. The `RoSession` instance is saved so that it can be terminated later, after the service has finished.

Note that the `RoSession` class automatically handles creating session IDs; the application is not required to set the session ID.

Listing 9-3 Creating and Using a RoSession

```
RoSession session = roApp.createSession();
CCR ccr = session.createCCR(RequestType.INITIAL);
CCA cca = ccr.sendAndWait();
sipAppSession.setAttribute("RoSession", session);
...
```

Handling Re-Auth-Request Messages

The OCS may initiate credit re-authorization by issuing a Re-Auth-Request (RAR) to the CTF. The application can register a session listener for handling this type of request. Upon receiving a RAR, the Diameter subsystem invoke the session listener on the applications corresponding `RoSession` object. The application should then respond to the OCS with an appropriate RAA message and initiate credit re-authorization to the CTF by sending a CCR with the CC-Request-Type AVP set to the value UPDATE_REQUEST, as described in section 5.5 of RFC 4006.

A session listener must implement the `SessionListener` interface and be serializable, or it must be an instance of `SipServlet`. A Servlet can register a listener as follows:

```
RoSession session = roApp.createSession();
session.addListener(new SessionListener() {
  public void rcvMessage(Message msg) {
    System.out.println("Got message: id = " msg.getSession().getId());
  }
});
```

Listing 9-4 shows sample `rcvMessage()` code for processing a Re-Auth-Request.

Listing 9-4 Managing a Re-Auth-Request

```
RoSession session = roApp.createSession();
session.addListener(new SessionListener() {
public void rcvMessage(Message msg) {
  if (req.getCommand() != Command.RE_AUTH_REQUEST) return;
  RoSession session = (RoSession) req.getSession();
  Answer ans = req.createAnswer();
  ans.setResultCode(ResultCode.LIMITED_SUCCESS); // Per RFC 4006 5.5
  ans.send();
  CCR ccr = session.createCCR(Ro.UPDATE_REQUEST);
  ... // Set CCR AVPs according to requested credit re-authorization
  ccr.send();
  CCA cca = (CCA) ccr.waitForAnswer();
}
```

In Listing 9-4, upon receiving the Re-Auth-Request the application sends an RAA with the result code DIAMETER_LIMITED_SUCCESS to indicate to the OCS that an additional CCR request is required in order to complete the procedure. The CCR is then sent to initiate credit re-authorization.

Note: Because the Diameter subsystem locks the call state before delivering the request to the corresponding `RoSession`, the call state remains locked while the handler processes the request.

Sending Credit-Control-Request Messages

The CCR class represents a Diameter Credit-Control-Request message, and can be used to send credit control requests to the OCF. For both ECUR (Event-Based Charging with Unit Reservation) and SCUR (Session-Based Charging with Unit Reservation), an instance of `RoSession` is used to create new CCR requests. You can also use `RoApplication` directly to create CCR messages for IEC (Immediate Event Charging). Listing 9-5 shows an example of how to create and send a CCR.

Listing 9-5 Creating and Sending a CCR

```
CCR ccr = session.createCCR(RequestType.INITIAL);
ccr.setServiceContextId("sample_id");
CCA cca = ccr.sendAndWait();
```

Once a CCR request is created, you can set whatever application- or service-specific AVPs that are required before sending the request using the `addAvp()` method. Because some of the same AVPs will need to be included in each new request for the session, it is also possible to set these AVPs on the session itself. Listing 9-6 shows a sample that sets:

- Subscription-Id to identify the user for the session
- Service-Identifier to indicate the service requested, and
- Requested-Service-Unit to specify the units requested.

A custom AVP is also added directly to the CCR request.

Listing 9-6 Setting AVPs in the CCR

```
session.setSubscriptionId(...);
session.setServiceIdentifier(...);
CCR ccr = session.createCCR(RequestType.INITIAL);
```

```
ccr.setRequestedServiceUnit(...);
ccr.addAvp(CUSTOM_MESSAGE, "This is a test");
ccr.send();
```

In this case, the same Subscription-Id and Service-Identifier are added to every new request for the session. The custom AVP "Custom-Message" is added to the message before it is sent out.

Handling Failures

Applications can examine the Result-Code AVP in CCA error responses from the OCF to detect the cause of a failure and take an appropriate action. Locally-generated errors, such as an unavailable peer or invalid route specification, cause the request send method to throw an `IOException` to with a detailed message indicating the nature of the failure.

Applications can also use the Diameter Timer Tx value for determining when the OCF fails to respond to a credit authorization request. Timer Tx has a default value of 10 seconds, but can be overridden using the `tx.timer` init parameter in the `RoApplication` configuration. Timer Tx starts when a CCR is sent to the OCF. The timer resets after the corresponding CCA is received.

If Tx expires before a corresponding CCA arrives, any call to `waitForAnswer` immediately returns null to indicate that the request has timed out. An application can then take action according to the value of the Credit-Control-Failure-Handling (CCFH) AVP in the request. See section 5.7, "Failure Procedures" in RFC 4006 for more details.

Listing 9-7 terminates the credit control session if timer Tx expires before receiving the CCA. If the CCA is received later by the Diameter subsystem, the message is ignored because the session longer exists.

Listing 9-7 Checking for Timer Tx Expiry

```
CCR ccr = session.createCCR(RequestType.INITIAL);
ccr.setCreditControlFailureHandling(RequestType.TERMINATION);
ccr.send();
CCA cca = ccr.waitForAnswer();
if (cca == null) {
  session.terminate();
}
```

CHAPTER 10

Developing Custom Profile Providers

The following sections describe how to use the profile service API to develop custom profile providers:

- "Overview of the Profile Service API" on page 10-1
- "Implementing Profile API Methods" on page 10-3
- "Configuring and Packaging Profile Providers" on page 10-3
- "Configuring Profile Providers Using the Administration Console" on page 10-6

Overview of the Profile Service API

WebLogic SIP Server includes a profile service API, `com.bea.wcp.profile.API`, that can be used to create profile provider implementations. A profile provider performs the work of accessing XML documents from a data repository using a defined protocol. Deployed SIP Servlets and other applications need not understand the underlying protocol or the data repository in which the document is stored; they simply reference profile data using a custom URL, and WebLogic SIP Server delegates the request processing to the correct profile provider.

The provider performs the necessary protocol operations for manipulating the document. All providers work with documents in XML DOM format, so client code can work with many different types of profile data in a common way.

You can also use the profile service API to create a custom provider for retrieving document schemas using another protocol. For example, a profile provider could be created to retrieve subscription data from an LDAP store or RDBMS.

Note: The Diameter Sh application also accesses profile data from a Home Subscriber Server using the Sh protocol. Profile. Although applications access this profile data using a simple URL, the Diameter applications are implemented using the Diameter base protocol implementation rather than the profile provider API.

