

Oracle® WebLogic Server SIP Container

Administrator's Guide

11g Release 1 (11.1.1)

E15459-02

January 2011

Oracle WebLogic Server SIP Container Administrator's Guide, 11g Release 1 (11.1.1)

E15459-02

Copyright © 2006, 2011, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	xi
Audience	xi
Documentation Accessibility	xi
Related Documents	xii
Conventions	xii
1 Configuring Oracle WebLogic SIP Container	
1.1 Oracle WebLogic SIP Container	1-1
1.1.1 WebLogic Server 10.3 Platform Supports SIP and Converged Applications	1-1
1.2 Shared Configuration Tasks	1-2
1.2.1 Shared Configuration Tasks for SIP Container and WebLogic Server	1-2
1.2.2 Oracle WebLogic Server SIP Container Configuration Overview	1-2
1.2.2.1 Diameter Configuration.....	1-3
1.2.3 Methods and Tools for Performing Configuration Tasks.....	1-4
1.2.3.1 Administration Console	1-4
1.2.3.2 WebLogic Scripting Tool (WLST)	1-5
1.2.3.3 Additional Configuration Methods	1-5
1.2.3.3.1 Editing Configuration Files.....	1-5
1.2.3.3.2 Custom JMX Applications	1-5
1.2.3.3.3 Setting Log Levels	1-5
1.2.4 Starting and Stopping Servers	1-5
1.2.5 Administration Server Best Practices.....	1-6
1.2.6 Common Configuration Tasks.....	1-7
2 Configuring SIP Servlet Container Properties	
2.1 Overview of SIP Container Configuration.....	2-1
2.2 Using the Administration Console to Configure Container Properties.....	2-1
2.2.1 Locking and Persisting the Configuration	2-2
2.3 Configuring Container Properties Using WLST (JMX)	2-3
2.3.1 Managing Configuration Locks.....	2-3
2.3.2 Locating the Oracle WebLogic Server SIP Container MBeans.....	2-4
2.4 WLST Configuration	2-4
2.4.1 Invoking WLST	2-4
2.4.2 Creating and Deleting MBeans.....	2-5
2.5 Configuring Timer Processing.....	2-5

2.5.1	Configuring Timer Affinity (Optional).....	2-5
2.5.2	Configuring NTP for Accurate SIP Timers	2-6

3 Managing Network Resources

3.1	Overview of Network Configuration	3-1
3.1.1	IPv4 and IPv6	3-2
3.2	Configuring Load Balancer Addresses	3-2
3.2.1	Multiple Load Balancers and DNS Load Balancers.....	3-2
3.3	Enabling Domain Name Service (DNS) Support	3-3
3.4	Configuring Network Channels for SIP or SIPS.....	3-4
3.4.1	Reconfiguring an Existing Channel	3-4
3.4.2	Creating a New SIP or SIPS Channel.....	3-4
3.4.3	Configuring Custom Timeout, MTU, and Other Properties	3-5
3.4.4	Configuring SIP Channels for DNS Machines	3-7
3.5	Configuring TCP and TLS Channels for Diameter Support.....	3-7
3.6	Configuring Engine Servers to Listen on Any IP Interface	3-7
3.7	Configuring Unique Listen Address Attributes for SIP Data Tier Replicas.....	3-8
3.8	Production Network Architectures and Configuration	3-8
3.8.1	Single-NIC Configurations with TCP and UDP Channels.....	3-9
3.8.1.1	Static Port Configuration for Outbound UDP Packets	3-10
3.8.2	DNS Server Configurations Overview	3-11
3.8.3	DNS Servers Listening On All Addresses (IP_ANY)	3-11
3.8.4	DNS Servers Listening on Multiple Subnets	3-11
3.8.4.1	Understanding the Route Resolver.....	3-12
3.8.4.2	IP Aliasing with DNS Hardware.....	3-13
3.8.5	Load Balancer Configurations	3-13
3.8.5.1	Single Load Balancer Configuration.....	3-13
3.8.5.2	Multiple Load Balancers and DNS Load Balancers.....	3-14
3.8.5.3	Network Address Translation Options.....	3-14
3.8.5.3.1	IP Masquerading Alternative to Source NAT.....	3-14
3.9	Example Network Configuration	3-15
3.9.1	Example Network Topology	3-15
3.9.2	Oracle WebLogic Server SIP Container Configuration.....	3-15
3.9.3	Load Balancer Configuration.....	3-16
3.9.3.1	NAT-based configuration.....	3-16
3.9.3.2	maddr-Based Configuration	3-21
3.9.3.3	rport-Based Configuration	3-22

4 Configuring SIP Data Tier Partitions and Replicas

4.1	Overview of SIP Data Tier Configuration.....	4-1
4.1.1	datatier.xml Configuration File	4-2
4.1.2	Configuration Requirements and Restrictions	4-2
4.2	Best Practices for Configuring and Managing SIP Data Tier Servers	4-3
4.3	Example SIP Data Tier Configurations and Configuration Files.....	4-3
4.3.1	SIP Data Tier with One Partition.....	4-4
4.3.2	SIP Data Tier with Two Partitions.....	4-4
4.3.3	SIP Data Tier with Two Partitions and Two Replicas	4-4

4.4	Storing Long-Lived Call State Data In A RDBMS.....	4-5
4.4.1	Requirements and Restrictions	4-5
4.4.2	Steps for Enabling RDBMS Call State Storage.....	4-6
4.4.3	Using the Configuration Wizard RDBMS Store Template	4-6
4.4.3.1	Modify the JDBC Datasource Connection Information	4-7
4.4.4	Configuring RDBMS Call State Storage by Hand.....	4-7
4.4.4.1	Configure JDBC Resources.....	4-7
4.4.4.2	Configure Oracle WebLogic Server SIP Container Persistence Options.....	4-8
4.4.4.3	Create the Database Schema	4-8
4.4.5	Using Persistence Hints in SIP Applications	4-9
4.5	Introducing Geo-Redundancy	4-9
4.5.1	Situations Best Suited to Use Geo-Redundancy.....	4-11
4.5.2	Situations Not Suited to Use Geo-Redundancy	4-11
4.5.3	Geo-Redundancy Considerations: Before Your Begin	4-11
4.6	Using Geographically-Redundant SIP Data Tiers.....	4-12
4.6.1	Example Domain Configurations.....	4-13
4.6.2	Requirements and Limitations.....	4-14
4.6.3	Steps for Configuring Geographic Persistence.....	4-15
4.6.4	Using the Configuration Wizard Templates for Geographic Persistence	4-15
4.6.4.1	Installing and Configuring the Primary Site	4-16
4.6.4.2	Installing the Secondary Site.....	4-16
4.6.5	Manually Configuring Geographical Redundancy	4-17
4.6.5.1	Configuring JDBC Resources (Primary and Secondary Sites)	4-17
4.6.5.2	Configuring Persistence Options (Primary and Secondary Sites).....	4-18
4.6.5.3	Configuring JMS Resources (Secondary Site Only).....	4-18
4.6.6	Understanding Geo-Redundant Replication Behavior	4-20
4.6.6.1	Call State Replication Process	4-20
4.6.6.2	Call State Processing After Failover.....	4-20
4.6.7	Removing Backup Call States	4-21
4.6.8	Monitoring Replication Across Regional Sites	4-21
4.6.9	Troubleshooting Geographical Replication	4-22
4.7	Caching SIP Data in the Engine Tier	4-22
4.7.1	Configuring Engine Tier Caching	4-22
4.7.2	Monitoring and Tuning Cache Performance.....	4-22
4.8	Monitoring and Troubleshooting SIP Data Tier Servers.....	4-23

5 Monitoring and Troubleshooting

5.1	Avoiding and Recovering from Server Failures.....	5-1
5.1.1	Failure Prevention and Automatic Recovery Features	5-1
5.1.1.1	Overload Protection	5-2
5.1.1.2	Redundancy and Failover for Clustered Services	5-2
5.1.1.3	Automatic Restart for Failed Server Instances	5-2
5.1.1.4	Managed Server Independence Mode.....	5-3
5.1.1.5	Automatic Migration of Failed Managed Servers	5-3
5.1.1.6	Geographic Redundancy for Regional Site Failures	5-3
5.1.2	Directory and File Backups for Failure Recovery	5-3
5.1.2.1	Enabling Automatic Configuration Backups	5-4

5.1.2.2	Storing the Domain Configuration Offline	5-4
5.1.2.3	Backing Up Server Start Scripts.....	5-5
5.1.2.4	Backing Up Logging Servlet Applications.....	5-5
5.1.2.5	Backing Up Security Data.....	5-5
5.1.2.5.1	Backing Up SerializedSystemIni.dat and Security Certificates	5-5
5.1.2.5.2	Backing Up the WebLogic LDAP Repository	5-5
5.1.2.6	Backing Up Additional Operating System Configuration Files	5-6
5.1.3	Restarting a Failed Administration Server.....	5-6
5.1.3.1	Restarting an Administration Server on the Same Machine	5-7
5.1.3.2	Restarting an Administration Server on Another Machine	5-7
5.1.4	Restarting Failed Managed Servers	5-8
5.2	Overview of Failover Detection	5-8
5.2.1	WssEchoServer Failure Detection	5-9
5.2.2	Forced Shutdown for Failed Replicas.....	5-9
5.3	Improving Failover Performance for Physical Network Failures.....	5-10
5.3.1	Starting WssEchoServer on SIP Data Tier Server Machines	5-10
5.3.2	Enabling and Configuring the Heartbeat Mechanism on Servers.....	5-11
5.4	Configuring SNMP	5-11
5.4.1	Browsing the MIB	5-12
5.4.2	Steps for Configuring SNMP	5-12
5.5	Understanding and Responding to SNMP Traps	5-12
5.5.1	Files for Troubleshooting.....	5-13
5.5.2	Trap Descriptions.....	5-13
5.5.2.1	connectionLostToPeer.....	5-13
5.5.2.1.1	Recovery Procedure	5-14
5.5.2.2	connectionReestablishedToPeer	5-14
5.5.2.2.1	Recovery Procedure	5-14
5.5.2.3	dataTierServerStopped	5-14
5.5.2.3.1	Recovery Procedure	5-14
5.5.2.4	overloadControlActivated, overloadControlDeactivated	5-14
5.5.2.4.1	Recovery Procedure	5-14
5.5.2.4.2	Additional Overload Information.....	5-14
5.5.2.5	replicaAddedToPartition.....	5-15
5.5.2.5.1	Recovery Procedure	5-15
5.5.2.6	replicaRemovedEnginesRegistration.....	5-15
5.5.2.6.1	Recovery Procedure	5-15
5.5.2.7	replicaRemovedFromPartition	5-15
5.5.2.7.1	Recovery Procedure	5-15
5.5.2.8	serverStopped	5-15
5.5.2.8.1	Recovery Procedure	5-15
5.5.2.8.2	Additional Shutdown Information.....	5-16
5.5.2.9	sipAppDeployed.....	5-16
5.5.2.9.1	Recovery Procedure	5-16
5.5.2.10	sipAppUndeployed.....	5-16
5.5.2.10.1	Recovery Procedure	5-16
5.5.2.11	sipAppFailedToDeploy	5-16
5.5.2.11.1	Recovery Procedure	5-16

5.6	Using the WebLogic Diagnostics Framework (WLDF).....	5-16
5.6.1	Data Collection and Logging	5-17
5.6.2	Watches and Notifications.....	5-17
5.6.3	Image Capture.....	5-18
5.6.4	Instrumentation.....	5-18
5.6.4.1	Configuring Server-Scoped Monitors.....	5-20
5.6.4.2	Configuring Application-Scoped Monitors.....	5-22
5.7	Logging SIP Requests and Responses.....	5-22
5.7.1	Defining Logging Servlets in sip.xml	5-23
5.7.2	Configuring the Logging Level and Destination	5-23
5.7.3	Specifying the Criteria for Logging Messages.....	5-23
5.7.3.1	Using XML Documents to Specify Logging Criteria.....	5-23
5.7.3.2	Using Servlet Parameters to Specify Logging Criteria.....	5-24
5.7.4	Specifying Content Types for Unencrypted Logging.....	5-26
5.7.5	Enabling Log Rotation and Viewing Log Files.....	5-26
5.7.6	trace-pattern.dtd Reference	5-26
5.7.7	Adding Tracing Functionality to SIP Servlet Code	5-28
5.7.8	Order of Startup for Listeners and Logging Servlets	5-29
5.8	Tuning JVM Garbage Collection for Production Deployments	5-29
5.8.1	Modifying JVM Parameters in Server Start Scripts	5-29
5.8.2	Tuning Garbage Collection with JRockit.....	5-30
5.8.3	Using Oracle JRockit Real Time (Deterministic Garbage Collection)	5-30
5.8.4	Using Oracle JRockit without Deterministic Garbage Collection.....	5-31
5.8.5	Tuning Garbage Collection with Sun JDK.....	5-31
5.9	Avoiding JVM Delays Caused By Random Number Generation.....	5-32

6 Configuring Diameter Client Nodes and Relay Agents

6.1	Overview of Diameter Protocol Configuration	6-1
6.2	Steps for Configuring Diameter Client Nodes and Relay Agents	6-2
6.3	Installing the Diameter Domain	6-2
6.4	Enabling the Diameter Console Extension.....	6-4
6.5	Creating TCP, TLS, and SCTP Network Channels for the Diameter Protocol	6-4
6.5.1	Configuring Two-Way SSL for Diameter TLS Channels	6-6
6.5.2	Configuring and Using SCTP for Diameter Messaging.....	6-6
6.6	Configuring Diameter Nodes.....	6-7
6.6.1	Creating a New Node Configuration (General Node Configuration).....	6-7
6.6.2	Configuring Diameter Applications	6-9
6.6.2.1	Configuring the Sh Client Application.....	6-10
6.6.2.2	Configuring the Rf Client Application	6-11
6.6.2.3	Configuring the Ro Client Application	6-11
6.6.2.4	Configuring a Diameter Relay Agent.....	6-11
6.6.2.5	Configuring the Sh and Rf Simulator Applications	6-13
6.6.2.6	Enabling Profile Service (using an Sh backend).....	6-13
6.6.3	Configuring Peer Nodes	6-14
6.6.4	Configuring Routes	6-14
6.7	Example Domain Configuration.....	6-15
6.8	Troubleshooting Diameter Configurations.....	6-18

7 Oracle WebLogic Server SIP Container Base Platform Topologies

7.1	Goals of the Oracle WebLogic Server SIP Container Base Platform	7-1
7.2	Load Balancer	7-2
7.3	Engine Tier	7-3
7.4	SIP Data tier	7-3
7.4.1	Example of Writing and Retrieving Call State Data	7-4
7.4.2	RDBMS Storage for Long-Lived Call State Data	7-4
7.5	Geographically-Redundant Installations	7-5
7.6	Example Hardware Configurations	7-5
7.7	Alternate Configurations	7-5

8 Upgrading Deployed SIP Applications

8.1	Overview of SIP Application Upgrades	8-1
8.2	Requirements and Restrictions for Upgrading Deployed Applications.....	8-2
8.3	Steps for Upgrading a Deployed SIP Application	8-3
8.4	Assign a Version Identifier	8-3
8.4.1	Defining the Version in the Manifest.....	8-3
8.5	Deploy the Updated Application Version.....	8-4
8.6	Undeploy the Older Application Version	8-4
8.7	Roll Back the Upgrade Process	8-5
8.8	Accessing the Application Name and Version Identifier	8-5
8.9	Using Administration Mode	8-5

A SIP Servlet Container Configuration Reference

A.1	Overview of sipserver.xml	A-1
A.2	Editing sipserver.xml	A-1
A.2.1	Steps for Editing sipserver.xml.....	A-2
A.3	XML Schema	A-2
A.4	Example sipserver.xml File.....	A-2
A.5	XML Element Description	A-2
A.5.1	enable-timer-affinity.....	A-2
A.5.2	overload.....	A-3
A.5.2.1	Selecting an Appropriate Overload Policy	A-4
A.5.2.2	Overload Control Based on Session Generation Rate	A-5
A.5.2.3	Overload Control Based on Capacity Constraints.....	A-5
A.5.2.4	Two Levels of Overload Protection	A-6
A.5.3	message-debug.....	A-6
A.5.4	proxy—Setting Up an Outbound Proxy Server	A-6
A.5.5	t1-timeout-interval.....	A-7
A.5.6	t2-timeout-interval.....	A-7
A.5.7	t4-timeout-interval.....	A-7
A.5.8	timer-b-timeout-interval	A-8
A.5.9	timer-f-timeout-interval	A-8
A.5.10	max-application-session-lifetime.....	A-8
A.5.11	enable-local-dispatch.....	A-8
A.5.12	cluster-loadbalancer-map	A-8