Figure 10-1 Profile Service API and Provider Implementation

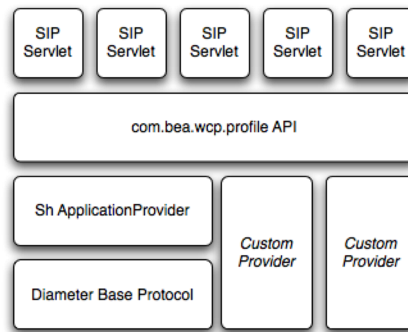

Each profile provider implemented using the API may enable the following operations against profile data:

- Creating new documents.
- Querying and updating existing documents.
- Deleting documents.
- Managing subscriptions for receiving notifications of profile document changes.

Clients that want to use a profile provider obtain a profile service instance via a Servlet context attribute. They then construct an appropriate URL and use that URL with one of the available profile API methods to work with profile data. The contents of the URL, combined with the configuration of profile providers, determines the provider implementation that WebLogic SIP Server to process the client's requests. "Using the Diameter Sh Interface Application" on page 7-1 provides additional examples about accessing a profile provider from within a SIP Servlet.

The sections that follow describe how to implement the profile service API interfaces in a custom profile provider.

Implementing Profile API Methods

A custom profile providers is implemented as a shared J2EE library (typically a simple JAR file) deployed to the engine tier cluster. The provider JAR file should include, at minimum, a class that implements com.bea.wcp.profile.ProfileServiceSpi. This interface inherits methods from `com.bea.wcp.profile.ProfileService` and defines new methods that are called during provider registration and unregistration.

In addition to the provider implementation, you must implement the com.bea.wcp.profile.ProfileSubscription interface if your provider supports subscription-based notification of profile data updates. A `ProfileSubscription` is returned to the client subscriber when the profile document is modified.

The WebLogic SIP Server JavaDoc describes each method of the profile service API in detail. Also keep in mind the following notes and best practices when implementing the profile service interfaces:

- The `putDocument`, `getDocument`, and `deleteDocument` methods each have two distinct method signatures. The basic version of a method passes only the document selector on which to operate. The alternate method signature also passes the address of the sender of the request for protocols that require explicit information about the requestor.
- The `subscribe` method has multiple method signatures to allow passing the sender's address, as well as for supporting time-based subscriptions.
- If you do not want to implement a method in `com.bea.wcp.profile.ProfileServiceSpi`, include a "no-op" method implementation that throws the OperationNotSupportedException.

com.bea.wcp.profile.ProfileServiceSpi defines provider methods that are called during registration and unregistration. Providers can create connections to data stores or perform any required initializing in the `register` method. The `register` method also supplies a `ProviderBean` instance, which includes any context parameters configured in the provider's configuration elements in `profile.xml`.

Providers should release any backing store connections, and clean up any state that they maintain, in the `unregister` method.

Configuring and Packaging Profile Providers

Providers must be deployed as a shared J2EE library, because all other deployed applications must be able to access the implementation. Creating Shared J2EE Libraries and Optional

Packages in the WebLogic Server 9.2 documentation describes how to assemble J2EE libraries. For most profile providers, you can simply package the implementation classes in a JAR file. Then register the library with WebLogic SIP Server using the instructions in Install a J2EE Library in the WebLogic Server 9.2 documentation.

After installing the provider as a library, you must also identify the provider class as a provider in a `profile.xml` file. The `name` element uniquely identifies a provider configuration, and the `class` element identifies the Java class that implements the profile service API interfaces. One or more context parameters can also be defined for the provider, which are delivered to the implementation class in the `register` method. For example, context parameters might be used to identify backing stores to use for retrieving profile data.

Listing 10-1 shows a sample configuration for a provider that accesses data using XCAP.

Listing 10-1 Provider Mapping in profile.xml

```
<profile-service xmlns="http://www.bea.com/ns/wlcp/wlss/profile/300"
                 xmlns:sec="http://www.bea.com/ns/weblogic/90/security"
                 xmlns:xsi="http://www.w3.org/2001/XMLSchema=instance"
                 xmlns:wls="http;//www.bea.com/ns/weblogic/90/security/wls
">
 <mapping>
   <map-by>provider-name</map-by>
 </mapping>
 <provider>
    <name>xcap</name>

<provider-class>com.mycompany.profile.XcapProfileProvider</provider-class>
    <param>
       <name>server</name>
       <value>example.com</name>
    </param>
```

```
    ...
  </provider>
</profile-service>
```

Mapping Profile Requests to Profile Providers

When an application makes a request using the profile API, WebLogic SIP Server must find a corresponding provider to process the request. By default, WebLogic SIP Server maps the prefix of the requested URL to a provider `name` element defined in `profile.xml`. For example, with the basic configuration shown in Listing 10-1, WebLogic SIP Server would map profile API requests beginning with `xcap://` to the provider class `com.mycompany.profile.XcapProfileProvider`.

Alternately, you can define a `mapping` entry in `profile.xml` that lists the prefixes corresponding to each named provider. Listing 10-2 shows a mapping with two alternate prefixes.

Listing 10-2 Mapping a Provider to Multiple Prefixes

```
...
<mapping>
   <map-by>prefix</map-by>
      <provider>
          <provider-name>xcap</provider-name>
          <doc-prefix>sip</doc-prefix>
          <doc-prefix>subscribe</doc-prefix>
      </provider>
   <by-prefix>
<mapping>
...
```

If the explicit mapping capabilities of `profile.xml` are insufficient, you can create a custom mapping class that implements the `com.bea.wcp.profile.ProfileRouter` interface, and then identify that class in the `map-by-router` element. Listing 10-3 shows an example configuration.

Listing 10-3 Using a Custom Mapping Class

```
...
<mapping>
   <map-by-router>
      <class>com.bea.wcp.profile.ExampleRouter</class>
   </map-by-router>
</mapping>
...
```

Configuring Profile Providers Using the Administration Console

You can optionally use the Administration Console to create or modify a `profile.xml` file. To do so, you must enable the profile provider console extension in the `config.xml` file for your domain.

Listing 10-4 Enabling the Profile Service Resource in config.xml

```
...
<custom-resource>
    <name>ProfileService</name>
    <target>AdminServer</target>
    <descriptor-file-name>custom/profile.xml</descriptor-file-name>

<resource-class>com.bea.wcp.profile.descriptor.resource.ProfileServiceReso
urce</resource-class>

<descriptor-bean-class>com.bea.wcp.profile.descriptor.beans.ProfileService
Bean</descriptor-bean-class>
  </custom-resource>
```

```
</domain>
```

The profile provider extension appears under the SipServer node in the left pane of the console, and enables you to configure new provider classes and mapping behavior.

CHAPTER 11

Using Content Indirection in SIP Servlets

The following sections describe how to develop SIP Servlets that work with indirect content specified in the SIP message body:

- "Overview of Content Indirection" on page 11-1
- "Using the Content Indirection API" on page 11-2

Overview of Content Indirection

Data provided by the body of a SIP message can be included either directly in the SIP message body, or indirectly by specifying an HTTP URL and metadata that describes the URL content. Indirectly specifying the content of the message body is used primarily in the following scenarios:

- When the message bodies include large volumes of data. In this case, content indirection can be used to transfer the data outside of the SIP network (using a separate connection or protocol).
- For bandwidth-limited applications. In this case, content indirection provides enough metadata for the application to determine whether or not it should retrieve the message body (potentially degrading performance or response time).

WebLogic SIP Server provides a simple API that you can use to work with indirect content specified in SIP messages.

Using the Content Indirection API

The content indirection API provided by WebLogic SIP Server helps you quickly determine if a SIP message uses content indirection, and to easily retrieve all metadata associated with the indirect content. The basic API consists of a utility class, `com.bea.wcp.sip.engine.server.ContentIndirectionUtil`, and an interface for accessing content metadata, `com.bea.wcp.sip.engine.server.ICParsedData`.

SIP Servlets can use the utility class to identify SIP messages having indirect content, and to retrieve an `ICParsedData` object representing the content metadata. The `ICParsedData` object has simple "getter" methods that return metadata attributes.

Additional Information

Complete details about content indirection are available in a draft document: http://www.ietf.org/internet-drafts/draft-ietf-sip-content-indirect-mech-05.txt.

See also the WebLogic SIP Server JavaDoc for additional documentation about the content indirection API.

CHAPTER 12

Securing SIP Servlet Resources

The following sections describe how to apply security constraints to SIP Servlet resources when deploying to WebLogic SIP Server:

- "Overview of SIP Servlet Security" on page 12-1
- "WebLogic SIP Server Role Mapping Features" on page 12-2
- "Using Implicit Role Assignment" on page 12-3
- "Assigning Roles Using security-role-assignment" on page 12-4
- "Assigning run-as Roles" on page 12-7
- "Role Assignment Precedence for SIP Servlet Roles" on page 12-8
- "Debugging Security Features" on page 12-9
- "weblogic.xml Deployment Descriptor Reference" on page 12-9

Overview of SIP Servlet Security

The SIP Servlet API specification defines a set of deployment descriptor elements that can be used for providing declarative and programmatic security for SIP Servlets. The primary method for declaring security constraints is to define one or more `security-constraint` elements in the `sip.xml` deployment descriptor. The `security-constraint` element defines the actual resources in the SIP Servlet, defined in `resource-collection elements`, that are to be protected. `security-constraint` also identifies the role names that are authorized to access the

resources. All role names used in the security-constraint are defined elsewhere in sip.xml in a security-role element.

SIP Servlets can also programmatically refer to a role name within the Servlet code, and then map the hard-coded role name to an alternate role in the sip.xml security-role-ref element during deployment. Roles must be defined elsewhere in a security-role element before they can be mapped to a hard-coded name in the security-role-ref element.

The SIP Servlet specification also enables Servlets to propagate a security role to a called Enterprise JavaBean (EJB) using the run-as element. Once again, roles used in the run-as element must be defined in a separate security-role element in sip.xml.

Chapter 14 in the SIP Servlet API specification provides more details about the types of security available to SIP Servlets. SIP Servlet security features are similar to security features available with HTTP Servlets; you can find additional information about HTTP Servlet security by referring to these sections in the WebLogic Server 9.2 documentation:

- Securing Web Applications in *Programming WebLogic Security* provides an overview of declarative and programmatic security models for Servlets.
- EJB Security-Related Deployment Descriptors in *Securing Enterprise JavaBeans (EJBs)* describes all security-related deployment descriptor elements for EJBs, including the run-as element used for propagating roles to called EJBs.

See also the example sip.xml excerpt in Listing 12-1, “Declarative Security Constraints in sip.xml,” on page 12-4.

WebLogic SIP Server Role Mapping Features

When you deploy a SIP Servlet, security-role definitions that were created for declarative and programmatic security must be assigned to actual principals and/or roles available in the Servlet container. WebLogic SIP Server uses the security-role-assignment element in weblogic.xml to help you map security-role definitions to actual principals and roles. security-role-assignment provides two different ways to map security roles, depending on how much flexibility you require for changing role assignment at a later time:

- The security-role-assignment element can define the complete list of principal names and roles that map to roles defined in sip.xml. This method defines the role assignment at deployment time, but at the cost of flexibility; to add or remove principals from the role, you must edit weblogic.xml and redeploy the SIP Servlet.

- The `externally-defined` element in `security-role-assignment` enables you to assign principal names and roles to a `sip.xml` role at any time using the Administration Console. When using the `externally-defined` element, you can add or remove principals and roles to a `sip.xml` role without having to redeploy the SIP Servlet.

Two additional XML elements can be used for assigning roles to a `sip.xml` `run-as` element: `run-as-principal-name` and `run-as-role-assignment`. These role assignment elements take precedence over `security-role-assignment` elements if they are used, as described in "Assigning run-as Roles" on page 12-7.

Optionally, you can choose to specify no role mapping elements in `weblogic.xml` to use implicit role mapping, as described in "Using Implicit Role Assignment" on page 12-3.

The sections that follow describe WebLogic SIP Server role assignment in more detail.

Using Implicit Role Assignment

With implicit role assignment, WebLogic SIP Server assigns a `security-role` name in `sip.xml` to a role of the exact same name, which should be configured in the WebLogic SIP Server security realm. To use implicit role mapping, you omit the `security-role-assignment` element in `weblogic.xml`, as well as any `run-as-principal-name`, and `run-as-role-assignment` elements use for mapping `run-as` roles.

When no role mapping elements are available in `weblogic.xml`, WebLogic SIP Server implicitly maps `sip.xml` `security-role` elements to roles having the same name. Note that implicit role mapping takes place regardless of whether the role name defined in `sip.xml` is actually available in the security realm. WebLogic SIP Server display a warning message anytime it uses implicit role assignment. For example, if you use the "everyone" role in `sip.xml` but you do not explicitly assign the role in `weblogic.xml`, the server displays the warning:

```
<Webapp: ServletContext(id=id,name=application,context-path=/context), the
role: everyone defined in web.xml has not been mapped to principals in
security-role-assignment in weblogic.xml. Will use the rolename itself as
the principal-name.>
```

You can ignore the warning message if the corresponding role has been defined in the WebLogic SIP Server security realm. The message can be disabled by defining an explicit role mapping in `weblogic.xml`.

Use implicit role assignment if you want to hard-code your role mapping at deployment time to a known principal name.

Assigning Roles Using security-role-assignment

The `security-role-assignment` element in `weblogic.xml` enables you to assign roles either at deployment time or at any time using the Administration Console. The sections that follow describe each approach.

Important Requirement for WebLogic SIP Server 3.1

If you specify a `security-role-assignment` element in `weblogic.xml`, WebLogic SIP Server 3.1 requires that you also define a duplicate `security-role` element in a `web.xml` deployment descriptor. This requirement applies even if you are deploying a pure SIP Servlet, which would not normally require a `web.xml` deployment descriptor (generally reserved for HTTP Web Applications).