A.5.13	default-behavior.....	A-9
A.5.14	default-servlet-name	A-10
A.5.15	retry-after-value	A-10
A.5.16	sip-security.....	A-10
A.5.17	route-header	A-10
A.5.18	engine-call-state-cache-enabled	A-11
A.5.19	server-header.....	A-11
A.5.20	server-header-value.....	A-11
A.5.21	persistence.....	A-12
A.5.22	use-header-form.....	A-12
A.5.23	enable-dns-srv-lookup	A-13
A.5.24	connection-reuse-pool.....	A-13
A.5.25	globally-routable-uri	A-14
A.5.26	domain-alias-name	A-15
A.5.27	enable-rport	A-15
A.5.28	image-dump-level.....	A-16
A.5.29	stale-session-handling.....	A-16
A.5.30	enable-contact-provisional-response	A-16
A.5.31	app-router	A-17
A.5.32	use-custom-app-router.....	A-17
A.5.33	app-router-config-data.....	A-17
A.5.34	custom-app-router-jar-file-name	A-17
A.5.35	default-application-name	A-18

B SIP Data Tier Configuration Reference

B.1	Overview of datatier.xml.....	B-1
B.2	Editing datatier.xml.....	B-1
B.3	XML Schema	B-1
B.4	Example datatier.xml File.....	B-1
B.5	XML Element Description	B-2

C Diameter Configuration Reference

C.1	Overview of diameter.xml.....	C-1
C.2	Graphical Representation	C-1
C.3	Editing diameter.xml.....	C-2
C.3.1	Steps for Editing diameter.xml	C-3
C.4	XML Schema	C-3
C.5	Example diameter.xml File.....	C-3
C.6	XML Element Description	C-3
C.6.1	configuration	C-3
C.6.2	target.....	C-4
C.6.3	host.....	C-4
C.6.4	realm	C-4
C.6.5	address	C-4
C.6.6	port.....	C-4
C.6.7	tls-enabled.....	C-4

C.6.8	sctp-enabled.....	C-5
C.6.9	debug-enabled.....	C-5
C.6.10	message-debug-enabled	C-5
C.6.11	application	C-5
C.6.11.1	class-name.....	C-5
C.6.11.2	param*	C-5
C.6.11.2.1	name	C-5
C.6.11.2.2	value	C-5
C.6.12	peer-retry-delay	C-5
C.6.13	allow-dynamic-peers.....	C-5
C.6.14	request-timeout	C-5
C.6.15	watchdog-timeout.....	C-6
C.6.16	supported-vendor-id+.....	C-6
C.6.17	include-origin-state.....	C-6
C.6.18	peer+	C-6
C.6.18.1	host.....	C-6
C.6.18.2	address	C-6
C.6.18.3	port.....	C-6
C.6.18.4	protocol	C-6
C.6.19	route.....	C-6
C.6.19.1	realm	C-7
C.6.19.2	application-id	C-7
C.6.19.3	action	C-7
C.6.19.4	server+.....	C-7
C.6.20	default-route.....	C-7
C.6.20.1	action	C-7
C.6.20.2	server+.....	C-7

D Startup Command Options

E Supported Platforms, Protocols, RFCs and Standards

E.1	Supported Configurations	E-1
E.2	Supported SIP Clients	E-2
E.3	Supported Load Balancer	E-2
E.4	Supported Databases.....	E-2
E.5	Overview of Oracle WebLogic Server SIP Container Standards Alignment.....	E-2
E.6	Java Sun Recommendation (JSR) Standards Compliance.....	E-2
E.7	IETF RFC Compliance.....	E-3
E.8	3GPP R6 Specification Conformance	E-11

F Using Oracle WebLogic SIP Container Export/Import

F.1	Export	F-1
F.1.1	Export the Database Data from the Current Environment.....	F-1
F.2	Import	F-3

Index

Preface

This book details conceptual, topology and configuration topics about Oracle WebLogic SIP Container. This Preface includes the following topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Documents](#)
- [Conventions](#)

Audience

The intended audience is system administrators who will set up and maintain Oracle WebLogic SIP Container for their organization.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle

technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at <http://www.fcc.gov/cgb/consumerfacts/trs.html>, and a list of phone numbers is available at <http://www.fcc.gov/cgb/dro/trsphonebk.html>.

Related Documents

For more information, see the following documents in the documentation set:

- *Oracle WebLogic SIP Container Installation Guide*
- *Oracle WebLogic SIP Container Developer's Guide*
- *Oracle Fusion Middleware 11g Release Notes*

Conventions

The following text conventions are used in this document:

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

Configuring Oracle WebLogic SIP Container

The following section provides an introduction to Oracle WebLogic SIP Container (OWLSC), and an overview of how to configure and manage Oracle WebLogic Server SIP Container deployments:

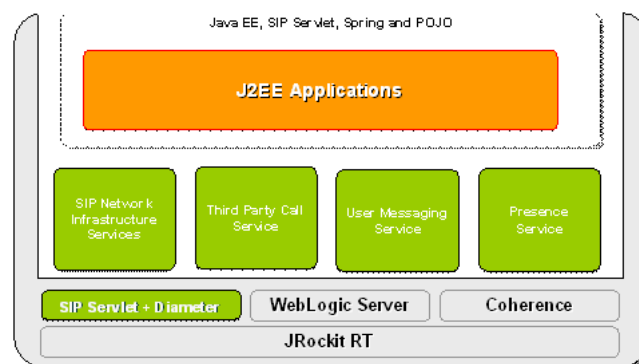
- [Section 1.1, "Oracle WebLogic SIP Container"](#)
- [Section 1.2, "Shared Configuration Tasks"](#)

1.1 Oracle WebLogic SIP Container

Oracle WebLogic Server SIP Container 11g (OWLSC) is a comprehensive platform designed to integrate communication services with enterprise services and applications. It includes easy to consume services to support interactions with key communication channels.

[Table 1–1](#) shows a simplified overview of Oracle WebLogic Server SIP Container; more details will be presented in later chapters.

Figure 1–1 OWLSC overview



1.1.1 WebLogic Server 10.3 Platform Supports SIP and Converged Applications

OWLSC extends the core WebLogic Server platform with a SIP Container compliant with JSR 289. This enables the development of J2EE applications that processes SIP in addition to HTTP for any advanced communications application. The platform enables the development of complementary communications services that integrate with SIP-based IP-PBXs as well as other SIP elements such as standard SIP clients.

Out of the box, OWLSC provides the key infrastructure applications to build a SIP-based network.

1.2 Shared Configuration Tasks

The following sections provide an overview of the configuration tasks that are common to both Oracle WebLogic SIP Container and Oracle WebLogic Server. These topics are included:

- [Section 1.2.1, "Shared Configuration Tasks for SIP Container and WebLogic Server"](#)
- [Section 1.2.2, "Oracle WebLogic Server SIP Container Configuration Overview"](#)
- [Section 1.2.3, "Methods and Tools for Performing Configuration Tasks"](#)
- [Section 1.2.4, "Starting and Stopping Servers"](#)
- [Section 1.2.5, "Administration Server Best Practices"](#)
- [Section 1.2.6, "Common Configuration Tasks"](#)

1.2.1 Shared Configuration Tasks for SIP Container and WebLogic Server

Oracle WebLogic Server SIP Container is based on the Oracle WebLogic Server 10g Release 3 application server, and many system-level configuration tasks are the same for both products. This guide addresses only those system-level configuration tasks that are unique to Oracle WebLogic Server SIP Container, such as tasks related to network and security configuration and cluster configuration for the engine and SIP data tiers.

HTTP server configuration and other basic configuration tasks such as server logging are addressed in Oracle WebLogic Server documentation. See *Oracle Fusion Middleware Getting Started With Installation for Oracle WebLogic Server* to get started.

1.2.2 Oracle WebLogic Server SIP Container Configuration Overview

The SIP Servlet container, SIP data tier replication, and Diameter protocol features of Oracle WebLogic Server SIP Container are implemented in the Oracle WebLogic Server 10g Release 3 product as custom resources. A pair of custom resources, `sipserver` and `datatier`, implement the engine tier SIP Servlet container functionality and SIP data tier replication functionality. In production deployments, both resources are generally installed. Specialized deployments may use only the `sipserver` resource in conjunction with a SIP-aware load balancer.

Another custom resource, `diameter`, provides Diameter base protocol functionality, and is required only for deployments that utilize one or more Diameter protocol applications.

The Oracle WebLogic SIP Container custom resource assignments are visible in the domain configuration file, `config.xml`, and should not be modified. [Example 1-1](#) shows the definitions for each resource. Note that the `sipserver` and `datatier` resources must each be targeted to the same servers or clusters; the resources are deployed to both the engine tier and SIP data tier cluster.

Example 1-1 Oracle WebLogic Server SIP Container Custom Resources

```
<custom-resource>
  <name>sipserver</name>
  <target>ORA_DATA_TIER_CLUST,ORA_ENGINE_TIER_CLUST</target>
  <descriptor-file-name>custom/sipserver.xml</descriptor-file-name>

  <resource-class>com.bea.wcp.sip.management.descriptor.resource.SipServerResource</
resource-class>
```

```

<descriptor-bean-class>com.bea.wcp.sip.management.descriptor.beans.SipServerBean</
descriptor-bean-class>
</custom-resource>
<custom-resource>
  <name>datatier</name>
  <target>ORA_DATA_TIER_CLUST,ORA_ENGINE_TIER_CLUST</target>
  <descriptor-file-name>custom/datatier.xml</descriptor-file-name>

<resource-class>com.bea.wcp.sip.management.descriptor.resource.DataTierResource</r
esource-class>

<descriptor-bean-class>com.bea.wcp.sip.management.descriptor.beans.DataTierBean</d
escriptor-bean-class>
  </custom-resource>
<custom-resource>
  <name>diameter</name>
  <target>ORA_ENGINE_TIER_CLUST</target>
  <deployment-order>200</deployment-order>
  <descriptor-file-name>custom/diameter.xml</descriptor-file-name>
  <resource-class>com.bea.wcp.diameter.DiameterResource</resource-class>

<descriptor-bean-class>com.bea.wcp.diameter.management.descriptor.beans.Configurat
ionBean</descriptor-bean-class>
</custom-resource>

```

The Oracle WebLogic Server SIP Container custom resources utilize the basic domain resources defined in `config.xml`, such as network channels, cluster and server configuration, and Java EE resources. However, Oracle WebLogic Server SIP Container-specific resources are configured in separate configuration files based on functionality:

- `sipserver.xml` configures SIP container properties and general Oracle WebLogic Server SIP Container engine tier functionality.
- `datatier.xml` identifies servers that participate as replicas in the SIP data tier, and also defines the number and layout of SIP data tier partitions.
- `diameter.xml` configures Diameter nodes and Diameter protocol applications used in the domain.
- `approuter.xml` configures Default Application Router. For more information on configuring DAR, see *Oracle WebLogic Server Installation Guide*.

Keep in mind that the domain configuration file, `config.xml`, defines all of the Managed Servers available in the domain. The `sipserver.xml`, `datatier.xml`, and `diameter.xml` configuration files included in the `sipserver` application determine the role of each server instance, such as whether they behave as SIP data tier replicas, engine tier nodes, or Diameter client nodes.

Configuration changes to SIP Servlet container properties can be applied dynamically (some SIP Servlet container properties may display a *Restart may be required* icon meaning that restart after making the change will be required) to a running server by using the Administration Console, or from the command line using the WLST utility. Configuration for SIP data tier nodes cannot be changed dynamically, so you must reboot SIP data tier servers in order to change the number of partitions or replicas.

1.2.2.1 Diameter Configuration

The Diameter protocol implementation is implemented as a custom resource separate from the SIP Servlet container functionality. The Diameter configuration file configures one or more Diameter protocol applications to provide Diameter node

functionality. Oracle WebLogic Server SIP Container provides the Diameter protocol applications to support the following node types:

- *Diameter Sh* interface client node (for querying a Home Subscriber Service)
- *Diameter Rf* interface client node (for offline charging)
- *Diameter Ro* interface client node (for online charging)
- *Diameter relay node*
- *HSS simulator node* (suitable for testing and development only, not for production deployment)

The Diameter custom resource is deployed only to domains having servers that must function as Diameter client nodes or relay agents, or to servers providing HSS simulation capabilities. The actual function of the server instance depends on the configuration defined in the `diameter.xml` file.

See [Chapter 6, "Configuring Diameter Client Nodes and Relay Agents"](#) for instructions to configure the Diameter Web Application in an Oracle WebLogic Server SIP Container domain. See *Oracle WebLogic Server SIP Container Developer's Guide* for information on developing Diameter applications.

1.2.3 Methods and Tools for Performing Configuration Tasks

Oracle WebLogic Server SIP Container provides several mechanisms for changing the configuration of the SIP Servlet container:

- [Section 1.2.3.1, "Administration Console"](#)
- [Section 1.2.3.2, "WebLogic Scripting Tool \(WLST\)"](#)
- [Section 1.2.3.3, "Additional Configuration Methods"](#)

1.2.3.1 Administration Console

Oracle WebLogic Server SIP Container provides Administration Console extensions that allow you to modify and SIP Servlet container, SIP Servlet domain, and Diameter configuration properties using a graphical user interface. The Administration Console extensions for Oracle WebLogic Server SIP Container are similar to the core console available in Oracle WebLogic Server 10g Release 3. All Oracle WebLogic Server SIP Container configuration and monitoring is provided via these nodes in the left pane of the console:

- **SipServer**—configures SIP Servlet container properties and other engine tier functionality. This extension also enables you to create new partitions, and view (but not modify) SIP data tier partitions and replicas. See [Chapter 2, "Configuring SIP Servlet Container Properties"](#) for more information about configuring the SIP Servlet container using the Administration Console.
- **Diameter**—configures Diameter nodes and applications.

Note: To learn more about using Oracle WebLogic Server Administration Console, see "Getting Started with Oracle WebLogic Server Administration Console" in *Oracle Fusion Middleware Administrator's Guide*.

1.2.3.2 WebLogic Scripting Tool (WLST)

The WebLogic Scripting Tool (WLST) enables you to perform interactive or automated (batch) configuration operations using a command-line interface. WLST is a JMX tool that can view or manipulate the MBeans available in a running Oracle WebLogic Server SIP Container domain. [Chapter 2, "Configuring SIP Servlet Container Properties"](#) provides instructions for modifying SIP Servlet container properties using WLST.

Note: To learn more about using WLST, see *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

1.2.3.3 Additional Configuration Methods

Most Oracle WebLogic Server SIP Container configuration is performed using either the Administration Console or WLST. The methods described in the following sections may also be used for certain configuration tasks.

1.2.3.3.1 Editing Configuration Files You may also edit `sipserver.xml`, `datatier.xml`, `diameter.xml`, and `approuter.xml` manually. If you edit configuration files manually, you must reboot all servers to apply the configuration changes.

1.2.3.3.2 Custom JMX Applications Oracle WebLogic Server SIP Container properties are represented by JMX-compliant MBeans. You can therefore program JMX applications to configure SIP container properties using the appropriate Oracle WebLogic Server SIP Container MBeans.

The general procedure for modifying Oracle WebLogic Server SIP Container MBean properties using JMX is described in [Chapter 2, "Configuring SIP Servlet Container Properties"](#) (WLST itself is a JMX-based application). For more information about the individual MBeans used to manage SIP container properties, see the *Oracle Fusion Middleware Communication Services Java API Reference*.

1.2.3.3.3 Setting Log Levels You can set log levels by manually editing the `logging.xml` file, by setting the `setLoggerLevel(String loggerName, String logLevel)` MBean, or through Oracle Enterprise Manager. For more information, see *Oracle Fusion Middleware 2 Day Administration Guide*.

1.2.4 Starting and Stopping Servers

Oracle WebLogic Server SIP Container start scripts use default values for many JVM parameters that affect performance. For example, JVM garbage collection and heap size parameters may be omitted, or may use values that are acceptable only for evaluation or development purposes. In a production system, you must rigorously profile your applications with different heap size and garbage collection settings in order to realize adequate performance.

Caution: When you configure a domain with multiple engine and SIP data tier servers, you must accurately synchronize all system clocks to a common time source (to within one or two milliseconds) in order for the SIP protocol stack to function properly.

Because a typical Oracle WebLogic Server SIP Container domain contains numerous engine and SIP data tier servers, with dependencies between the different server types, you should generally follow this sequence when starting up a domain:

1. **Start the Administration Server for the domain.** Start the Administration Server in order to provide the initial configuration to engine and SIP data tier servers in the domain. The Administration Server can also be used to monitor the startup/shutdown status of each Managed Server. You generally start the Administration Server by using either the `startWebLogic.cmd` script installed with the Configuration Wizard, or a custom startup script.
2. **Start SIP data tier servers in each partition.** The engine tier cannot function until servers in the SIP data tier are available to manage call state data. Although all replicas in each partition need not be available to begin processing requests, at least one replica in each configured partition must be available in order to manage the concurrent call state. All replicas should be started and available before opening the system to production network traffic.