Note: If you specify a `security-role-assignment` in `weblogic.xml` but there is no corresponding `security-role` element in `web.xml`, WebLogic SIP Server generates the error message:

```
The security-role-assignment references an invalid security-role:
rolename
```

The server then implicitly maps the `security-role` defined in `sip.xml` to a role of the same name, as described in "Using Implicit Role Assignment" on page 12-3.

For example, Listing 12-1 shows a portion of a `sip.xml` deployment descriptor that defines a security constraint with the role, `roleadmin`. Listing 12-2 shows that a `security-role-assignment` element has been defined in `weblogic.xml` to assign principals and roles to `roleadmin`. In WebLogic SIP Server, this Servlet *must* be deployed with a `web.xml` deployment descriptor that also defines the `roleadmin` role, as shown in Listing 12-3.

If the `web.xml` contents were not available, WebLogic SIP Server would use implicit role assignment and assume that the `roleadmin` role was defined in the security realm; the principals and roles assigned in `weblogic.xml` would be ignored.

Listing 12-1 Declarative Security Constraints in sip.xml

```
...
  <security-constraint>
      <resource-collection>
      <resource-name>RegisterRequests</resource-name>
```

```
      <servlet-name>registrar</servlet-name>
    </resource-collection>
    <auth-constraint>
      <role-name>roleadmin</role-name>
    </auth-constraint>
  </security-constraint>

  <security-role>
    <role-name>roleadmin</role-name>
  </security-role>
...
```

Listing 12-2 Example security-role-assignment in weblogic.xml

```
<weblogic-web-app>
  <security-role-assignment>
       <role-name>roleadmin</role-name>
       <principal-name>Tanya</principal-name>
       <principal-name>Fred</principal-name>
       <principal-name>system</principal-name>
  </security-role-assignment>
</weblogic-web-app>
```

Listing 12-3 Required security-role Element in web.xml

```
<!DOCTYPE web-app
    PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
    "http://java.sun.com/dtd/web-app_2_3.dtd">
<web-app>
```

```
  <security-role>
    <role-name>roleadmin</role-name>
  </security-role>
</web-app>
```

Assigning Roles at Deployment Time

A basic `security-role-assignment` element definition in `weblogic.xml` declares a mapping between a `security-role` defined in `sip.xml` and one or more principals or roles available in the WebLogic SIP Server security realm. If the `security-role` is used in combination with the `run-as` element in `sip.xml`, WebLogic SIP Server assigns the first principal or role name specified in the `security-role-assignment` to the `run-as` role.

Listing 12-2, “Example security-role-assignment in weblogic.xml,” on page 12-5 shows an example `security-role-assignment` element. This example assigns three users to the `roleadmin` role defined in Listing 12-1, “Declarative Security Constraints in sip.xml,” on page 12-4. To change the role assignment, you must edit the `weblogic.xml` descriptor and redeploy the SIP Servlet.

Dynamically Assigning Roles Using the Administration Console

The `externally-defined` element can be used in place of the `<principal-name>` element to indicate that you want the security roles defined in the `role-name` element of `sip.xml` to use mappings that you assign in the Administration Console. The `externally-defined` element gives you the flexibility of not having to specify a specific security role mapping for each security role at deployment time. Instead, you can use the Administration Console to specify and modify role assignments at anytime.

Additionally, because you may elect to use this element for some SIP Servlets and not others, it is not necessary to select the **ignore roles and polices from DD** option for the security realm. (You select this option in the **On Future Redeploys:** field on the **General** tab of the **Security->Realms->myrealm** control panel on the Administration Console.) Therefore, within the same security realm, deployment descriptors can be used to specify and modify security for some applications while the Administration Console can be used to specify and modify security for others.

Note: When specifying security role names, observe the following conventions and restrictions:

- The proper syntax for a security role name is as defined for an Nmtoken in the Extensible Markup Language (XML) recommendation available on the Web at: http://www.w3.org/TR/REC-xml#NT-Nmtoken.
- Do not use blank spaces, commas, hyphens, or any characters in this comma-separated list: \t, < >, #, |, &, ~, ?, ( ), { }.
- Security role names are case sensitive.
- The BEA suggested convention for security role names is that they be singular.

Listing 12-4 shows an example of using the `externally-defined` element with the `roleadmin` role defined in Listing 12-1, “Declarative Security Constraints in sip.xml,” on page 12-4. To assign existing principals and roles to the `roleadmin` role, the Administrator would use the WebLogic SIP Server Administration Console.

See Users, Groups, and Security Roles in the WebLogic Server 9.2 documentation for information about adding and modifying security roles using the Administration Console.

Listing 12-4 Example externally-defined Element in weblogic.xml

```
<weblogic-web-app>
     <security-role-assignment>
          <role-name>webuser</role-name>
          <externally-defined/>
     </security-role-assignment>
</weblogic-web-app>
```

Assigning run-as Roles

The `security-role-assignment` described in “Assigning Roles Using security-role-assignment” on page 12-4 can be also be used to map `run-as` roles defined in `sip.xml`. Note, however, that two additional elements in `weblogic.xml` take precedence over the `security-role-assignment` if they are present: `run-as-principal-name` and `run-as-role-assignment`.

`run-as-principal-name` specifies an existing principle in the security realm that is used for all `run-as` role assignments. When it is defined within the `servlet-descriptor` element of

weblogic.xml, run-as-principal-name takes precedence over any other role assignment elements for run-as roles.

run-as-role-assignment specifies an existing role or principal in the security realm that is used for all run-as role assignments, and is defined within the weblogic-web-app element.

See "weblogic.xml Deployment Descriptor Reference" on page 12-9 for more information about individual weblogic.xml descriptor elements. See also "Role Assignment Precedence for SIP Servlet Roles" on page 12-8 for a summary of the role mapping precedence for declarative and programmatic security as well as run-as role mapping.

Role Assignment Precedence for SIP Servlet Roles

WebLogic SIP Server provides several ways to map sip.xml roles to actual roles in the SIP Container during deployment. For declarative and programmatic security defined in sip.xml, the order of precedence for role assignment is:

1. If weblogic.xml assigns a sip.xml role in a security-role-assignment element, the security-role-assignment is used.

   Note: WebLogic SIP Server also requires a role definition in web.xml in order to use a security-role-assignment. See "Important Requirement for WebLogic SIP Server 3.1" on page 12-4.

2. If no security-role-assignment is available (or if the required web.xml role assignment is missing), implicit role assignment is used.

For run-as role assignment, the order of precedence for role assignment is:

1. If weblogic.xml assigns a sip.xml run-as role in a run-as-principal-name element defined within servlet-descriptor, the run-as-principal-name assignment is used.

   Note: WebLogic SIP Server also requires a role definition in web.xml in order to assign roles with run-as-principal-name. See "Important Requirement for WebLogic SIP Server 3.1" on page 12-4.

2. If weblogic.xml assigns a sip.xml run-as role in a run-as-role-assignment element, the run-as-role-assignment element is used.