You generally start each SIP data tier server by using either the `startManagedWebLogic.cmd` script installed with the Configuration Wizard, or a custom startup script. `startManagedWebLogic.cmd` requires that you specify the name of the server to startup, as well as the URL of the Administration Server for the domain, as in:

```
startManagedWebLogic.cmd datanode0-0 t3://adminhost:7001
```

3. **Start engine tier servers.** After the SIP data tier servers have started, you can start servers in the engine tier and begin processing client requests. As with SIP data tier servers, engine tier servers are generally started using the `startManagedWebLogic.cmd` script or a custom startup script.

Following the above startup sequence ensures that all Managed Servers use the latest SIP Servlet container and SIP data tier configuration. This sequence also avoids engine tier error messages that are generated when servers in the SIP data tier are unavailable.

1.2.5 Administration Server Best Practices

The Administration Server in a Oracle WebLogic Server SIP Container installation is required for configuring, deploying, and monitoring services and applications.

Note: If an Administration Server fails due to a hardware, software, or network problem, only management, deployment, and monitoring operations are affected. **Managed Servers do not require the Administration Server for continuing operation; Java EE applications and SIP features running on Managed Server instances continue to function even if the Administration Server fails.**

Oracle recommends the following best practices for configuring Administration Server and Managed Server instances in your Oracle WebLogic Server SIP Container domain:

- Run the Administration Server instance on a dedicated machine. The Administration Server machine should have a memory capacity similar to Managed Server machines, although a single CPU is generally acceptable for administration purposes.
- Configure all Managed Server instances to use Managed Server Independence. This feature allows the Managed Servers to restart even if the Administration

Server is unreachable due to a network, hardware, or software failure. See *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* for more information.

- Configure the Node Manager utility to automatically restart all Managed Servers in the Oracle WebLogic Server SIP Container domain. See *Oracle Fusion Middleware Oracle WebLogic Scripting Tool* for more information.

Should an Administration Server instance or machine fail, remember that only configuration, deployment, and monitoring features are affected, but Managed Servers continue to operate and process client requests. Potential losses incurred due to an Administration Server failure include:

- Loss of in-progress management and deployment operations.
- Loss of ongoing logging functionality.
- Loss of SNMP trap generation for WebLogic Server instances (as opposed to Oracle WebLogic Server SIP Container instances). On Managed Servers, Oracle WebLogic Server SIP Container traps are generated even in the absence of the Administration Server.

To resume normal management activities, restart the failed Administration Server instance as soon as possible.

1.2.6 Common Configuration Tasks

General administration and maintenance of Oracle WebLogic Server SIP Container requires that you manage both WebLogic Server configuration properties and Oracle WebLogic Server SIP Container container properties. These common configuration tasks are summarized in [Table 1-1](#).

Table 1-1 Common Oracle WebLogic Server SIP Container Configuration Tasks

Task	Description
Chapter 2	<ul style="list-style-type: none"> ■ Configuring SIP Container Properties using the Administration Console ■ Using WLST to perform batch configuration
Chapter 4, "Configuring SIP Data Tier Partitions and Replicas"	<ul style="list-style-type: none"> ■ Assigning Oracle WebLogic Server SIP Container instances to the SIP data tier partitions ■ Replicating call state using multiple SIP data tier instances
Chapter 3, "Managing Network Resources"	<ul style="list-style-type: none"> ■ Configuring WebLogic Server network channels to handling SIP and HTTP traffic ■ Setting up multi-homed server hardware ■ Configuring load balancers for use with Oracle WebLogic Server SIP Container

Configuring SIP Servlet Container Properties

The following sections describe how to configure SIP Container features in the engine tier of an Oracle WebLogic Server SIP Container deployment:

- [Section 2.1, "Overview of SIP Container Configuration"](#)
- [Section 2.2, "Using the Administration Console to Configure Container Properties"](#)
- [Section 2.3, "Configuring Container Properties Using WLST \(JMX\)"](#)
- [Section 2.4, "WLST Configuration"](#)
- [Section 2.5, "Configuring Timer Processing"](#)

2.1 Overview of SIP Container Configuration

You can configure SIP Container properties either by using a JMX utility such as the Administration Console or WebLogic Scripting Tool (WLST), or by programming a custom JMX application. [Section 2.2, "Using the Administration Console to Configure Container Properties"](#) describes how to configure container properties using the Administration Console graphical user interface.

[Section 2.3, "Configuring Container Properties Using WLST \(JMX\)"](#) describes how to directly access JMX MBeans to modify the container configuration. All examples use WLST to illustrate JMX access to the configuration MBeans.

2.2 Using the Administration Console to Configure Container Properties

The Administration Console included with Oracle WebLogic SIP Container enables you to configure and monitor core WebLogic Server functionality as well as the SIP Servlet container functionality provided with Oracle WebLogic SIP Container. To configure or monitor SIP Servlet features using the Administration Console, see "Getting Started with Oracle WebLogic Server Administration Console" in *Oracle Fusion Middleware Administrator's Guide*.

Table 2–1 Oracle WebLogic Server SIP Container Configuration and Monitoring Pages

Page	SubPage	Function
Configuration	General	Configure SIP timer values, session timeout duration, default Oracle WebLogic Server SIP Container behavior (proxy or user agent), server header format, call state caching, DNS name resolution, timer affinity, domain aliases, rport support, and diagnostic image format.
Configuration	Application Router	Configure custom Application Router (AR) class name, configuration, or default application.
Configuration	Proxy	Configure proxy routing URIs and proxy policies.

Table 2–1 (Cont.) Oracle WebLogic Server SIP Container Configuration and Monitoring Pages

Page	SubPage	Function
Configuration	Overload Protection	Configure the conditions for enabling and disabling automatic overload controls.
Configuration	Message Debug	Enable or disable SIP message logging on a development system.
Configuration	SIP Security	Identify trusted hosts for which authentication is not performed.
Configuration	Persistence	Configure persistence options for storing long-lived session data in an RDBMS, or for replicating long-lived session data to a remote, geographically-redundant site.
Configuration	Data Tier	View the current configuration of SIP data tier servers. You can also add, delete and configure partitions here.
Configuration	LoadBalancer Map	Configure the mapping of multiple clusters to internal virtual IP addresses during a software upgrade.
Configuration	Targets	Configure the list of servers or clusters that receive the engine tier configuration. The target server list determines which servers and/or clusters provide SIP Servlet container functionality.
Configuration	Connection Pools	Configure connection reuse pools to minimize communication overhead with a Session Border Control (SBC) function or Serving Call Session Control Function (S-CSCF).
Monitoring	General	View runtime information about messages and sessions processed in engine tier servers.
Monitoring	SIP Applications	View runtime session information for deployed SIP applications.
Monitoring	Data Tier Information	View runtime information about the current status and the work performed by servers in the SIP data tier.

2.2.1 Locking and Persisting the Configuration

In order to modify information on any of the Oracle WebLogic Server SIP Container configuration pages, the configuration must be locked. Locking a configuration prevents other Administrators from modifying the configuration at the same time. Locking is automatically enabled in Production domains; it can be enabled or disabled in Development domains.

To make changes:

1. Locate the Change Center in the upper left corner of the Administration Console.
2. Click **Lock & Edit** to lock the editable configuration hierarchy for the domain. This enables you to make changes using the Administration Console.
3. Make the changes you desire on the relevant page of the Console and click **Save** on each page where you make a change.
4. When you have finished making all the desired changes, click **Activate Changes** in the Change Center.

Note: Some changes you make in the Administration Console take place immediately when you activate them. Other changes require you to restart the server or module affected by the change. These latter changes are called non-dynamic changes. Non-dynamic changes are indicated in the Administration Console with a warning icon.

If an edit is made to a non-dynamic configuration setting, no edits to dynamic configuration settings will take effect until after you restart the server.

For more information on using Oracle WebLogic Server Administration Console, see *Oracle Fusion Middleware Administrator's Guide*.

2.3 Configuring Container Properties Using WLST (JMX)

The WebLogic Scripting Tool (WLST) is a utility that you can use to observe or modify JMX MBeans available on a WebLogic Server or Oracle WebLogic Server SIP Container instance. Full documentation for WLST is available in *Oracle Fusion Middleware Oracle WebLogic Scripting Tool*.

Before using WLST to configure a Oracle WebLogic Server SIP Container domain, set you environment to add required Oracle WebLogic Server SIP Container classes to your classpath. Use either a domain environment script or the `setWLSEnv.sh` script located in `MIDDLEWARE_HOME/server/bin` where `MIDDLEWARE_HOME` is the root of your Oracle WebLogic Server SIP Container installation.

2.3.1 Managing Configuration Locks

Table 2–1 summarizes the WLST methods used to lock a configuration and apply changes.

Table 2–2 MBean Method Summary

Method	Description
<code>activate</code>	Writes the current configuration MBean attributes (the current SIP Servlet container configuration) to the <code>sipserver.xml</code> configuration file and applies changes to the running servers.
<code>cancelEdit</code>	Cancels an edit session, releasing the edit lock, and discarding all unsaved changes. This operation can be called by any user with administrator privileges, even if the user did not start the edit session.
<code>cd</code>	Navigate the hierarchy of configuration or runtime beans.
<code>connect</code>	Connect WLST to a WebLogic Server instance.
<code>edit</code>	Starts an edit session.
<code>save</code>	Writes the current configuration MBean attributes (the current SIP Servlet container configuration) to a temporary configuration file.
<code>set</code>	Set the specified attribute value for the current configuration bean.
<code>stopEdit</code>	Releases the lock obtained for modifying SIP container properties and rolls back any pending MBean changes, discarding any temporary files.

Here is an example of using the commands to modify the T1 Timer interval:

Example 2–1 Modifying T1 Timer Interval

```
connect()
edit()
cd('CustomResources/sipserver/Resource/sipserver')
set('T1TimeoutInterval',505)
activate()
```

2.3.2 Locating the Oracle WebLogic Server SIP Container MBeans

All SIP Servlet container configuration MBeans are located in the "serverConfig" MBean tree, accessed using the `serverConfig()` command in WLST. Within this bean tree, individual configuration MBeans can be accessed using the path:

```
CustomResources/sipserver/Resource/sipserver
```

For example, to browse the default Proxy MBean for a Oracle WebLogic Server SIP Container domain you would enter these WLST commands:

```
serverConfig()
cd('CustomResources/sipserver/Resource/sipserver/Proxy')
ls()
```

Runtime MBeans are accessed in the *custom* MBean tree, accessed using the `custom()` command in WLST. Runtime MBeans use the path:

```
mydomain:Location=myserver,Name=myserver,Type=mbeantype
```

Certain configuration settings, such as proxy and overload protection settings, are defined by default in `sipserver.xml`. Configuration MBeans are generated for these settings when you boot the associated server, so you can immediately browse the Proxy and OverloadProtection MBeans. Other configuration settings are not configured by default and you will need to create the associated MBeans before they can be accessed. See [Section 2.4.2, "Creating and Deleting MBeans"](#).

2.4 WLST Configuration

The following sections describe WLST scripts and commands for configuring SIP Servlet container properties.

2.4.1 Invoking WLST

To use WLST with Oracle WebLogic Server SIP Container, you must ensure that all Oracle WebLogic Server SIP Container JAR files are included in your classpath. Follow these steps:

1. Set your Oracle WebLogic Server SIP Container environment:

```
cd MIDDLEWARE_HOME/user_projects/domains/mydomain/bin
./setDomainEnv.sh
```

2. Start WLST:

```
java weblogic.WLST
```

3. Connect to the Administration Server for your Oracle WebLogic Server SIP Container domain:

```
connect('system','weblogic','t3://myadminserver:<portnumber>')
```


2.4.2 Creating and Deleting MBeans

The `SipServer` MBean represents the entire contents of the `sipserver.xml` configuration file. In addition to having several attributes for configuring SIP timers and SIP application session timeouts, `SipServer` provides helper methods to help you create or delete MBeans representing proxy settings and overload protection controls. See [Table 2-1, "Oracle WebLogic Server SIP Container Configuration and Monitoring Pages"](#) for a listing of other helper methods in `SipServer`; see also *Oracle Fusion Middleware Communication Services Java API Reference*.

2.5 Configuring Timer Processing

As engine tier servers add new call state data to the SIP data tier, SIP data tier instances queue and maintain the complete list of SIP protocol timers and application timers associated with each call. Engine tier servers periodically poll partitions in the SIP data tier to determine which timers have expired, given the current time. By default, multiple engine tier polls to the SIP data tier are staggered to avoid contention on the timer tables. Engine tier servers then process all expired timers using threads allocated in the `sip.timer.Default` execute queue.

2.5.1 Configuring Timer Affinity (Optional)

With the default timer processing mechanism, a given engine tier server processes all timers that are currently due to fire, regardless of whether or not that engine was involved in processing the calls associated with those timers. However, some deployment scenarios require that a timer is processed on the same engine server that last modified the call associated with that timer. One example of this scenario is a hot standby system that maintains a secondary engine that should not process any call data until another engine fails. Oracle WebLogic Server SIP Container enables you to configure timer affinity in such scenarios.

When you enable timer affinity, replicas request that each engine tier server periodically poll the SIP data tier for processed timers. When polling the SIP data tier, an engine processes only those timers associated with calls that were last modified by that engine, or timers for calls that have no owner.

Note: When an engine tier server fails, any call states that were last modified by that engine no longer have an owner. Expired timers that have no owner are processed by the next engine server that polls the SIP data tier.

To enable timer affinity:

1. Access the Administration Console for your domain.
2. Click **Lock & Edit** (if enabled in your development environment) to obtain a configuration lock.
3. Select the `SipServer` node in the left pane. The right pane of the console provides two levels of tabbed pages that are used for configuring and monitoring Oracle WebLogic Server SIP Container.
4. Select the **Configuration > General** tab in the right pane.
5. Select *Enable Timer Affinity*.
6. Click **Save** to save your configuration changes.

Note that the Enable Timer Affinity setting is persisted in `sipserver.xml` in the element, `enable-timer-affinity`.

2.5.2 Configuring NTP for Accurate SIP Timers

In order for the SIP protocol stack to function properly, all engine and SIP data tier servers must accurately synchronize their system clocks to a common time source, to within one or two milliseconds. Large differences in system clocks cause a number of severe problems such as:

- SIP timers firing prematurely on servers with fast clock settings.
- Poor distribution of timer processing in the engine tier. For example, one engine tier server may process all expired timers, whereas other engine tier servers process no timers.

Oracle recommends using a Network Time Protocol (NTP) client or daemon on each Oracle WebLogic Server SIP Container instance and synchronizing to a common NTP server.

Caution: You must accurately synchronize server system clocks to a common time source (to within one or two milliseconds) in order for the SIP protocol stack to function properly. Because the initial T1 timer value of 500 milliseconds controls the retransmission interval for INVITE request and responses, and also sets the initial values of other timers, even small differences in system clock settings can cause improper SIP protocol behavior. For example, an engine tier server with a system clock 250 milliseconds faster than other servers will process more expired timers than other engine tier servers, will cause retransmits to begin in half the allotted time, and may force messages to timeout prematurely.

Managing Network Resources

The following sections describe how to configure network resources for use with Oracle WebLogic Server SIP Container:

- [Section 3.1, "Overview of Network Configuration"](#)
- [Section 3.2, "Configuring Load Balancer Addresses"](#)
- [Section 3.3, "Enabling Domain Name Service \(DNS\) Support"](#)
- [Section 3.4, "Configuring Network Channels for SIP or SIPS"](#)
- [Section 3.5, "Configuring TCP and TLS Channels for Diameter Support"](#)
- [Section 3.6, "Configuring Engine Servers to Listen on Any IP Interface"](#)
- [Section 3.7, "Configuring Unique Listen Address Attributes for SIP Data Tier Replicas"](#)

3.1 Overview of Network Configuration

The default HTTP network configuration for each Oracle WebLogic Server SIP Container instance is determined from the Listen Address and Listen Port setting for each server. However, Oracle WebLogic Server SIP Container does not support the SIP protocol over HTTP. The SIP protocol is supported over the UDP and TCP transport protocols. SIPS is also supported using the TLS transport protocol.

To enable UDP, TCP, or TLS transports, you configure one or more *network channels* for a Oracle WebLogic Server SIP Container instance. A network channel is a configurable WebLogic Server resource that defines the attributes of a specific network connection to the server instance. Basic channel attributes include:

- The protocols supported by the connection
- The listen address (DNS name or IP address) of the connection
- The port number used by the connection
- (optional) The port number used by outgoing UDP packets
- The public listen address (load balancer address) to embed in SIP headers when the channel is used for an outbound connection.

You can assign multiple channels to a single Oracle WebLogic Server SIP Container instance to support multiple protocols or to utilize multiple interfaces available with DNS server hardware. You cannot assign the same channel to multiple server instances.

When you configure a new network channel for the SIP protocol, both the UDP and TCP transport protocols are enabled on the specified port. You cannot create a SIP

channel that supports only UDP transport or only TCP transport. When you configure a network channel for the SIPS protocol, the server uses the TLS transport protocol for the connection.

As you configure a new SIP Server domain, you will generally create multiple SIP channels for communication to each engine tier server in your system. Engine tier servers can communicate to SIP data tier replicas using the configured Listen Address attributes for the replicas. Note, however, that replicas must use unique Listen Addressees in order to communicate with one another.

Note: If you configure multiple replicas in a SIP data tier cluster, you must configure a unique Listen Address for each server (a unique DNS name or IP address). If you do not specify a unique Listen Address, the replica service binds to the default *localhost* address and multiple replicas cannot locate one another.

3.1.1 IPv4 and IPv6

If your operating system and hardware support IPv6, you can also configure Oracle WebLogic Server SIP Container to use IPv6 for network communication. IPv6 for SIP traffic is enabled by configuring a network channel with an IPv6 address. You must configure an IPv6 SIP channel on each engine tier server that will support IPv6 traffic. See [Section 3.4.2, "Creating a New SIP or SIPS Channel"](#) for instructions.

Note that each SIP network channel configured on an engine supports either IPv6 or IPv4 traffic. You cannot mix IPv4 and IPv6 traffic on a single channel. A single engine can be configured with both an IPv4 and IPv6 channel to support multiple, separate networks.

It is also possible for Oracle WebLogic Server SIP Container engine and SIP data tier nodes to communicate on IPv4 (or IPv6) while supporting the other protocol version for external SIP traffic. To configure engine and SIP data tier nodes on an IPv6 network, simply specify IPv6 listen addresses for each server instance.

See [Section 3.4.2, "Creating a New SIP or SIPS Channel"](#) for information about configuring IPv4 and IPv6 network channels.

3.2 Configuring Load Balancer Addresses

If your system uses one or more load balancers to distribute connections to the engine tier, you must configure SIP network channels to include a load balancer address as the *external listen address*. When a SIP channel has an external listen address that differs from the channel's primary listen address, Oracle WebLogic Server SIP Container embeds the host and port number of the external address in SIP headers such as Response. In this way, subsequent communication for the call is directed to the public load balancer address, rather than the local engine tier server address (which may not be accessible to external clients).

If a network channel does not have a configured external listen address, the primary listen address is embedded into SIP headers.

3.2.1 Multiple Load Balancers and DNS Load Balancers

If your system uses two load balancers, you must define two channels on each engine tier server (one for each network connection to each load balancer) and assign the external listen address to the corresponding load balancer. When a particular network interface on the engine tier server is selected for outbound traffic, the network channel

associated with that Network Interface Card's (NIC) address is examined to determine the external listen address to embed in SIP headers.

If your system uses a DNS load balancer having two public addresses, you must also define a pair of channels to configure both public addresses. If the engine tier server has only one NIC, you must define a second, logical address on the NIC to configure a dedicated channel for the second public address. In addition, you must configure your IP routing policies to define which logical address is associated with each public load balancer address.

3.3 Enabling Domain Name Service (DNS) Support

Oracle WebLogic Server SIP Container supports DNS for resolving the transport, IP address and port number of a proxy required to send a SIP message. This matches the behavior described in RFC 3263 (<http://www.ietf.org/rfc/rfc3263.txt>). DNS may also be used when routing responses in order to resolve the IP address and port number of a destination.

Note: Because DNS resolution is performed within the context of SIP message processing, any DNS performance problems result in increased latency. Oracle recommends using a caching DNS server in a production environment to minimize potential performance problems.

To configure DNS support:

1. Log in to the Administration Console for the Oracle WebLogic Server SIP Container domain you want to configure.
2. Select the *SipServer* node in the left pane of the Console.
3. Select the **Configuration > General** tab in the right pane.
4. Select **Enable DNS Server Lookup**.
5. Click **Save** to save your changes.

When you enable DNS lookup, the server can use DNS to:

- Discover a proxy server's transport, IP address, and port number when a request is sent to a SIP URI.
- Resolve an IP address and/or port number during response routing, depending on the contents of the *Sent-by* field.

For proxy discovery, Oracle WebLogic Server SIP Container uses DNS resolution only once per SIP transaction to determine transport, IP, and port number information. All retransmissions, ACKs, or CANCEL requests are delivered to the same address and port using the same transport. For details about how DNS resolution takes place, see RFC 3263 (<http://www.ietf.org/rfc/rfc3263.txt>).

When a proxy is required to send a response message, Oracle WebLogic Server SIP Container uses DNS lookup to determine the IP address and/or port number of the destination, depending on the information provided in the *Sent-by* field and through header.

3.4 Configuring Network Channels for SIP or SIPS

When you create a new domain using the Configuration Wizard, Oracle WebLogic Server SIP Container instances are configured with a default network channel supporting the SIP protocol over UDP and TCP. This default channel is configured to use Listen Port 5060, but specifies no Listen Address. Follow the instructions in [Section 3.4.1, "Reconfiguring an Existing Channel"](#) to change the default channel's listen address or listen port settings. See [Section 3.4.2, "Creating a New SIP or SIPS Channel"](#) to create a new channel resource to support additional protocols or additional network interfaces.

3.4.1 Reconfiguring an Existing Channel

You cannot change the protocol supported by an existing channel. To reconfigure an existing listen address/port combination to use a different network protocol, you must delete the existing channel and create a new channel using the instructions in [Section 3.4.2, "Creating a New SIP or SIPS Channel"](#).

To configure a channel:

1. Log in to the Administration Console for the Oracle WebLogic Server SIP Container domain you want to configure.
2. In the left pane, select the **Environment > Servers** tab.
3. In the right pane, select the name of the server you want to configure.
4. Select the **Protocols > Channels** tab to display the configured channels.
5. To delete an existing channel, select it in the table and click **Delete**.
6. To reconfigure an existing channel:
 - a. Select the channel's name from the channel list (for example, the default `sip` channel).
 - b. Edit the *Listen Address* or *Listen Port* fields to correspond to the address of a NIC or logical address on the associated engine tier machine.

Note: The channel must be disabled before you can modify the listen address or listen port.

- c. Edit the External Listen Address or External Listen Port fields to match the public address of a load balancer in the system.
 - d. Edit the advanced channel attributes as necessary (see [Section 3.4.2, "Creating a New SIP or SIPS Channel"](#) for details.)
7. Click **Save**.

3.4.2 Creating a New SIP or SIPS Channel

To create a new SIP or SIPS channel to the configuration of a Oracle WebLogic Server SIP Container instance:

1. Log in to the Administration Console for the Oracle WebLogic Server SIP Container domain you want to configure.
2. In the left pane, select the **Environment > Servers** tab.
3. In the right pane, select the name of the server you want to configure.

4. Select the **Protocols > Channels** tab to display the configured channels.
5. Click **New** to configure a new channel.
6. Fill in the new channel fields as follows:
 - **Name:** Enter an administrative name for this channel, such as *SIPS-Channel-eth0*.
 - **Protocol:** Select either *SIP* to support UDP and TCP transport, or *SIPS* to support TLS transport. Note that a SIP channel cannot support only UDP or only TCP transport on the configured port.
7. Click **Next**.
8. Fill in the new channel's addressing fields as follows:
 - **Listen Address:** Enter the IP address or DNS name for this channel. On a DNS machine, enter the exact IP address of the interface you want to configure, or a DNS name that maps to the exact IP address.
 - **Listen Port:** Enter the port number used to communicate through this channel. The combination of Listen Address and Listen Port must be unique across all channels configured for the server. SIP channels support both UDP and TCP transport on the configured port.
 - **External Listen Address** and **External Listen Port:** Edit these fields to match the public address of a load balancer associated with this channel. When the server selects an interface or logical address to use for outbound network traffic, Oracle WebLogic Server SIP Container examines the channel that was configured with the same primary *Listen Address*; if the *External Listen Address* of this channel differs, the external address is embedded into SIP message headers for further call traffic. See [Section 3.2, "Configuring Load Balancer Addresses"](#).
9. Click **Next**.
10. Optionally click **Show** to display and edit advanced channel properties, such as *connection timeout* values. Keep in mind the following restrictions and suggestions for advanced channel properties:
 - **Outbound Enabled**—This attribute cannot be unchecked, because all SIP and SIPS channels can originate network connections.
 - **HTTP Enabled for This Protocol**—This attribute cannot be selected for SIP and SIPS channels, because Oracle WebLogic Server SIP Container does not support HTTP transport SIP protocols.
 - **Maximum Message Size**—This attribute specifies the maximum TCP message size that the server allows on a connection from this channel. Oracle WebLogic Server SIP Container shuts off any connection where the messages size exceeds the configured value. The default size of 10,000,000 bytes is large. If you are concerned about preventing Denial Of Service (DOS) attacks against the server, reduce this attribute to a value that is compatible with your deployed services.
11. Click **Finish**.

3.4.3 Configuring Custom Timeout, MTU, and Other Properties

SIP channels can be further configured using one or more custom channel properties. The custom properties cannot be set using the Administration Console. Instead, you

must use a text editor to add the properties to a single, `custom-property` stanza in the channel configuration portion of the `config.xml` file for the domain.

Oracle WebLogic Server SIP Container provides the following custom properties that affect the transport protocol of SIP channels:

- `TcpConnectTimeoutMillis`—Specifies the amount of time Oracle WebLogic Server SIP Container waits before it declares a destination address (for an outbound TCP connection) as unreachable. The property is applicable only to SIP channels; Oracle WebLogic Server SIP Container ignores this attribute value for SIPS channels. A value of 0 disables the timeout completely. A default value of *3000 milliseconds* is used if you do not specify the custom property.
- `SctpConnectTimeoutMillis`—Specifies the amount of time Oracle WebLogic Server SIP Container waits before it declares a destination address (for an outbound SCTP connection) as unreachable. The property is applicable only to SCTP channels (for Diameter traffic). A value of 0 disables the timeout completely. A default value of *3000 milliseconds* is used if you do not specify the custom property. See [Section 3.8.1.1, "Static Port Configuration for Outbound UDP Packets"](#) for information about creating SCTP channels for Diameter.
- `SourcePorts`—Configures one or more static port numbers that a server uses for originating UDP packets.

Caution: Oracle does not recommend using the `SourcePorts` custom property in most configurations because it degrades performance. Configure the property only in cases where you must specify the exact ports that Oracle WebLogic Server SIP Container uses to originate UDP packets.

- `Mtu`—Specifies the Maximum Transmission Unit (MTU) value for this channel. A value of -1 uses the default MTU size for the transport.
- `EnabledProtocolVersions`—Specifies the version of the SSL protocol to use with this channel when Oracle WebLogic Server SIP Container acts as an SSL client. When acting as an SSL client, by default the channel requires TLS V1.0 as the supported protocol. The server can be configured to use SSL V3.0 as well, if that is the highest version that the SSL peer servers support. You can set one of the following values for this property:
 - `TLS1`, the default, configures the channel to send and accept only TLS V1.0 messages. Peers must respond with a TLS V1.0 message, or the SSL connection is dropped.
 - `SSL3` configures the channel to send and accept only SSL V3.0 messages. Peers must respond with an SSL V3.0 message, or the SSL connection is dropped.
 - `ALL` supports either TLS V1.0 or SSL V3.0 messages. Peers must respond with a TLS V1.0 or SSL V3.0 message, or the SSL connection is dropped.

To configure a custom property, use a text editor to modify the `config.xml` file directly, or use a JMX client such as WLST to add the custom property. When editing `config.xml` directly, ensure that you add only one `custom-properties` element to the end of a channel's configuration stanza. Separate multiple custom properties within the same element using semicolons (;) as shown in [Example 3–1](#).

Example 3–1 Setting Custom Properties

```
<network-access-point>
```



```

<name>sip</name>
<protocol>sip</protocol>
<listen-port>5060</listen-port>
<public-port>5060</public-port>
<http-enabled-for-this-protocol>>false</http-enabled-for-this-protocol>
<tunneling-enabled>>false</tunneling-enabled>
<outbound-enabled>>true</outbound-enabled>
<enabled>>true</enabled>
<two-way-ssl-enabled>>false</two-way-ssl-enabled>
<client-certificate-enforced>>false</client-certificate-enforced>

<custom-properties>EnabledProtocolVersions=ALL;Mtu=1000;SourcePorts=5060</custom-p
roperties>
</network-access-point>

```

3.4.4 Configuring SIP Channels for DNS Machines

If you are configuring a server that has multiple network interfaces (a *DNS* server), you must configure a separate network channel for each IP address used by Oracle WebLogic Server SIP Container. Oracle WebLogic Server SIP Container uses the listen address and listen port values for each channel when embedding routing information into SIP message system headers.

Note: If you do not configure a channel for a particular IP address on a DNS machine, that IP address cannot be used when populating Via, Contact, and Record-Route headers.

3.5 Configuring TCP and TLS Channels for Diameter Support

The Oracle WebLogic Server SIP Container Diameter implementation supports the Diameter protocol over the TCP or TLS transport protocols. To enable incoming Diameter connections on a server, you configure a dedicated network channel using the protocol type *diameter* for TCP transport, or *diameters* for both TCP and TLS transport. The Diameter implementation application may automatically upgrade Diameter connections to use TLS as described in the Diameter specification (RFC 3558).

See [Chapter 6, "Configuring Diameter Client Nodes and Relay Agents"](#) for more information about configuring network channels for Diameter protocol support.

3.6 Configuring Engine Servers to Listen on Any IP Interface

To configure Oracle WebLogic Server SIP Container to listen for UDP traffic on any available IP interface, create a new SIP channel and specify 0.0.0.0 (or :: for IPv6 networks) as the listen address. Note that you must still configure at least one additional channel with an explicit IP address to use for outgoing SIP messages. (For DNS machines, each interface used for outgoing messages must have a configured channel.)

Note: If you configure a SIP channel without specifying the channel listen address, but you do configure a listen address for the server itself, then the SIP channel inherits the server listen address. In this case the SIP channel *does not* listen on IP_ANY.

Note: Using the 0.0.0.0 configuration affects only UDP traffic on Linux platforms. Oracle WebLogic Server SIP Container only creates TCP and HTTP listen threads corresponding to the configured host name of the server, and localhost. If multiple addresses are mapped to the host name, Oracle WebLogic Server SIP Container displays warning messages upon startup. To avoid this problem and listen on all addresses, specify the :: address, which encompasses all available addresses for both IPv6 and IPv4 for HTTP and TCP traffic as well.

3.7 Configuring Unique Listen Address Attributes for SIP Data Tier Replicas

Each replica in the SIP data tier must bind to a unique Listen Address attribute (a unique DNS name or IP address) in order to contact one another as peers. Follow these instructions for each replica to assign a unique Listen Address:

1. Access the Administration Console for the Oracle WebLogic Server SIP Container domain.
2. Select **Environment > Servers** from the left pane.
3. In the right pane, select the name of the server to configure.
4. Select the **Configuration > General** tab.
5. Enter a unique DNS name or IP address in the *Listen Address* field.
6. Click **Save**.

3.8 Production Network Architectures and Configuration


Most production (Enterprise Deployment) installations of Oracle WebLogic Server SIP Container contain one or more of the following characteristics:

- Multiple engine tier servers arranged in a cluster.
- Multiple network channels per engine tier server instance, in support of multiple SIP transport protocols or multiple Network Interface Cards (NICs) on DNS hardware.
- One or more load balancers, or a DNS load balancer, performing server failover and possibly Network Address Translation (NAT) for source or destination network packets.

A combination of these network elements can make it difficult to understand how elements interact with one another, and how a particular combination of elements or configuration options affects the contents of a SIP message or transport protocol datagram.

The sections that follow attempt to describe common Oracle WebLogic Server SIP Container network architectures and explain how servers are configured in each architecture. The sections also explain how information in SIP messages and transport datagrams is affected by each configuration. [Figure 3-1](#) shows the typical Open Systems Interconnect (OSI) model layers that can be affected by different network configurations.

Figure 3–1 OSI Layers Affected by Oracle WebLogic Server SIP Container Network Configuration



Layer 3 (Network) and Layer 4 (Transport) contain the source or destination IP address and port numbers for both outgoing and incoming transport datagrams. Layer 7 (Application) may also be affected because the SIP protocol specifies that certain SIP headers include addressing information for contacting the sender of a SIP message.

3.8.1 Single-NIC Configurations with TCP and UDP Channels

In a simple network configuration for a server having a single NIC, one or more network channels may be created to support SIP messages over UDP and TCP, or SIPs over TLS. It is helpful to understand how this simple configuration affects information in the OSI model, so that you can understand how more complex configurations involving DNS hardware and load balancers affect the same information.

Figure 3–2 Single-NIC Network Channel Configuration




Figure 3–2 shows a single engine tier server instance with a single NIC. The server is configured with one network channel supporting SIP over UDP and TCP. (SIP channels always support both UDP and TCP transports; you cannot support only one of the two.) Figure 3–2 also shows two clients communicating with the server, one over UDP and one over TCP.