   Note: WebLogic SIP Server also requires a role definition in web.xml in order to assign roles with run-as-role-assignment. See "Important Requirement for WebLogic SIP Server 3.1" on page 12-4.

3. If `weblogic.xml` assigns a `sip.xml run-as` role in a `security-role-assignment` element, the `security-role-assignment` is used.

   Note: WebLogic SIP Server also requires a role definition in `web.xml` in order to use a `security-role-assignment`. See "Important Requirement for WebLogic SIP Server 3.1" on page 12-4.

4. If no `security-role-assignment` is available (or if the required `web.xml` role assignment is missing), implicit role assignment is used.

Debugging Security Features

If you want to debug security features in SIP Servlets that you develop, specify the `-Dweblogic.Debug=wlss.Security startup` option when you start WebLogic SIP Server. Using this debug option causes WebLogic SIP Server to display additional security-related messages in the standard output.

weblogic.xml Deployment Descriptor Reference

The `weblogic.xml` DTD contains detailed information about each of the role mapping elements discussed in this section. See http://www.bea.com/servers/wls810/dtd/weblogic810-web-jar.dtd for the complete DTD. See also weblogic.xml Deployment Descriptor Elements in the WebLogic Server 9.2 documentation.

CHAPTER 13

Developing SIP Servlets Using Eclipse

The following sections describe how to use Eclipse to develop SIP Servlets for use with WebLogic SIP Server:

- "Overview" on page 13-1
- "Setting Up the Development Environment" on page 13-2
- "Building and Deploying the Project" on page 13-6
- "Debugging SIP Servlets" on page 13-6

Overview

This document provides detailed instructions for using the Eclipse IDE as a tool for developing and deploying SIP Servlets with WebLogic SIP Server. The full development environment requires the following components, which you must obtain and install before proceeding:

- WebLogic SIP Server 3.1
- JDK 1.4.2
- Ant (installed with WebLogic SIP Server 3.1)
- Eclipse version 3.1
- CVS client and server (required only for version control)

SIP Servlet Organization

Building a SIP Servlet produces a Web Archive (WAR file or directory) as an end product. A basic SIP Servlet WAR file contains the subdirectories and contents described in Figure 13-1.

Figure 13-1 SIP Servlet WAR Contents

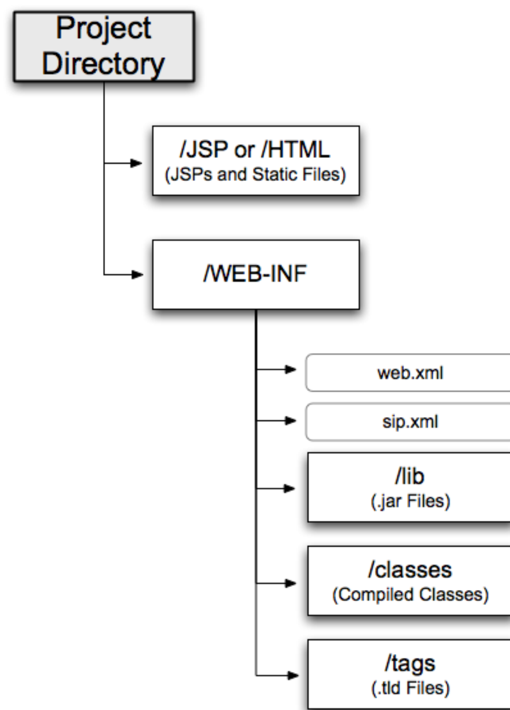

Setting Up the Development Environment

Follow these steps to set up the development environment for a new SIP Servlet project:

1. Create a new WebLogic SIP Server Domain.
2. Create a new Eclipse project.
3. Create an Ant build file.

The sections that follow describe each step in detail.

Creating a WebLogic SIP Server Domain

In order to deploy and test your SIP Servlet, you need access to a WebLogic SIP Server domain that you can reconfigure and restart as necessary. Follow the instructions in Create an Administrative Domain in the *Installation Guide* to create a new domain using the Configuration Wizard. When generating a new domain:

- Select Development Mode as the startup mode for the new domain.
- Select Sun SDK 1.4.2 as the SDK for the new domain.

Configure the Default Eclipse JVM

The latest versions of Eclipse use the version 1.5 JRE by default. Follow these steps to configure Eclipse to use the version 1.4.2 JRE installed with WebLogic SIP Server:

1. Start Eclipse.
2. Select Window->Preferences
3. Expand the Java category in the left pane, and select Installed JREs.
4. Click Add... to add the new JRE.
5. Enter a name to use for the new JRE in the JRE name field.
6. Click the Browse... button next to the JRE home directory field. Then navigate to the BEA_HOME/jdk142_08 directory and click OK.
7. Click OK to add the new JRE.
8. Select the check box next to the new JRE to make it the default.
9. Click OK to dismiss the preferences dialog.

Creating a New Eclipse Project

Follow these steps to create a new Eclipse project for your SIP Servlet development, adding the WebLogic SIP Server libraries required for building and deploying the application:

1. Start Eclipse.
2. Select File->New->Project...

3. Select Java Project and click Next.
4. Enter a name for your project in the Project Name field.
5. In the Location field, select Create project in workspace if you have not yet begun writing the SIP Servlet code. If you already have source code available in another location, Select Create project at external location and specify the directory. Click Next.
6. Click the Libraries tab and follow these steps to add required JARs to your project:
   a. Click Add External JARs...
   b. Use the JAR selection dialog to add the `BEA_HOME/sipserver31/server/lib/weblogic.jar` file to your project.
   c. Click Add External JARs... once again.
   d. Use the JAR selection dialog to add the `BEA_HOME/sipserver31/telco/auxlib/sipservlet.jar` file to your project.
   e. (Optional.) If your application needs to access WebLogic SIP Server MBeans using JMX, also use the JAR selection dialog to add `BEA_HOME/sipserver31/telco/lib/wcp_sip_core.jar` to your project.
7. Add any additional JAR files that you may require for your project.
8. Click Finish to create the new project. Eclipse displays your new project name in the Package Explorer.
9. Right-click on the name of your project and use the New->Folder command to recreate the directory structure shown in Figure 13-1, “SIP Servlet WAR Contents,” on page 13-2.

Creating an Ant Build File

Follow these steps to create an Ant build file that you can use for building and deploying your project:

1. Right-click on the name of your project in Eclipse, and select New->File
2. Enter the name `build.xml` and click Finish. Eclipse opens the empty file in a new window.
3. Copy the sample text from Listing 13-1, substituting your domain name and application name for *myDomain* and *myApplication*.