For the TCP transport, the outgoing datagram (delivered from Oracle WebLogic Server SIP Container to the UA) contains the following information:

- Layer 3 includes the source IP address specified by the network channel (10.1.1.10 in the example above).
- Layer 4 includes the source port number allocated by the underlying operating system.

Incoming TCP datagrams (delivered from the UA to Oracle WebLogic Server SIP Container) contain the following information:

- Layer 3 includes the destination IP address specified by the network channel (10.1.1.10).
- Layer 4 contains the destination port number specified by the network channel (5060).

For outgoing UDP datagrams, the OSI layer information contains the same information as with TCP transport. For incoming UDP datagrams, the OSI layer information is the same as TCP except in the case of incoming datagram Layer 4 information. For incoming UDP datagrams, Layer 4 contains either:

- The destination port number specified by the network channel (5060), or
- The ephemeral port number previously allocated by Oracle WebLogic Server SIP Container.

By default Oracle WebLogic Server SIP Container allocates ports from the ephemeral port number range of the underlying operating system for outgoing UDP datagrams. Oracle WebLogic Server SIP Container allows external connections to use an ephemeral point as the destination port number, in addition to the port number configured in the network channel. In other words, Oracle WebLogic Server SIP Container automatically listens on all ephemeral ports that the server allocates. You can optionally disable Oracle WebLogic Server SIP Container's use of ephemeral port numbers by specifying the following option when starting the server:

```
-Dwlss.udp.listen.on.ephemeral=false
```

You can determine Oracle WebLogic Server SIP Container's use of a particular ephemeral port by examining the server log file:

```
<Nov 30, 2005 12:00:00 AM PDT> <Notice> <WebLogicServer> <BEA-000202> <Thread "SIP Message Processor (Transport UDP)" listening on port 35993.>
```

3.8.1.1 Static Port Configuration for Outbound UDP Packets

Oracle WebLogic Server SIP Container network channels provide a `SourcePorts` attribute that you can use to configure one or more static ports that a server uses for originating UDP packets.

Caution: Oracle does not recommend using the `SourcePorts` custom property in most configurations because it degrades performance. Configure the property only in cases where you must specify the exact ports that Oracle WebLogic Server SIP Container uses to originate UDP packets.

To configure the `SourcePorts` property, use a JMX client such as WLST or directly modify a network channel configuration in `config.xml` to include the custom property. `SourcePorts` defines an array of port numbers or port number ranges. Do not include spaces in the `SourcePorts` definition; use only port numbers, hyphens ("-") to designate ranges of ports, and commas (",") to separate ranges or individual ports. See [Example 3-2](#) for an example configuration.

Example 3–2 Static Port Configuration for Outgoing UDP Packets

```

<network-access-point>
  <name>sip</name>
  <protocol>sip</protocol>
  <listen-port>5060</listen-port>
  <public-port>5060</public-port>
  <http-enabled-for-this-protocol>>false</http-enabled-for-this-protocol>
  <tunneling-enabled>>false</tunneling-enabled>
  <outbound-enabled>>true</outbound-enabled>
  <enabled>>true</enabled>
  <two-way-ssl-enabled>>false</two-way-ssl-enabled>
  <client-certificate-enforced>>false</client-certificate-enforced>
  <custom-properties>SourcePorts=5060</custom-properties>
</network-access-point>

```

3.8.2 DNS Server Configurations Overview

Engine tier servers in a production deployment frequently utilize DNS server hardware, having two or more NICs. DNS hardware is typically used for one of the following purposes:

- To provide redundant network connections within the same subnet. Having multiple NICs ensures that one or more network connections are available to communicate with SIP data tier servers or the Administration Server, even if a single NIC fails.
- To support SIP communication across two or more different subnets. For example Oracle WebLogic Server SIP Container may be configured to proxy SIP requests from UAs in one subnet to UAs in a second subnet, when the UAs cannot directly communicate across subnets.

The configuration requirements and OSI layer information differ depending on the use of DNS hardware in your system. When multiple NICs are used to provide redundant connections within a subnet, servers are generally configured to listen on all available addresses (IP_ANY) as described in [Section 3.8.3, "DNS Servers Listening On All Addresses \(IP_ANY\)"](#).

When using multiple NICs to support different subnets, you must configure multiple network on the server for each different NIC as described in [Section 3.8.4, "DNS Servers Listening on Multiple Subnets"](#).


3.8.3 DNS Servers Listening On All Addresses (IP_ANY)

The simplest DNS configuration enables a Oracle WebLogic Server SIP Container instance to listen on all available NICs (physical NICs as well as logical NICs), sometimes described as IP_ANY. To accomplish this, you configure a single network channel and specify a channel listen address of 0.0.0.0.

Note that you must configure the 0.0.0.0 address directly on the server's network channel. If you specify no IP address in the channel, the channel inherits the listen address configured for the server instance itself. See [Section 3.6, "Configuring Engine Servers to Listen on Any IP Interface"](#).

3.8.4 DNS Servers Listening on Multiple Subnets

Multiple NICs can also be used in engine tier servers to listen on multiple subnets. The most common configuration uses Oracle WebLogic Server SIP Container to proxy SIP traffic from one subnet to another where no direct access between subnets is permitted. [Figure 3–3](#) shows this configuration.

Figure 3–3 DNS Configuration for Proxying between Subnets

To configure the Oracle WebLogic Server SIP Container instance in [Figure 3–3](#) you must define a separate network channel for each NIC used on the server machine. [Example 3–3](#) shows the `config.xml` entries that define channels for the sample configuration.

Example 3–3 Sample Network Channel Configuration for NICs on Multiple Subnets

```
<NetworkAccessPoint ListenAddress="10.1.1.10" ListenPort="5060" Name="sipchannelA"
Protocol="sip"/>
<NetworkAccessPoint ListenAddress="10.2.1.10" ListenPort="5060" Name="sipchannelB"
Protocol="sip"/>
```

3.8.4.1 Understanding the Route Resolver

When Oracle WebLogic Server SIP Container is configured to listen on multiple subnets, a feature called the *route resolver* is responsible for the following activities:

- Populating OSI Layer 7 information (SIP system headers such as Via, Contact, and so forth) with the correct address.
- Populating OSI Layer 3 information with the correct source IP address.

For example, in the configuration shown in [Figure 3–3](#), Oracle WebLogic Server SIP Container must add the correct subnet address to SIP system headers and transport datagrams in order for each UA to continue processing SIP transactions. If the wrong subnet is used, replies cannot be delivered because neither UA can directly access the other UA's subnet.

The route resolver works by determining which NIC the operating system will use to send a datagram to a given destination, and then examining the network channel that is associated with that NIC. It then uses the address configured in the selected network channel to populate SIP headers and Layer 3 address information.

For example, in the configuration shown in [Figure 3–3](#), an INVITE message sent from Oracle WebLogic Server SIP Container to UAC B would have a destination address of 10.2.1.16. The operating system would transmit this message using NIC B, which is configured for the corresponding subnet. The route resolver associates NIC B with the configured `sipchannelB` and embeds the channel's IP address (10.2.1.10) in the VIA header of the SIP message. UAC B then uses the VIA header to direct subsequent

messages to the server using the correct IP address. A similar process is used for UAC A, to ensure that messages are delivered only on the correct subnet.

3.8.4.2 IP Aliasing with DNS Hardware

IP aliasing assigns multiple logical IP addresses to a single NIC, and is configured in the underlying server operating system. If you configure IP aliasing and all logical IP addresses are within the same subnet, you can simply configure Oracle WebLogic Server SIP Container to listen on all addresses as described in [Section 3.8.3, "DNS Servers Listening On All Addresses \(IP_ANY\)"](#).

If you configure IP aliasing to create multiple logical IP addresses on different subnets, you must configure a separate network channel for each logical IP address. In this configuration, Oracle WebLogic Server SIP Container treats all logical addresses as separate physical interfaces (NICs) and uses the route resolver to populate OSI Layer 4 and Layer 7 information based on the configured channel.


3.8.5 Load Balancer Configurations

In addition to providing failover capabilities and distributing the client load across multiple servers, a load balancer is also an important tool for configuring the network information transmitted between clients and servers. The sections that follow describe common load balancer configurations used with Oracle WebLogic Server SIP Container.

3.8.5.1 Single Load Balancer Configuration

The most common load balancer configuration utilizes a single load balancer that gates access to a cluster of engine tier servers, as shown in [Figure 3-4](#).

Figure 3-4 Single Load Balancer Configuration



To configure Oracle WebLogic Server SIP Container for use with a single load balancer as in [Figure 3-4](#), configure one or more network channels for each server, and configure the public address of each channel with the Virtual IP address of the load balancer. In this configuration, Oracle WebLogic Server SIP Container embeds the load balancer IP address in SIP message system headers to ensure that clients can reach the

cluster for subsequent replies. [Chapter 3, "Managing Network Resources"](#) presents detailed steps for configuring network channels with load balancer addresses.

Note: Although some load balancing switches can automatically re-route all SIP messages in a given call to the same engine tier server, this functionality is not required with Oracle WebLogic Server SIP Container installations. See [Section 7.7, "Alternate Configurations"](#) for more information.

3.8.5.2 Multiple Load Balancers and DNS Load Balancers

Multiple load balancers (or a DNS load balancer) can be configured to provide several virtual IP addresses for a single Oracle WebLogic Server SIP Container cluster. To configure Oracle WebLogic Server SIP Container for use with a DNS load balancer, you create a dedicated network channel for each load balancer or local server NIC, and set the channel's public address to the virtual IP address of the appropriate load balancer. In this configuration, the route resolver associates a configured channel with the NIC used for originating SIP messages. The public address of the selected channel is then used for populating SIP system messages. See [Section 3.8.4.1, "Understanding the Route Resolver"](#).

3.8.5.3 Network Address Translation Options

In the most common case, a load balancer is configured using destination NAT to provide a public IP address that clients use for communicating with one or more internal (private) Oracle WebLogic Server SIP Container addresses. Load balancers may also be configured using source NAT, which modifies the Layer 3 address information originating from a private address to match the virtual IP address of the load balancer itself.

With the default route resolver behavior, an Oracle WebLogic Server SIP Container engine originates UDP packets having a source IP address that matches the address of a local NIC (the private address). This can be problematic for applications that try to respond directly to the Layer 3 address embedded in the transport packet, because the local server address may not be publicly accessible. If your applications exhibit this problem, Oracle recommends that you configure the load balancer to perform source NAT to change the transport Layer 3 address to a publicly-accessible virtual IP address.

3.8.5.3.1 IP Masquerading Alternative to Source NAT If you choose not to enable source NAT on your load balancer, Oracle WebLogic Server SIP Container provides limited IP masquerading functionality. To use this functionality, configure a logical address on each engine tier server using the public IP address of the load balancer for the cluster. (Duplicate the same logical IP address on each engine tier server machine). When a local server interface matches the IP address of a configured load balancer (defined in the public address of a network channel), Oracle WebLogic Server SIP Container uses that interface to originate SIP UDP messages, and the Layer 3 address contains a public address.

Caution: Using the Oracle WebLogic Server SIP Container IP masquerading functionality can lead to network instability, because it requires duplicate IP addresses on multiple servers. Production deployments must use a load balancer configured for source NAT, rather than IP masquerading, to ensure reliable network performance.

You can disable IP masquerading functionality by using the startup option:

```
-Dwlss.udp.lb.masquerade=false
```

3.9 Example Network Configuration

Oracle WebLogic Server SIP Container is compatible with load balancers that are not SIP-aware, meaning that they do not consider existing SIP dialogues when routing requests to servers. This document demonstrates load balancer and Oracle WebLogic Server SIP Container configuration, as well as SIP and Network Address Translation (NAT) interactions in various configurations.

The following sections describe a sample network configuration for Oracle WebLogic Server SIP Container using a non-SIP-aware load balancer:

For more information about implementation-dependent issues surrounding NAT see the IETF document, *NAT Behavioral Requirements for Unicast UDP* at <http://tools.ietf.org/wg/behave/draft-ietf-behave-nat-udp/>.


3.9.1 Example Network Topology

Figure 3–5 shows the sample network topology described in this section. An Oracle WebLogic Server SIP Container cluster, consisting of engines WLSS 1 and WLSS 2, is configured on private IP network 10.1/16 (an internal 16-bit subnet). The cluster's public IP address is 1.2.3.4, which is the virtual IP address configured on the load balancer.

The User Agent, UAC A, with IP address 2.3.4.5 never sees the internal IP addresses configured for the Oracle WebLogic Server SIP Container cluster. Instead, it sends requests to, and receives responses from 1.2.3.4.

The sections that follow discuss configuring the Oracle WebLogic Server SIP Container cluster and load balancer for this example system.

Figure 3–5 Example Network Topology



3.9.2 Oracle WebLogic Server SIP Container Configuration

The Oracle WebLogic Server SIP Container cluster configuration specifies the public address as 1.2.3.4, and the public port as 5060 (see [Section 3.2, "Configuring Load Balancer Addresses"](#)) for each engine. The default route on both Oracle WebLogic Server SIP Container engines points to the load balancer's 10.1/16 network interface: 10.1.3.4. The Oracle WebLogic Server SIP Container (servers WLSS 1 and WLSS 2) routing table is shown in [Example 3–4](#).

Example 3–4 Oracle WebLogic Server SIP Container Routing Table

```

$ /sbin/route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
10.1.0.0 * 255.255.0.0 U 0 0 0 eth0
default 10.1.3.4 0.0.0.0 UG 0 0

```

3.9.3 Load Balancer Configuration

The load balancer is configured with a virtual IP address of 1.2.3.4, and two real servers, WLSS 1 and WLSS 2, having addresses 10.1.1.1 and 10.1.1.2, respectively. The load balancer also has an internal IP address of 10.1.3.4 configured on the 10.1/16 network. The UAC address, 2.3.4.5, is reachable from the load balancer by static route configuration on the load balancer. The load balancer routing table is shown in [Example 3–5](#).

Example 3–5 Load balancer Routing Table

```

$ /sbin/route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
10.1.0.0 * 255.255.0.0 U 0 0 0 eth1
1.2.0.0 * 255.255.0.0 U 0 0

```

Because the SIP protocol specification (RFC 3261) dictates the destination IP address and UDP port numbers that user agents must use when sending requests or responses, the NAT configuration of the load balancer must be done in a way that does not violate RFC 3261 requirements. Three setup options can be used to accomplish this configuration:

- [Section 3.9.3.1, "NAT-based configuration"](#)
- [Section 3.9.3.2, "maddr-Based Configuration"](#)
- [Section 3.9.3.3, "rport-Based Configuration"](#)


The sections that follow describe each approach.

3.9.3.1 NAT-based configuration

The default UDP NAT behavior for load balancers is to perform destination IP address translation in the public > private network direction, and source IP address translation in the private > public network direction. This means setting up destination address translation in the UAC > Oracle WebLogic Server SIP Container (2.3.4.5 > 1.2.3.4) direction without source address translation, and source address translation in the Oracle WebLogic Server SIP Container > UAC (10.1/16 > 2.3.4.5) direction without destination address translation.

[Figure 3–6](#) illustrates the UDP packet flow for a SUBSCRIBE/200OK transaction.

Figure 3-6 SUBSCRIBE Sequence



Note that the source and destination IP addresses of the UDP packets are shown in blue. In the UAC > Oracle WebLogic Server SIP Container direction, the load balancer translates the destination IP address but not the source IP address. In the Oracle WebLogic Server SIP Container > UAC direction, the load balancer translates the source IP address but not the destination IP address.

The complete message trace (including IP and UDP headers, as well as the SIP payload) for the sequence from [Figure 3-6](#) is shown in [Example 3-6](#) below.

Example 3-6 Complete SUBSCRIBE Message Trace

No.	Time	Source	Destination	Protocol	Info
1	1.425250	2.3.4.5	1.2.3.4	SIP	Request: SUBSCRIBE sip:subscribe@1.2.3.4:5060

```

Internet Protocol, Src: 2.3.4.5 (2.3.4.5), Dst: 1.2.3.4 (1.2.3.4)
User Datagram Protocol, Src Port: 9999 (9999), Dst Port: sip (5060)
Session Initiation Protocol
Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0
Message Header
Via: SIP/2.0/UDP 2.3.4.5:9999;branch=1
From: sipp <sip:sipp@2.3.4.5>;tag=1
To: sut <sip:subscribe@1.2.3.4:5060>
Call-ID: 1-25923@2.3.4.5
Cseq: 1 SUBSCRIBE
Contact: sip:sipp@2.3.4.5:9999
Max-Forwards: 70
Event: ua-profile
Expires: 10
Content-Length: 0
    
```

No.	Time	Source	Destination	Protocol	Info
2	2.426250	2.3.4.5	10.1.1.1	SIP	Request: SUBSCRIBE sip:subscribe@1.2.3.4:5060

```

Internet Protocol, Src: 2.3.4.5 (2.3.4.5), Dst: 10.1.1.1 (10.1.1.1)
User Datagram Protocol, Src Port: 9999 (9999), Dst Port: sip (5060)
Session Initiation Protocol
Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0
Message Header
Via: SIP/2.0/UDP 2.3.4.5:9999;branch=1
From: sipp <sip:sipp@2.3.4.5>;tag=1
To: sut <sip:subscribe@1.2.3.4:5060>
    
```

```

Call-ID: 1-25923@2.3.4.5
Cseq: 1 SUBSCRIBE
Contact: sip:sipp@2.3.4.5:9999
Max-Forwards: 70
Event: ua-profile
Expires: 10
Content-Length: 0


No.      Time      Source      Destination      Protocol Info
3 3.430903 10.1.1.1    2.3.4.5          SIP           Status: 200
OK

Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 2.3.4.5 (2.3.4.5)
User Datagram Protocol, Src Port: 42316 (42316), Dst Port: 9999 (9999)
Session Initiation Protocol
  Status-Line: SIP/2.0 200 OK
  Message Header
    To: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
    Content-Length: 0
    Contact:
<sip:app-12eomtm5h5f77@1.2.3.4:5060;transport=udp;wlsscids=1ae4479ac6ff71>
  CSeq: 1 SUBSCRIBE
  Call-ID: 1-25923@2.3.4.5

```

If Oracle WebLogic Server SIP Container subsequently sends a NOTIFY request to the UAC, the sequence shown in [Figure 3-7](#) takes place:

Figure 3-7 NOTIFY Sequence



As in the previous sequence, the IP address translation takes place in the Oracle WebLogic Server SIP Container > UAC direction for the source IP address, and UAC > Oracle WebLogic Server SIP Container direction for the destination IP address.

Note that this setup does not require the load balancer to maintain session state information or to be SIP-aware. The complete message trace from [Figure 3-7](#) is shown in [Example 3-7](#) below.

Example 3-7 Complete NOTIFY Message Trace

```

No.      Time      Source      Destination      Protocol Info
1 5.430952 10.1.1.1    2.3.4.5          SIP           Request:
NOTIFY sip:sipp@2.3.4.5:9999

Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 2.3.4.5 (2.3.4.5)
User Datagram Protocol, Src Port: 42316 (42316), Dst Port: 9999 (9999)
Session Initiation Protocol

```

```

Request-Line: NOTIFY sip:sipp@2.3.4.5:9999 SIP/2.0
Message Header
  To: sipp <sip:sipp@2.3.4.5>;tag=1
  Content-Length: 0
  Contact:
<sip:app-12eomtm5h5f77@1.2.3.4:5060;transport=udp;wlsscid=1ae4479ac6ff71>
  CSeq: 1 NOTIFY
  Call-ID: 1-25923@2.3.4.5
  Via: SIP/2.0/UDP
1.2.3.4:5060;wlsscid=1ae4479ac6ff71;branch=z9hG4bKc5e4c3b4c22be517133ab749adeece4e
  From: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
  Max-Forwards: 70

```

No.	Time	Source	Destination	Protocol	Info
	2 6.430952	1.2.3.4	2.3.4.5	SIP	Request: NOTIFY

sip:sipp@2.3.4.5:9999

```

Internet Protocol, Src: 1.2.3.4 (1.2.3.4), Dst: 2.3.4.5 (2.3.4.5)
User Datagram Protocol, Src Port: 2222 (2222), Dst Port: 9999 (9999)
Session Initiation Protocol

```

```

Request-Line: NOTIFY sip:sipp@2.3.4.5:9999 SIP/2.0
Message Header
  To: sipp <sip:sipp@2.3.4.5>;tag=1
  Content-Length: 0
  Contact:
<sip:app-12eomtm5h5f77@1.2.3.4:5060;transport=udp;wlsscid=1ae4479ac6ff71>
  CSeq: 1 NOTIFY
  Call-ID: 1-25923@2.3.4.5
  Via: SIP/2.0/UDP
1.2.3.4:5060;wlsscid=1ae4479ac6ff71;branch=z9hG4bKc5e4c3b4c22be517133ab749adeece4e
  From: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
  Max-Forwards: 70

```

No.	Time	Source	Destination	Protocol	Info
	3 7.431367	2.3.4.5	1.2.3.4	SIP	Status: 200 OK

```

Internet Protocol, Src: 2.3.4.5 (2.3.4.5), Dst: 1.2.3.4 (1.2.3.4)
User Datagram Protocol, Src Port: 9999 (9999), Dst Port: sip (5060)
Session Initiation Protocol

```


```

Status-Line: SIP/2.0 200 OK
Message Header
  Via: SIP/2.0/UDP
1.2.3.4:5060;wlsscid=1ae4479ac6ff71;branch=z9hG4bKc5e4c3b4c22be517133ab749adeece4e
  From: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
  To: sipp <sip:sipp@2.3.4.5>;tag=1;tag=1
  Call-ID: 1-25923@2.3.4.5
  CSeq: 1 NOTIFY
  Contact: <sip:2.3.4.5:9999;transport=UDP>

```

Caution: If NAT is performed on both the source (SNAT) and destination IP addresses, the configuration does not work because the load balancer usually relies on a specific destination port number value to be sent in responses to requests. That port number value is dictated by RFC 3261, and must come from the Via header, which presents a conflict with load balancer's NAT requirements. RFC 3261 requires that responses to SIP requests be sent to the IP address used to send the request (unless maddr is present in the Via, as described in [Section 3.9.3.2, "maddr-Based Configuration"](#)). Consequently, in [Figure 3–8](#) below, Step 3, Oracle WebLogic Server SIP Container sends a 200 OK response to the load balancer internal IP address (10.1.3.4) and port 5060. That response is then dropped.

Figure 3–8 Source and Destination NAT



The complete message trace from [Figure 3–8](#) is show in [Example 3–8](#) below.

Example 3–8 Complete Failing SUBSCRIBE Message Trace

No.	Time	Source	Destination	Protocol	Info
1	1.425250	2.3.4.5	1.2.3.4	SIP	Request: SUBSCRIBE sip:subscribe@1.2.3.4:5060

```
Internet Protocol, Src: 2.3.4.5 (2.3.4.5), Dst: 1.2.3.4 (1.2.3.4)
User Datagram Protocol, Src Port: 9999 (9999), Dst Port: sip (5060)
Session Initiation Protocol
Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0
Message Header
Via: SIP/2.0/UDP 2.3.4.5:9999;branch=1
From: sipp <sip:sipp@2.3.4.5>;tag=1
To: sut <sip:subscribe@1.2.3.4:5060>
Call-ID: 1-25923@2.3.4.5
Cseq: 1 SUBSCRIBE
Contact: sip:sipp@2.3.4.5:9999
Max-Forwards: 70
Event: ua-profile
Expires: 10
Content-Length: 0
```

No.	Time	Source	Destination	Protocol	Info
2	2.426250	10.1.3.4	10.1.1.1	SIP	Request: SUBSCRIBE sip:subscribe@1.2.3.4:5060

```

Internet Protocol, Src: 10.1.3.4 (10.1.3.4), Dst: 10.1.1.1 (10.1.1.1)
User Datagram Protocol, Src Port: 2222 (2222), Dst Port: sip (5060)
Session Initiation Protocol
  Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0
  Message Header
    Via: SIP/2.0/UDP 2.3.4.5:9999;branch=1
    From: sipp <sip:sipp@2.3.4.5>;tag=1
    To: sut <sip:subscribe@1.2.3.4:5060>
    Call-ID: 1-25923@2.3.4.5
    Cseq: 1 SUBSCRIBE
    Contact: sip:sipp@2.3.4.5:9999
    Max-Forwards: 70
    Event: ua-profile
    Expires: 10
    Content-Length: 0
  
```

No.	Time	Source	Destination	Protocol	Info
3	3.430903	10.1.1.1	10.1.3.4	SIP	Status: 200 OK

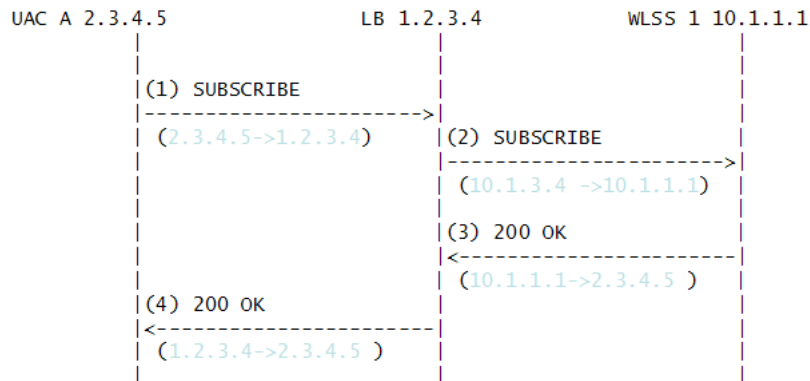
```

Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 10.1.3.4 (10.1.3.4)
User Datagram Protocol, Src Port: 42316 (42316), Dst Port: 9999 (9999)
Session Initiation Protocol
  
```

3.9.3.2 maddr-Based Configuration

When the maddr parameter is present in the Via header, the response is sent to the IP address specified in the maddr rather than to the received IP address (even when SNAT is enabled). In the example below, the UAC specifies a maddr set to 2.3.4.5 in the Via header. Consequently the response from the SIP server makes it to the UAC.

Figure 3–9 maddr Sequence



The complete message trace from [Figure 3–9](#) is shown in [Example 3–9](#) below.

Example 3–9 Complete maddr Message Trace

No.	Time	Source	Destination	Protocol	Info
1	1.425250	2.3.4.5	1.2.3.4	SIP	Request: SUBSCRIBE sip:subscribe@1.2.3.4:5060

```

Internet Protocol, Src: 2.3.4.5 (2.3.4.5), Dst: 1.2.3.4 (1.2.3.4)
User Datagram Protocol, Src Port: 9999 (9999), Dst Port: sip (5060)
Session Initiation Protocol
  
```

```
Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0
Message Header
  Via: SIP/2.0/UDP 2.3.4.5:9999;maddr=2.3.4.5;branch=1
  From: sipp <sip:sipp@2.3.4.5>;tag=1
  To: sut <sip:subscribe@1.2.3.4:5060>
  Call-ID: 1-25923@2.3.4.5
  Cseq: 1 SUBSCRIBE
  Contact: sip:sipp@2.3.4.5:9999
  Max-Forwards: 70
  Event: ua-profile
  Expires: 10
  Content-Length: 0
```

No.	Time	Source	Destination	Protocol	Info
	2 2.426250	10.1.3.4	10.1.1.1	SIP	Request:
SUBSCRIBE sip:subscribe@1.2.3.4:5060					

Internet Protocol, Src: 10.1.3.4 (10.1.3.4), Dst: 10.1.1.1 (10.1.1.1)
 User Datagram Protocol, Src Port: 2222 (2222), Dst Port: sip (5060)
 Session Initiation Protocol

```
Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0
Message Header
  Via: SIP/2.0/UDP 2.3.4.5:9999;maddr=2.3.4.5;branch=1
  From: sipp <sip:sipp@2.3.4.5>;tag=1
  To: sut <sip:subscribe@1.2.3.4:5060>
  Call-ID: 1-25923@2.3.4.5
  Cseq: 1 SUBSCRIBE
  Contact: sip:sipp@2.3.4.5:9999
  Max-Forwards: 70
  Event: ua-profile
  Expires: 10
  Content-Length: 0
```

No.	Time	Source	Destination	Protocol	Info
	3 3.430903	10.1.1.1	2.3.4.5	SIP	Status: 200
OK					

Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 2.3.4.5 (2.3.4.5)
 User Datagram Protocol, Src Port: 42316 (42316), Dst Port: 9999 (9999)
 Session Initiation Protocol

```
Status-Line: SIP/2.0 200 OK
```

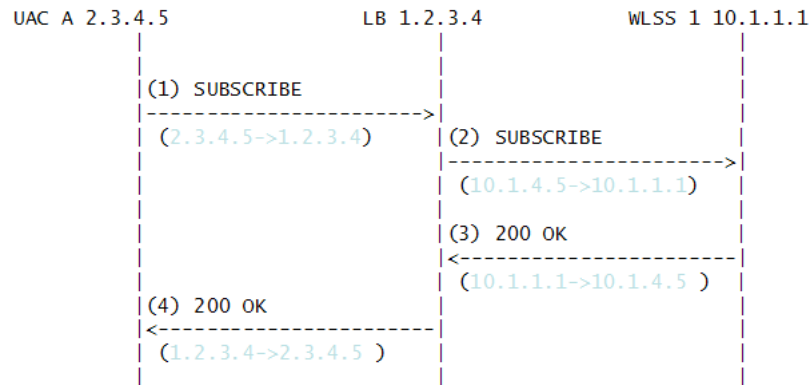
```
Message Header
```

```
To: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
Content-Length: 0
Contact:
```

```
<sip:app-12eomtm5h5f77@1.2.3.4:5060;transport=udp;wlsscrid=1ae4479ac6ff71>
```

3.9.3.3 rport-Based Configuration

RFC 3581 improves SIP and NAT interactions by allowing the client to request that the server send responses to a UDP port number from the request rather than from the Via. In order for both SUBSCRIBE and NOTIFY to work correctly, both the UAC as well as Oracle WebLogic Server SIP Container must support RFC 3581. [Figure 3-10](#) illustrates the SUBSCRIBE flow.

Figure 3–10 rport SUBSCRIBE Sequence

The complete message trace from [Figure 3–10](#) is shown in [Example 3–10](#) below.

Example 3–10 Complete Message Trace for rport SUBSCRIBE

No.	Time	Source	Destination	Protocol	Info
1	1.425250	2.3.4.5	1.2.3.4	SIP	Request: SUBSCRIBE sip:subscribe@1.2.3.4:5060

Internet Protocol, Src: 2.3.4.5 (2.3.4.5), Dst: 1.2.3.4 (1.2.3.4)
 User Datagram Protocol, Src Port: 9999 (9999), Dst Port: sip (5060)
 Session Initiation Protocol

Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0

Message Header

```

Via: SIP/2.0/UDP 2.3.4.5:9999;rport;branch=1
From: sipp <sip:sipp@2.3.4.5>;tag=1
To: sut <sip:subscribe@1.2.3.4:5060>
Call-ID: 1-25923@2.3.4.5
Cseq: 1 SUBSCRIBE
Contact: sip:sipp@2.3.4.5:9999
Max-Forwards: 70
Event: ua-profile
Expires: 10
Content-Length: 0
  
```

No.	Time	Source	Destination	Protocol	Info
2	2.426250	10.1.3.4	10.1.1.1	SIP	Request: SUBSCRIBE sip:subscribe@1.2.3.4:5060

Internet Protocol, Src: 10.1.3.4 (10.1.3.4), Dst: 10.1.1.1 (10.1.1.1)
 User Datagram Protocol, Src Port: 2222 (2222), Dst Port: sip (5060)
 Session Initiation Protocol

Request-Line: SUBSCRIBE sip:subscribe@1.2.3.4:5060 SIP/2.0

Message Header

```

Via: SIP/2.0/UDP 2.3.4.5:9999;rport;branch=1
From: sipp <sip:sipp@2.3.4.5>;tag=1
To: sut <sip:subscribe@1.2.3.4:5060>
Call-ID: 1-25923@2.3.4.5
Cseq: 1 SUBSCRIBE
Contact: sip:sipp@2.3.4.5:9999
Max-Forwards: 70
Event: ua-profile
Expires: 10
Content-Length: 0
  
```

```


No.      Time      Source      Destination      Protocol Info
3 3.430903 10.1.1.1    10.1.3.4         SIP      Status: 200
OK

Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 10.1.3.4 (10.1.3.4)
User Datagram Protocol, Src Port: 42316 (42316), Dst Port: 2222 (2222)
Session Initiation Protocol
  Status-Line: SIP/2.0 200 OK
  Message Header
    To: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
    Content-Length: 0
    Contact:
<sip:app-12eomt5h5f77@1.2.3.4:5060;transport=udp;wlsscid=1ae4479ac6ff71>
  CSeq: 1 SUBSCRIBE
  Call-ID: 1-25923@2.3.4.5
  
```

Figure 3–11 illustrates the NOTIFY flow.

Note that while source address NAT is enabled for both directions (**UAS > Oracle WebLogic Server SIP Container** and **Oracle WebLogic Server SIP Container > UA**), the load balancer can correctly identify the destination address in Step 3 by relying on receiving responses on the same port number as the one used to send requests. This implies that the load balancer maintains state.

Figure 3–11 *rport NOTIFY Sequence*



The complete message trace from Figure 3–11 is shown in Example 3–11 below.