Listing 13-1 Ant Build File Contents

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<project default="all">
  <property environment="env"/>
  <property name="beahome" value="${env.BEA_HOME}"/>
  <target name="all" depends="compile,install"/>
  <target name="compile">
    <mkdir dir="WEB-INF/classes"/>
    <javac destdir="WEB-INF/classes" srcdir="src" debug="true"
debuglevel="lines,vars,source">
      <classpath>
        <pathelement path="${weblogic.jar}"/>
      </classpath>
    </javac>
  </target>
  <target name="install">
    <jar
destfile="${beahome}/user_projects/domains/myDomain/applications/myApplica
tion.war">
      <zipfileset dir="WEB-INF" prefix="WEB-INF"/>
      <zipfileset dir="WEB-INF" includes="*.html"/>
      <zipfileset dir="WEB-INF" includes="*.jsp"/>
    </jar>
  </target>
</project>
```

4. Close the `build.xml` file and save your changes.

5. Verify that the `build.xml` file is valid by selecting Window->Show View->Ant and dragging the `build.xml` file into the Ant view. Correct any problems before proceeding.
6. Right-click on the project name and select Properties.
7. Select the Builders property in the left column, and click New.
8. Select the Ant Build tool type and click OK to add an Ant builder.
9. In the Buildfile field, click Browse Workspace and select the `build.xml` file you created.
10. In the Base Directory field, click Browse Workspace and select the top-level directory for your project.
11. Click the JRE tab and choose Separate JRE in the Runtime JRE field. Use the drop-down list or the Installed JREs... button to select an installed version 1.4.2 JRE.
12. Click the Environment tab, and Click New. Enter a new name/value pair to define the `BEA_HOME` variable. The `BEA_HOME` variable must point to the home directory of the WebLogic SIP Server directory. For example:
    - Name: BEA_HOME
    - Value: c:\bea
13. Click OK to add the new Ant builder to the project.
14. De-select Java Builder in the builder list to remove the Java builder from the project.
15. Click OK to finish configuring Builders for the project.

Building and Deploying the Project

The `build.xml` file that you created compiles your code, packages the WAR, and copies the WAR file to the `/applications` subdirectory of your development domain. WebLogic SIP Server automatically deploys valid applications located in the `/applications` subdirectory.

Debugging SIP Servlets

In order to debug SIP Servlets, you must enable certain debug options when you start WebLogic SIP Server. Follow these steps to add the required debug options to the script used to start WebLogic SIP Server:

1. Use a text editor to open the `StartWebLogic.cmd` script for your development domain.

2. Beneath the line that reads:

   ```
   set JAVA_OPTIONS=
   ```

   Enter the following line:

   ```
   set DEBUG_OPTS=-Xdebug
   -Xrunjdwp:transport=dt_socket,address=9000,server=y,suspend=n
   ```

3. In the last line of the file, add the `%DEBUG_OPTS%` variable in the place indicated below:

   ```
   "%JAVA_HOME%\bin\java" %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %DEBUG_OPTS%

   -Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER%

   -Dweblogic.management.password=%WLS_PW%

   -Dweblogic.management.server=%ADMIN_URL%

   -Djava.security.policy="%WL_HOME%\server\lib\weblogic.policy"
   weblogic.Server
   ```

4. Save the file and use the script to restart WebLogic SIP Server.

CHAPTER 14

Enabling Message Logging

The following sections describe how to use WebLogic SIP Server message logging features on a development system:

- "Overview" on page 14-1
- "Enabling Message Logging" on page 14-2
- "Specifying Content Types for Unencrypted Logging" on page 14-5
- "Example Message Log Configuration and Output" on page 14-6
- "Configuring Log File Rotation" on page 14-8

Overview

Message logging records SIP and Diameter messages (both requests and responses) received by WebLogic SIP Server. You can use the message log in a development environment to check how external SIP requests and SIP responses are received. By outputting the distinguishable information of SIP dialogs such as Call-IDs from the application log, and extracting relevant SIP messages from the message log, you can also check SIP invocations from HTTP servlets and so forth.

WARNING: The message logging functionality logs *all* SIP requests and responses; do not enable this feature in a production system. In a production system, you can instead configure one or more logging Servlets, which enable you to specify additional criteria for determining which messages to log. See Logging SIP Requests and Responses in the *Operations Guide*.

When you enable message logging, WebLogic SIP Server records log records in the Managed Server log file associated with each engine tier server instance by default. You can optionally log the messages in a separate, dedicated log file, as described in "Configuring Log File Rotation" on page 14-8.

Enabling Message Logging

You enable and configure message logging by adding a `message-debug` element to the `sipserver.xml` configuration file. WebLogic SIP Server provides two different methods of configuring the information that is logged:

- Specify a predefined logging level (terse, basic, or full), or
- Identify the exact portions of the SIP message that you want to include in a log record, in a specified order

The sections that follow describe each method of configuring message logging functionality using elements in the `sipserver.xml` file. Note that you can also set these elements using the Administration Console, in the Configuration->Message Debug tab of the SIP Server console extension node.

Specifying a Predefined Logging Level

The optional `level` element in `message-debug` specifies a predefined collection of information to log for each SIP request and response. The following levels are supported:

- `terse`—Logs only the `domain` setting, logging Servlet name, logging `level`, and whether or not the message is an incoming message.
- `basic`—Logs the `terse` items plus the SIP message status, reason phrase, the type of response or request, the SIP method, the **From** header, and the **To** header.
- `full`—Logs the `basic` items plus all SIP message headers plus the timestamp, protocol, request URI, request type, response type, content type, and raw content.

Listing 14-1 shows a configuration entry that specifies the `full` logging level.

Listing 14-1 Sample Message Logging Level Configuration in sipserver.xml