Example 3–11 *Complete Message Trace for rport NOTIFY*

```

No.      Time      Source      Destination      Protocol Info
1 5.430952 10.1.1.1    2.3.4.5          SIP      Request:
NOTIFY sip:sipp@2.3.4.5:9999

Internet Protocol, Src: 10.1.1.1 (10.1.1.1), Dst: 2.3.4.5 (2.3.4.5)
User Datagram Protocol, Src Port: 42316 (42316), Dst Port: 9999 (9999)
Session Initiation Protocol
  Request-Line: NOTIFY sip:sipp@2.3.4.5:9999 SIP/2.0
  Message Header
    To: sipp <sip:sipp@2.3.4.5>;tag=1
    Content-Length: 0
    Contact:
<sip:app-12eomt5h5f77@1.2.3.4:5060;transport=udp;wlsscid=1ae4479ac6ff71>
  CSeq: 1 NOTIFY
  Call-ID: 1-25923@2.3.4.5
  Via: SIP/2.0/UDP
  
```

```
1.2.3.4:5060;wlsscid=1ae4479ac6ff71;branch=z9hG4bKc5e4c3b4c22be517133ab749adeece4e
;rport
```

```
From: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
Max-Forwards: 70
```

No.	Time	Source	Destination	Protocol	Info
2	6.430952	1.2.3.4	2.3.4.5	SIP	Request: NOTIFY sip:sipp@2.3.4.5:9999

```
Internet Protocol, Src: 1.2.3.4 (1.2.3.4), Dst: 2.3.4.5 (2.3.4.5)
User Datagram Protocol, Src Port: 2222 (2222), Dst Port: 9999 (9999)
Session Initiation Protocol
```

```
Request-Line: NOTIFY sip:sipp@2.3.4.5:9999 SIP/2.0
```

```
Message Header
```

```
To: sipp <sip:sipp@2.3.4.5>;tag=1
Content-Length: 0
Contact:
```

```
<sip:app-12eomtm5h5f77@1.2.3.4:5060;transport=udp;wlsscid=1ae4479ac6ff71>
```

```
CSeq: 1 NOTIFY
Call-ID: 1-25923@2.3.4.5
Via: SIP/2.0/UDP
```

```
1.2.3.4:5060;wlsscid=1ae4479ac6ff71;branch=z9hG4bKc5e4c3b4c22be517133ab749adeece4e
;rport
```

```
From: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
Max-Forwards: 70
```

No.	Time	Source	Destination	Protocol	Info
3	7.431367	2.3.4.5	1.2.3.4	SIP	Status: 200 OK

```
Internet Protocol, Src: 2.3.4.5 (2.3.4.5), Dst: 1.2.3.4 (1.2.3.4)
User Datagram Protocol, Src Port: 9999 (9999), Dst Port: (2222)
Session Initiation Protocol
```

```
Status-Line: SIP/2.0 200 OK
```

```
Message Header
```

```
Via: SIP/2.0/UDP
```

```
1.2.3.4:5060;wlsscid=1ae4479ac6ff71;branch=z9hG4bKc5e4c3b4c22be517133ab749adeece4e
;rport
```

```
From: sut <sip:subscribe@1.2.3.4:5060>;tag=82722c03
To: sipp <sip:sipp@2.3.4.5>;tag=1;tag=1
Call-ID: 1-25923@2.3.4.5
CSeq: 1 NOTIFY
Contact: <sip:2.3.4.5:9999;transport=UDP
```

Configuring SIP Data Tier Partitions and Replicas

The following sections describe how to configure Oracle WebLogic Server SIP Container instances that make up the SIP data tier cluster of a deployment:

- [Section 4.1, "Overview of SIP Data Tier Configuration"](#)
- [Section 4.2, "Best Practices for Configuring and Managing SIP Data Tier Servers"](#)
- [Section 4.3, "Example SIP Data Tier Configurations and Configuration Files"](#)
- [Section 4.4, "Storing Long-Lived Call State Data In A RDBMS"](#)
- [Section 4.5, "Introducing Geo-Redundancy"](#)
- [Section 4.6, "Using Geographically-Redundant SIP Data Tiers"](#)
- [Section 4.7, "Caching SIP Data in the Engine Tier"](#)
- [Section 4.8, "Monitoring and Troubleshooting SIP Data Tier Servers"](#)

4.1 Overview of SIP Data Tier Configuration

The Oracle WebLogic Server SIP Container SIP data tier is a cluster of server instances that manages the application call state for concurrent SIP calls. The SIP data tier may manage a single copy of the call state or multiple copies as needed to ensure that call state data is not lost if a server machine fails or network connections are interrupted.

The SIP data tier cluster is arranged into one or more *partitions*. A partition consists of one or more SIP data tier server instances that manage the same portion of concurrent call state data. In a single-server Oracle WebLogic Server SIP Container installation, or in a two-server installation where one server resides in the engine tier and one resides in the SIP data tier, all call state data is maintained in a single partition. Multiple partitions are required when the size of the concurrent call state exceeds the maximum size that can be managed by a single server instance. When more than one partition is used, the concurrent call state is split among the partitions, and each partition manages an separate portion of the data. For example, with a two-partition SIP data tier, one partition manages the call state for half of the concurrent calls (for example, calls A through M) while the second partition manages the remaining calls (N through Z).

In most cases, the maximum call state size that can be managed by an individual server corresponds to the Java Virtual Machine limit of approximately 1.6GB per server.

Additional servers can be added within the same partition to manage copies of the call state data. When multiple servers are members of the same partition, each server

manages a copy of the same portion of the call data, referred to as a *replica* of the call state. If a server in a partition fails or cannot be contacted due to a network failure, another replica in the partition supplies the call state data to the engine tier. Oracle recommends configuring two servers in each partition for production installations, to guard against machine or network failures. A partition can have a maximum of three replicas for providing additional redundancy.

4.1.1 `datatier.xml` Configuration File

The `datatier.xml` configuration file, located in the `config/custom` subdirectory of the domain directory, identifies SIP data tier servers and also defines the partitions and replicas used to manage the call state. If a server's name is present in `datatier.xml`, that server loads Oracle WebLogic Server SIP Container SIP data tier functionality at boot time. (Server names that do not appear in `datatier.xml` act as engine tier nodes, and instead provide SIP Servlet container functionality configured by the `sipserver.xml` configuration file.)

The sections that follow show examples of the `datatier.xml` contents for common SIP data tier configurations.

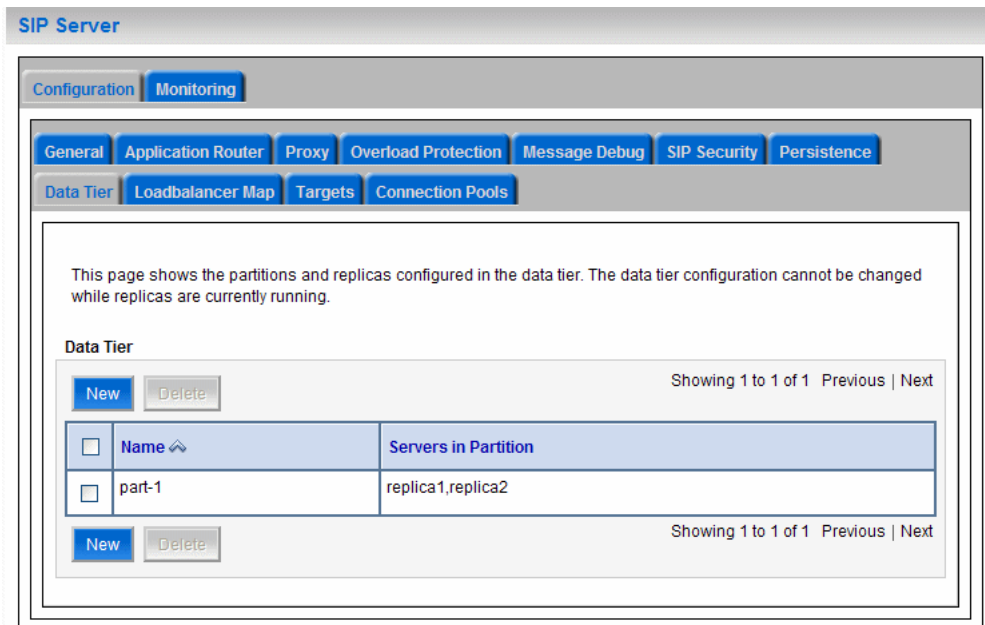
4.1.2 Configuration Requirements and Restrictions

All servers that participate in the SIP data tier should be members of the same WebLogic Server cluster. The cluster configuration enables each server to monitor the status of other servers. Using a cluster also enables you to easily target the `sipserver` and `datatier` custom resources to all servers for deployment.

For high reliability, you can configure up to three replicas within a partition.

You cannot change the SIP data tier configuration while replicas or engine tier nodes are running. You must restart servers in the domain in order to change SIP data tier membership or reconfigure partitions or replicas.

You can view the current SIP data tier configuration (and configure the data tier) using the Configuration > Data Tier page (SipServer node) of the Administration Console, as shown in [Figure 4-1](#).

Figure 4–1 Administration Console Display of SIP Data Tier Configuration (Read-Only)

4.2 Best Practices for Configuring and Managing SIP Data Tier Servers

Adding replicas can increase reliability for the system as a whole, but keep in mind that each additional server in a partition requires additional network bandwidth to manage the replicated data. With three replicas in a partition, each transaction that modifies the call state updates data on three different servers.

To ensure high reliability when using replicas, always ensure that server instances in the same partition reside on different machines. Hosting two or more replicas on the same machine leaves all of the hosted replicas vulnerable to a machine or network failure.

SIP data tier servers can have one of three different statuses:

- **ONLINE**—indicates that the server is available for managing call state transactions.
- **OFFLINE**—indicates that the server is shut down or unavailable.
- **ONLINE_LOCK_AUTHORITY_ONLY**—indicates that the server was rebooted and is currently being updated (from other replicas) with the current call state data. A recovering server cannot yet process call state transactions, because it does not maintain a full copy of the call state managed by the partition.

If you need to take a SIP data tier server instance offline for scheduled maintenance, make sure that at least one other server in the same partition is *active*. If you shut down an *active* server and all other servers in the partition are *offline* or *recovering*, you will lose a portion of the active call state.

Oracle WebLogic Server SIP Container automatically divides the call state evenly over all configured partitions.

4.3 Example SIP Data Tier Configurations and Configuration Files

The sections that follow describe some common Oracle WebLogic Server SIP Container installations that utilize a separate SIP data tier.

4.3.1 SIP Data Tier with One Partition

A single-partition, single-server SIP data tier represents the simplest data tier configuration. [Example 4-1](#) shows a SIP data tier configuration for a single-server deployment.

Example 4-1 SIP Data Tier Configuration for Small Deployment

```
<?xml version="1.0" encoding="UTF-8"?>
<data-tier xmlns="http://www.bea.com/ns/wlcp/wlss/300">
  <partition>
    <name>part-1</name>
    <server-name>replica1</server-name>
  </partition>
</data-tier>
```

To add a replica to an existing partition, simply define a second `server-name` entry in the same partition. For example, the `datatier.xml` configuration file shown in [Example 4-2](#) recreates a two-replica configuration.

Example 4-2 SIP Data Tier Configuration for Small Deployment with Replication

```
<?xml version="1.0" encoding="UTF-8"?>
<data-tier xmlns="http://www.bea.com/ns/wlcp/wlss/300">
  <partition>
    <name>Partition0</name>
    <server-name>DataNode0-0</server-name>
    <server-name>DataNode0-1</server-name>
  </partition>
</data-tier>
```

4.3.2 SIP Data Tier with Two Partitions

Multiple partitions can be easily created by defining multiple `partition` entries in `datatier.xml`, as shown in [Example 4-3](#).

Example 4-3 Two-Partition SIP Data Tier Configuration

```
<?xml version="1.0" encoding="UTF-8"?>
<data-tier xmlns="http://www.bea.com/ns/wlcp/wlss/300">
  <partition>
    <name>Partition0</name>
    <server-name>DataNode0-0</server-name>
  </partition>
  <partition>
    <name>Partition1</name>
    <server-name>DataNode1-0</server-name>
  </partition>
</data-tier>
```

4.3.3 SIP Data Tier with Two Partitions and Two Replicas

Replicas of the call state can be added by defining multiple SIP data tier servers in each partition. [Example 4-4](#) shows the `datatier.xml` configuration file used to define a system having two partitions with two servers (replicas) in each partition.

Example 4-4 SIP Data Tier Configuration for Small Deployment

```
<?xml version="1.0" encoding="UTF-8"?>
<data-tier xmlns="http://www.bea.com/ns/wlcp/wlss/300">
```



```

<partition>
  <name>Partition0</name>
  <server-name>DataNode0-0</server-name>
  <server-name>DataNode0-1</server-name>
</partition>
<partition>
  <name>Partition1</name>
  <server-name>DataNode1-0</server-name>
  <server-name>DataNode1-1</server-name>
</partition>
</data-tier>

```

4.4 Storing Long-Lived Call State Data In A RDBMS

Oracle WebLogic Server SIP Container enables you to store long-lived call state data in an Oracle or MySQL RDBMS in order to conserve RAM. When you enable RDBMS persistence, by default the SIP data tier persists a call state's data to the RDBMS after the call dialog has been established, and at subsequent dialog boundaries, retrieving or deleting the persisted call state data as necessary to modify or remove the call state.

Oracle also provides an API for application designers to provide "hints" as to when the SIP data tier should persist call state data. These hints can be used to persist call state data to the RDBMS more frequently, or to disable persistence for certain calls.

Note that Oracle WebLogic Server SIP Container only uses the RDBMS to supplement the SIP data tier's in-memory replication functionality. To improve latency performance when using an RDBMS, the SIP data tier maintains SIP timers in memory, along with call states being actively modified (for example, in response to a new call being set up). Call states are automatically persisted only after a dialog has been established and a call is in progress, at subsequent dialog boundaries, or in response to persistence hints added by the application developer.

When used in conjunction with an RDBMS, the SIP data tier selects one replica server instance to process all call state writes (or deletes) to the database. Any available replica can be used to retrieve call states from the persistent store as necessary for subsequent reads.

RDBMS call state storage can be used in combination with an engine tier cache, if your domain uses a SIP-aware load balancer to manage connections to the engine tier. See [Section 4.7, "Caching SIP Data in the Engine Tier"](#).

4.4.1 Requirements and Restrictions

Enable RDBMS call state storage only when all of the following criteria are met:

- The call states managed by your system are typically long-lived.
- The size of the call state to be stored is large. Very large call states may require a significant amount of RAM in order to store the call state.
- Latency performance is not critical to your deployed applications.

The latency requirement, in particular, must be well understood before choosing to store call state data in an RDBMS. The RDBMS call state storage option measurably increases latency for SIP message processing, as compared to using a SIP data tier cluster. If your system must handle a large number of short-lived SIP transactions with brief response times, Oracle recommends storing all call state data in the SIP data tier.

Note: RDBMS persistence is designed only to reduce the RAM requirements in the SIP data tier for large, long-lived call states. The persisted data cannot be used to restore a failed SIP data tier partition or replica.

4.4.2 Steps for Enabling RDBMS Call State Storage

In order to use the RDBMS call state storage feature, your Oracle WebLogic Server SIP Container domain must include the necessary JDBC configuration, SIP Servlet container configuration, and a database having the schema required to store the call state. You can automate much of the required configuration by using the Configuration Wizard to set up a new domain with the RDBMS call state template. See [Section 4.4.3, "Using the Configuration Wizard RDBMS Store Template"](#).

If you have an existing Oracle WebLogic Server SIP Container domain, or you want to configure the RDBMS store on your own, see [Section 4.4.4, "Configuring RDBMS Call State Storage by Hand"](#) for instructions to configure JDBC and Oracle WebLogic Server SIP Container to use an RDBMS store.

4.4.3 Using the Configuration Wizard RDBMS Store Template

The Configuration Wizard provides a simple template that helps you easily begin using and testing the RDBMS call state store. Follow these steps to create a new domain from the template:

1. Start the Configuration Wizard application (config.sh)
2. Accept the default selection, *Create a new WebLogic domain*, and click **Next**.
3. Select *Base this domain on an existing template*, and click **Browse** to display the Select a Template dialog.
4. Select the template named `replicateddomain.jar`, and click **OK**.
5. Click **Next**.
6. Enter the username and password for the Administrator of the new domain, and click **Next**.
7. Select a JDK to use, and click **Next**.
8. Select **No** to keep the settings defined in the source template file, and click **Next**.
9. Click **Create** to create the domain.

The template creates a new domain with two engine tier servers in a cluster, two SIP data tier servers in a cluster, and an Administration Server (AdminServer). The engine tier cluster includes the following resources and configuration:

- A JDBC datasource, `wlss.callstate.datasource`, required for storing long-lived call state data. Note that you must modify this configuration to configure the datasource for your own RDBMS server. See [Section 4.4.3.1, "Modify the JDBC Datasource Connection Information"](#).
 - A persistence configuration (shown in the SipServer node, Configuration > Persistence tab of the Administration Console) that defines default handling of persistence hints for both RDBMS and geographical redundancy.
10. Click **Done** to exit the configuration wizard.
 11. Follow the steps under [Section 4.4.3.1, "Modify the JDBC Datasource Connection Information"](#) to create the necessary tables in your RDBMS.