```
<message-debug>
   <level>full</level>
```

```
</message-debug>
```

Customizing Log Records

WebLogic SIP Server also enables you to customize the exact content and order of each message log record. To configure a custom log record, you provide a `format` element that defines a log record `pattern` and one or more `tokens` to log in each record.

Note: If you specify a `format` element with a `<level>full</level>` element (or with the `level` element undefined) in `message-debug`, WebLogic SIP Server uses "full" message debugging and ignores the `format` entry. The `format` entry can be used in combination with either the "terse" or "basic" `message-debug` levels.

Table 14-1 describes the nested elements used in the `format` element.

Table 14-1 Nested format Elements

param-name	param-value Description
`pattern`	Specifies the pattern used to format a message log entry. The format is defined by specifying one or more integers, bracketed by "{" and "}". Each integer represents a `token` defined later in the `format` definition.
`token`	A string token that identifies a portion of the SIP message to include in a log record. `Table 14-2` provides a list of available string tokens. You can define multiple `token` elements as needed to customize your log records.

Table 14-2 describes the string `token` values used to specify information in a message log record:

Table 14-2 Available Tokens for Message Log Records

Token	Description	Example or Type
%call_id	The Call-ID header. It is blank when forwarding.	43543543
%content	The raw content.	Byte array
%content_length	The content length.	String value
%content_type	The content type.	String value
%cseq	The CSeq header. It is blank when forwarding.	INVITE 1

Table 14-2 Available Tokens for Message Log Records

Token	Description	Example or Type
%date	The date when the message was received. (“yyyy/MM/dd” format)	2004/05/16
%exception	The class name of the exception occurred when calling the AP. Detailed information is recorded to the run-time log.	NullPointerException
%from	The From header (all). It is blank when forwarding.	sip:foo@bea.com;tag=438943
%from_addr	The address portion of the From header.	foo@bea.com
%from_port	The port number portion of the From header.	7002
%from_tag	The tag parameter of the From header. It is blank when forwarding.	12345
%from_uri	The SIP URI part of the From header. It is blank when forwarding.	sip:foo@bea.com
%headers	A List of message headers stored in a 2-element array. The first element is the name of the header, while the second is a list of all values for the header.	List of headers
%io	Whether the message is incoming or not.	TRUE
%method	The name of the SIP method. It records the method name to invoke when forwarding.	INVITE
%msg	Summary Call ID	String value
%mtype	The type of receiving.	SIPREQ
%protocol	The protocol used.	UDP
%reason	The response reason.	OK
%req_uri	The request URI. This token is only available for the SIP request.	sip:foo@bea.com
%status	The response status.	200
%time	The time when the message was received. (“HH:mm:ss” format)	18:05:27

Table 14-2 Available Tokens for Message Log Records

Token	Description	Example or Type
%timestampmillis	Time stamp in milliseconds.	9295968296
%to	The To header (all). It is blank when forwarding.	sip:foo@bea.com;tag=438943
%to_addr	The address portion of the To header.	foo@bea.com
%to_port	The port number portion of the To header.	7002
%to_tag	The tag parameter of the To header. It is blank when forwarding.	12345
%to_uri	The SIP URI part of the To header. It is blank when forwarding.	sip:foo@bea.com

See "Example Message Log Configuration and Output" on page 14-6 for an example `sipserver.xml` file that defines a custom log record using two tokens.

Specifying Content Types for Unencrypted Logging

By default WebLogic SIP Server uses String format (UTF-8 encoding) to log the content of SIP messages having a text or application/sdp Content-Type value. For all other Content-Type values, WebLogic SIP Server attempts to log the message content using the character set specified in the `charset` parameter of the message, if one is specified. If no `charset` parameter is specified, or if the `charset` value is invalid or unsupported, WebLogic SIP Server uses Base-64 encoding to encrypt the message content before logging the message.

If you want to avoid encrypting the content of messages under these circumstances, specify a list of String-representable Content-Type values using the `string-rep` element in `sipserver.xml`. The `string-rep` element can contain one or more `content-type` elements to match. If a logged message matches one of the configured `content-type` elements, WebLogic SIP Server logs the content in String format using UTF-8 encoding, regardless of whether or not a `charset` parameter is included.

Note: You do not need to specify text/* or application/sdp content types as these are logged in String format by default.

Listing 14-2 shows a sample `message-debug` configuration that logs String content for three additional Content-Type values, in addition to text/* and application/sdp content.

Listing 14-2 Logging String Content for Additional Content Types

```
<message-debug>
  <level>full</level>
  <string-rep>
    <content-type>application/msml+xml</content-type>
    <content-type>application/media_control+xml</content-type>
    <content-type>application/media_control</content-type>
  </string-rep>
</message-debug>
```

Example Message Log Configuration and Output

Listing 14-3 shows a sample message log configuration in `sipserver.xml`. Listing 14-4, "Sample Message Log Output," on page 14-6 shows sample output from the Managed Server log file.

Listing 14-3 Sample Message Log Configuration in sipserver.xml

```
<message-debug>
   <format>
      <pattern>{0} {1}</pattern>
      <token>%headers</token>
      <token>%content</token>
   </format>
</message-debug>
```

Listing 14-4 Sample Message Log Output

```
####<Aug 10, 2005 7:12:08 PM PDT> <Info> <WLSS.Trace> <jiri.bea.com>
<myserver> <ExecuteThread: '11' for queue: 'sip.transport.Default'> <<WLS
```

```
Kernel>> <> <BEA- 331802> <SIP Tracer: logger Message: To: sut
<sip:invite@10.32.5.230:5060> <mailto:sip:invite@10.32.5.230:5060>

Content-Length: 136

Contact: user:user@10.32.5.230:5061

CSeq: 1 INVITE

Call-ID: 59.3170.10.32.5.230@user.call.id

From: user <sip:user@10.32.5.230:5061> <mailto:sip:user@10.32.5.230:5061>
;tag=59

Via: SIP/2.0/UDP 10.32.5.230:5061

Content-Type: application/sdp

Subject: Performance Test

Max-Forwards: 70

 v=0

o=user1 53655765 2353687637 IN IP4 127.0.0.1

s=-

c=IN IP4       127.0.0.1

t=0 0

m=audio 10000 RTP/AVP 0

a=rtpmap:0 PCMU/8000

>

####<Aug 10, 2005 7:12:08 PM PDT> <Info> <WLSS.Trace> <jiri.bea.com>
<myserver> <ExecuteThread: '11' for queue: 'sip.transport.Default'> <<WLS
Kernel>> <> <BEA- 331802> <SIP Tracer: logger Message: To: sut
<sip:invite@10.32.5.230:5060> <mailto:sip:invite@10.32.5.230:5060>

Content-Length: 0

CSeq: 1 INVITE

Call-ID: 59.3170.10.32.5.230@user.call.id

Via: SIP/2.0/UDP 10.32.5.230:5061
```

```
From: user <sip:user@10.32.5.230:5061> <mailto:sip:user@10.32.5.230:5061>
;tag=59

Server: BEA WebLogic SIP Server 3.1.0.0

 >
```

Configuring Log File Rotation

Message log entries for SIP and Diameter messages are stored in the main WebLogic SIP Server log file by default. You can optionally store the messages in a dedicated log file. Using a separate file makes it easier to locate message logs, and also enables you to use WebLogic SIP Server's log rotation features to better manage logged data.

Log rotation is configured using several elements nested within the main `message-debug` element in `sipserver.xml`. As with the other XML elements described in this section, you can also configure values using the Configuration->Message Debug tab of the SIP Server Administration Console extension.

Table 14-3 describes each element. Note that a server restart is necessary in order to initiate independent logging and log rotation.

Table 14-3 XML Elements for Configuring Log Rotation

Element	Description
`logging-enabled`	Determines whether a separate log file is used to store message debug log messages. By default, this element is set to false and messages are logged in the general WebLogic SIP Server log file.
`file-min-size`	Configures the minimum size, in kilobytes, after which the server automatically rotate log messages into another file. This value is used when the `rotation-type` element is set to `bySize`.
`log-filename`	Defines the name of the log file for storing messages. By default, the log files are stored under *`domain_home`*`/servers/`*`server_name`*`/logs`.
`rotation-type`	Configures the criterion for moving older log messages to a different file. This element may have one of the following values: • `bySize`—This default setting rotates log messages based on the specified `file-min-size`. • `byTime`—This setting rotates log messages based on the specified `rotation-time`. • `none`—Disables log rotation.

Table 14-3 XML Elements for Configuring Log Rotation

`number-of-files-limited`	Specifies whether or not the server places a limit on the total number of log files stored after a log rotation. By default, this element is set to false.
`file-count`	Configures the maximum number of log files to keep when `number-of-files-limited` is set to true.
`rotate-log-on-startup`	Determines whether the server should rotate the log file at server startup time.
`log-file-rotation-dir`	Configures a directory in which to store rotated log files. By default, rotated log files are stored in the same directory as the active log file.
`rotation-time`	Configures a start time for log rotation when using the `byTime` log rotation criterion.
`file-time-span`	Specifies the interval, in hours, after which the log file is rotated. This value is used when the `rotation-type` element is set to `byTime`.
`date-format-pattern`	Specifies the pattern to use for rending dates in log file entries. The value of this element must conform to the `java.text.SimpleDateFormat` class.

Listing 14-5 shows a sample `message-debug` configuration using log rotation.

Listing 14-5 Sample Log Rotation Configuration

```
<?xml version='1.0' encoding='UTF-8'?>
<sip-server xmlns="http://www.bea.com/ns/wlcp/wlss/300"
xmlns:sec="http://www.bea.com/ns/weblogic/90/security"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:wls="http://www.bea.com/ns/weblogic/90/security/wls">
  <message-debug>
    <logging-enabled>true</logging-enabled>
    <file-min-size>500</file-min-size>
    <log-filename>sip-messages.log</log-filename>
    <rotation-type>byTime</rotation-type>
    <number-of-files-limited>true</number-of-files-limited>
```

```
    <file-count>5</file-count>
    <rotate-log-on-startup>false</rotate-log-on-startup>
    <log-file-rotation-dir>old_logs</log-file-rotation-dir>
    <rotation-time>00:00</rotation-time>
    <file-time-span>20</file-time-span>
    <date-format-pattern>MMM d, yyyy h:mm a z</date-format-pattern>
  </message-debug>
</sip-server>
```

CHAPTER 15

Generating SNMP Traps from Application Code

The following sections describe how to use the WebLogic SIP Server `SipServletSnmpTrapRuntimeMBean` to generate SNMP traps from within a SIP Servlet:

- "Overview" on page 15-1
- "Requirement for Accessing SipServletSnmpTrapRuntimeMBean" on page 15-2
- "Obtaining a Reference to SipServletSnmpTrapRuntimeMBean" on page 15-3
- "Generating a SNMP Trap" on page 15-5

See Configuring SNMP in the *Operations Guide* for information about configuring SNMP in a WebLogic SIP Server domain.

Overview

WebLogic SIP Server includes a runtime MBean, `SipServletSnmpTrapRuntimeMBean,` that enables applications to easily generate SNMP traps. The WebLogic SIP Server MIB contains seven new OIDs that are reserved for traps generated by an application. Each OID corresponds to a severity level that the application can assign to a trap, in order from the least severe to the most severe:

- Info
- Warning
- Error

- Notice
- Critical
- Alert
- Emergency

To generate a trap, an application simply obtains an instance of the `SipServletSnmpTrapRuntimeMBean` and then executes a method that corresponds to the desired trap severity level (`sendInfoTrap()`, `sendWarningTrap()`, `sendErrorTrap()`, `sendNoticeTrap()`, `sendCriticalTrap()`, `sendAlertTrap()`, and `sendEmergencyTrap()`). Each method takes a single parameter—the String value of the trap message to generate.

For each SNMP trap generated in this manner, WebLogic SIP Server also automatically transmits the Servlet name, application name, and WebLogic SIP Server instance name associated with the calling Servlet.

Requirement for Accessing SipServletSnmpTrapRuntimeMBean

In order to obtain a `SipServletSnmpTrapRuntimeMBean`, the calling SIP Servlet must be able to perform MBean lookups from the Servlet context. To enable this functionality, you must assign a WebLogic SIP Server administrator `role-name` entry to the `security-role` and `run-as` role elements in the `sip.xml` deployment descriptor. Listing 15-1 shows a sample `sip.xml` file with the required role elements highlighted.

Listing 15-1 Sample Role Requirement in sip.xml

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE sip-app
   PUBLIC "-//Java Community Process//DTD SIP Application 1.0//EN"
   "http://www.jcp.org/dtd/sip-app_1_0.dtd">
<sip-app>
  <display-name>My SIP Servlet</display-name>
  <distributable/>
```

```
  <servlet>
    <servlet-name>myservlet</servlet-name>
    <servlet-class>com.mycompany.MyServlet</servlet-class>
    <run-as>
      <role-name>weblogic</role-name>
    </run-as>
  </servlet>
  <servlet-mapping>
    <servlet-name>myservlet</servlet-name>
    <pattern>
      <equal>
       <var>request.method</var>
       <value>INVITE</value>
      </equal>
    </pattern>
  </servlet-mapping>
  <security-role>
    <role-name>weblogic</role-name>
  </security-role>
</sip-app>
```

Obtaining a Reference to SipServletSnmpTrapRuntimeMBean

Any SIP Servlet that generates SNMP traps must first obtain a reference to the `SipServletSnmpTrapRuntimeMBean`. Listing 15-2 shows the sample code for a method to obtain the MBean.

Listing 15-2 Sample Method for Accessing SipServletSnmpTrapRuntimeMBean

```
public SipServletSnmpTrapRuntimeMBean getServletSnmpTrapRuntimeMBean() {
    MBeanHome localHomeB = null;
    SipServletSnmpTrapRuntimeMBean ssTrapMB = null;

    try
    {
      Context ctx = new InitialContext();
      localHomeB = (MBeanHome)ctx.lookup(MBeanHome.LOCAL_JNDI_NAME);
      ctx.close();
    } catch (NamingException ne){
      ne.printStackTrace();
    }

    Set set = localHomeB.getMBeansByType("SipServletSnmpTrapRuntime");
    if (set == null || set.isEmpty()) {
      try {
        throw new ServletException("Unable to lookup type
'SipServletSnmpTrapRuntime'");
      } catch (ServletException e) {
        e.printStackTrace();
      }
    }
    ssTrapMB = (SipServletSnmpTrapRuntimeMBean) set.iterator().next();
    return ssTrapMB;
}
```

Generating a SNMP Trap

In combination with the method shown in Listing 15-2, Listing 15-3 demonstrates how a SIP Servlet would use the MBean instance to generate an SNMP trap in response to a SIP INVITE.

Listing 15-3 Generating a SNMP Trap

```
public class MyServlet extends SipServlet {
  private SipServletSnmpTrapRuntimeMBean sipServletSnmpTrapMb = null;

  public MyServlet () {
  }

  public void init (ServletConfig sc) throws ServletException {
    super.init (sc);
    sipServletSnmpTrapMb = getServletSnmpTrapRuntimeMBean();
  }

  protected void doInvite(SipServletRequest req) throws IOException {
    sipServletSnmpTrapMb.sendInfoTrap("Rx Invite from " +
req.getRemoteAddr() + "with call id" + req.getCallId());
  }
}
```