12. Follow the steps under [Section 4.4.4.3, "Create the Database Schema"](#) to create the necessary tables in your RDBMS.

4.4.3.1 Modify the JDBC Datasource Connection Information

After installing the new domain, modify the template JDBC datasource to include connection information for your RDBMS server:

1. Use your browser to access the URL `http://address:port/console` where *address* is the Administration Server's listen address and *port* is the listen port.
2. From the scrolling list in the Domain Structures panel (left side of the page), select `Services > Data Sources`.
3. Select the data source named `wlss.callstate.datasource` in the right pane.
4. Select the `Configuration > Connection Pool` tab in the right pane.
5. Modify the following connection pool properties:
 - **URL:** Modify the URL to specify the host name and port number of your RDBMS server.
 - **Properties:** Modify the value of the user, portNumber, SID, and serverName properties to match the connection information for your RDBMS.
 - **Password and Confirm Password:** Enter the password of the RDBMS user you specified.
6. Click **Save** to save your changes.
7. Select the **Targets** tab in the right pane.
8. On the Select Targets page, select the name of your SIP data tier cluster (for example, `BEA_DATA_TIER_CLUST`), then click **Save**.
9. Click **Save**.
10. Follow the steps under [Section 4.4.4.3, "Create the Database Schema"](#) to create the necessary tables in your RDBMS.

4.4.4 Configuring RDBMS Call State Storage by Hand

To change an existing Oracle WebLogic Server SIP Container domain to store call state data in an Oracle or MySQL RDBMS, you must configure the required JDBC datasource, edit the Oracle WebLogic Server SIP Container configuration, and add the required schema to your database. Follow the instructions in the sections below to configure an Oracle Database.

4.4.4.1 Configure JDBC Resources

Follow these steps to create the required JDBC resources in your domain:

1. Boot the Administration Server for the domain if it is not already running.
2. Access the Administration Console for the domain.
3. From the scrolling list in the Domain Structures panel (left side of the page), select `Services > Data Sources`.
4. Click **New** to create a new data source.
5. Fill in the fields of the Create a New JDBC Data Source page as follows:
 - **Name:** Enter `wlss.callstate.datasource`

- **JNDI Name:** Enter `wlss.callstate.datasource`.
 - **Database Type:** Select "Oracle."
 - **Database Driver:** Select an appropriate JDBC driver from the Database Driver list. Note that some of the drivers listed in this field may not be installed by default on your system. Install third-party drivers as necessary using the instructions from your RDBMS vendor.
6. Click **Next**.
 7. Fill in the fields of the Connection Properties tab using connection information for the database you want to use. Click **Next** to continue.
 8. Click **Test Configuration** to test your connection to the RDBMS, or click **Next** to continue.
 9. On the **Select Targets** page, select the name of your SIP data tier cluster (for example, `BEA_DATA_TIER_CLUST`).
 10. Click **Finish** to save your changes.

4.4.4.2 Configure Oracle WebLogic Server SIP Container Persistence Options

Follow these steps to configure the Oracle WebLogic Server SIP Container persistence options to use an RDBMS call state store:

1. Boot the Administration Server for the domain if it is not already running.
2. Access the Administration Console for the domain.
3. Select the `SipServer` node in the left pane.
4. Select the **Configuration > Persistence** tab in the right pane.
5. In the **Default Handling** drop-down menu, select either "db" or "all." It is acceptable to select "all" because geographically-redundant replication is only performed if the `Geo Site ID` and `Geo Remote T3 URL` fields have been configured.
6. Click **Save** to save your changes.

4.4.4.3 Create the Database Schema

Oracle WebLogic Server SIP Container includes a SQL script, `callstate.sql`, that you can use to create the tables necessary for storing call state information. The script is installed to the `user_staged_config` subdirectory of the domain directory when you configure a replicated domain using the Configuration Wizard. The script is also available in the `WLSS_HOME/common/templates/scripts/db/oracle` directory.

The contents of the `callstate.sql` SQL script are shown in [Example 4-5](#).

Example 4-5 `callstate.sql` Script for Call State Storage Schema

```
drop table callstate;

create table callstate (
  key1 int,
  key2 int,
  bytes blob default empty_blob(),
  constraint pk_callstate primary key (key1, key2)
);
```

Follow these steps to execute the script commands using SQL*Plus:

1. Move to the Oracle WebLogic Server SIP Container `utils` directory, in which the SQL Script is stored:

```
cd ~/bea/wlcserver_10.3/common/templates/scripts/db/oracle
```

2. Start the SQL*Plus application, connecting to the Oracle database in which you will create the required tables. Use the same username, password, and connect to the same database that you specified when configuring the JDBC driver in [Section 4.4.4.1, "Configure JDBC Resources"](#). For example:

```
sqlplus username/password@connect_identifier
where connect_identifier connects to the database identified in the JDBC
connection pool.
```

3. Execute the Oracle WebLogic Server SIP Container SQL script, `callstate.sql`:

```
START callstate.sql
```

4. Exit SQL*Plus:

```
EXIT
```

4.4.5 Using Persistence Hints in SIP Applications

Oracle WebLogic Server SIP Container provides a simple API to provide "hints" as to when the SIP data tier should persist call state data. You can use the API to disable persistence for specific calls or SIP requests, or to persist data more frequently than the default setting (at SIP dialog boundaries).

To use the API, simply obtain a `WlssSipApplicationSession` instance and use the `setPersist` method to enable or disable persistence. Note that you can enable or disable persistence either to an RDBMS store, or to as geographically-redundant Oracle WebLogic Server SIP Container installation (see [Section 4.6, "Using Geographically-Redundant SIP Data Tiers"](#)).

For example, some SIP-aware load balancing products use the SIP `OPTIONS` message to determine if a SIP Server is active. To avoid persisting these messages to an RDBMS and to a geographically-redundant site, a Servlet might implement a `doOptions` method to echo the request and turn off persistence for the message, as shown in [Example 4-6](#).

Example 4-6 Disabling RDBMS Persistence for Option Methods

```
protected void doOptions(SipServletRequest req) throws IOException {
    WlssSipApplicationSession session =
        (WlssSipApplicationSession) req.getApplicationSession();
    session.setPersist(WlssSipApplicationSession.PersistenceType.DATABASE,
        false);
    session.setPersist(WlssSipApplicationSession.PersistenceType.GEO_REDUNDANCY,
        false);
    req.createResponse(200).send();
}
```

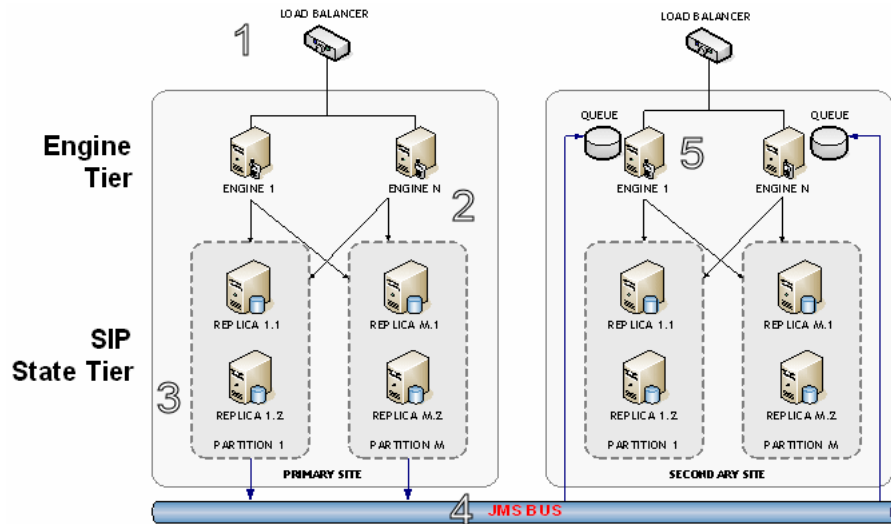
4.5 Introducing Geo-Redundancy

Geo-Redundancy ensures uninterrupted transactions and communications for providers, using geographically-separated SIP server deployments.

A primary site can process various SIP transactions and communications and upon determining a transaction boundary, replicate the state data associated with the

transaction being processed, to a secondary site. Upon failure of the primary site, calls are routed from the failed primary site to a secondary site for processing. Similarly, upon recovery, the calls are re-routed back to the primary site.

Figure 4–2 Geo-Redundancy



In the preceding figure, Geo-Redundancy is portrayed. The process proceeds in this manner:

1. Call is initiated on a primary OCCAS Cluster site, call setup and processing occurs normally.
2. Call is replicated as usual to the site's SIP State Tier, and becomes eligible for replication to a secondary site.
3. A single replica in the SIP State Tier then places the call state data to be replicated on a JMS queue configured.
4. Call is transmitted to one of the available engines using JMS over WAN.
5. Engines at the secondary site monitor their local queue for new messages. Upon receiving a message, an Engine in the secondary site OCCAS Cluster persists the call state data and assigns it the site ID value of the primary site.

Table 4–1 Geographic Redundancy flow

Normal Operation	failover
When a session is initiated on a primary OCCAS site, call setup and processing occurs normally.	Global LB policy updated to begin routing calls - primary site to secondary site.
When a SIP transaction boundary is reached, the call is replicated (in-memory) to the site's data tier, and becomes eligible for replication to a secondary site.	Once complete, the secondary site begins processing requests for the backed-up call state data.
A single replica in the data tier then places the call state data to be replicated on a JMS queue configured on the replica site.	When a requests hit secondary site engine retrieves the data and activates the call state, taking ownership for the call.
Data is transmitted to one of the available engines round-robin fashion.	Sets the site ID associated with the call to zero (making it appear local).
Engines at the secondary site monitor their local queue for new messages.	Activates all dormant timers present in the call state.

Table 4–1 (Cont.) Geographic Redundancy flow

Normal Operation	failover
Upon receiving a message, an engine on the secondary site persists the call state data and assigns it the site ID value of the primary site.	By default, call states are activated only for individual calls, and only after those calls are requested on the backup site.
The site ID distinguishes replicated call state data on the secondary site from any other call state data actively managed by the secondary site.	Servlets can use the <code>WlssipApplicationSession.getGeoSiteId()</code> method to examine the site ID associated with a call.
Timers in replicated call state data remain dormant on the secondary site, so that timer processing does not become a bottleneck to performance.	Any non-zero value for the site ID indicates that the Servlet is working with call state data that was replicated from another site.

4.5.1 Situations Best Suited to Use Geo-Redundancy

The following situations are best suited to take advantage of Geo-Redundancy:

- Your application uses SIP dialog states that are long-lived (dialog states that typically last 30 seconds or longer, such as SUBSCRIBE dialogs or conferences)
- Your application would reasonably be able to reconstruct the session (re-INVITE, expire SUBSCRIBE dialogs to trigger re-subscriptions, and so on) from the state that has been replicated
- The link between two OCCAS clusters or sites is low-bandwidth (<1Gb/s each direction) or high (or variable) latency (>5ms 95%)

4.5.2 Situations Not Suited to Use Geo-Redundancy

Geo-Redundancy should not be used in these situations:

- A high-capacity link between sites is available
- Your application does not reach SIP dialog steady-states that are likely to last longer than the time it would take to re-route all traffic to the secondary site in the event of catastrophic failure (15-30 seconds)
- If the application session is likely to be terminated by the user before the application could re-construct the session (most users will disconnect their calls before the session can be re-established from the secondary site)
- The volume of session state objects created by the application is greater than the site interconnect can support

4.5.3 Geo-Redundancy Considerations: Before Your Begin

Keep in mind the following considerations when planning your Geo-Redundancy:

- Dimension the system for the site link!
- Each dialog state is ~25KB on the wire (25600 bits)
- A typical B2BUA is two (2) dialogs
- Aim for 25% utilization (or less, depending on the specific equipment and topology of the site) to accommodate “jitter” and sustained latency on the link


For example, a 100 Mb/s link can handle approximately 1000 call states per second, and a typical B2BUA (in the default configuration) generates 4 states during the call (two for each dialog). So, a 100 Mb/s link will support a single OWLSC cluster dimensioned for a peak arrival rate (call rate) of 250 CPS.

- Geo-Redundancy is not *transparent* to the application; in most cases the application must be designed to use `SetPersist()` appropriately, and the developer must consider the volume of state that the application will queue for replication between sites
- `SetPersist()` should be used within the application code to selectively identify dialog states that will be long-lived
- Given the time it generally takes to route traffic to a secondary site, any application that replicates state more frequently will unnecessarily saturate the JMS queue and site interconnect
- Tuning of JMS to the specific application environment is required: Serialization options, message batching, reliable delivery options and queue size are all variable, depending on the specific application and site characteristics
- Geo-Redundancy default behavior is to replicate all dialog state changes when Geo-Redundancy is enabled for the container (this is *not* recommended for production deployments)
- Given the time it generally takes to route traffic to a secondary site, any application that replicates state more frequently will unnecessarily saturate the site interconnect
- `SetPersist()` should be used within the application code to selectively identify dialog states that will be long-lived (longer than ~20-30 seconds would be a reasonable threshold)

4.6 Using Geographically-Redundant SIP Data Tiers


The basic call state replication functionality available in the Oracle WebLogic Server SIP Container SIP data tier provides excellent failover capabilities for a single site installation. However, the active replication performed within the SIP data tier requires high network bandwidth in order to meet the latency performance needs of most production networks. This bandwidth requirement makes a single SIP data tier cluster unsuitable for replicating data over large distances, such as from one regional data center to another.

The Oracle WebLogic Server SIP Container geographic persistence feature enables you to replica call state transactions across multiple Oracle WebLogic Server SIP Container installations (multiple Administrative domains or "sites"). A geographically-redundant configuration minimizes dropped calls in the event of a catastrophic failure of an entire site, for example due to an extended, regional power outage.

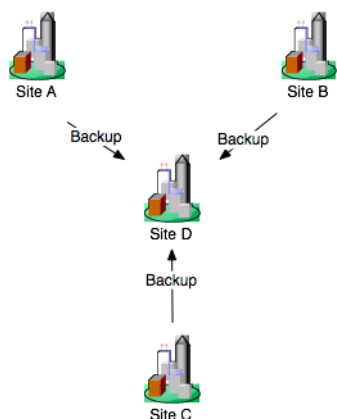
Figure 4–3 Oracle WebLogic Server SIP Container Geographic Persistence

4.6.1 Example Domain Configurations

A secondary Oracle WebLogic Server SIP Container domain that persists data from another domain may itself process SIP traffic, or it may exist solely as an active standby domain. In the most common configuration, two sites are configured to replicate each other's call state data, with each site processing its own local SIP traffic. The administrator can then use either domain as the "secondary" site should one of domains fail.

Figure 4–4 Common Geographically-Redundant Configuration

An alternate configuration utilizes a single domain that persists data from multiple, other sites, acting as the secondary for those sites. Although the secondary site in this configuration can also process its own, local SIP traffic, keep in mind that the resource requirements of the site may be considerable because of the need to persist active traffic from several other installations.

Figure 4–5 Alternate Geographically-Redundant Configuration

When using geographic persistence, a single replica in the primary site places modified call state data on a distributed JMS queue. By default, data is placed on the queue only at SIP dialog boundaries. (A custom API is provided for application developers that want to replicate data using a finer granularity, as described in [Section 4.4.5, "Using Persistence Hints in SIP Applications"](#).) In a secondary site, engine tier servers use a message listener to monitor the distributed queue to receive messages and write the data to its own SIP data tier cluster. If the secondary site uses an RDBMS to store long-lived call states (recommended), then all data writes from the distribute queue go directly to the RDBMS, rather than to the in-memory storage of the SIP data tier.

4.6.2 Requirements and Limitations

The Oracle WebLogic Server SIP Container geographically-redundant persistence feature is most useful for sites that manage long-lived call state data in an RDBMS. Short-lived calls may be lost in the transition to a secondary site, because Oracle WebLogic Server SIP Container may choose to collect data for multiple call states before replicating between sites.

You must have a reliable, site-aware load balancing solution that can partition calls between geographic locations, as well as monitor the health of a given regional site. Oracle WebLogic Server SIP Container provides no automated functionality for detecting the failure of an entire domain, or for failing over to a secondary site. It is the responsibility of the Administrator to determine when a given site has "failed," and to redirect that site's calls to the correct secondary site. Furthermore, the site-aware load balancer must direct all messages for a given callId to a single home site (the "active" site). If, after a failover, the failed site is restored, the load balancer must continue directing calls to the active site and not partition calls between the two sites.

During a failover to a secondary site, some calls may be dropped. This can occur because Oracle WebLogic Server SIP Container generally queues call state data for site replication only at SIP dialog boundaries. Failures that occur before the data is written to the queue result in the loss of the queued data.

Also, Oracle WebLogic Server SIP Container replicates call state data across sites only when a SIP dialog boundary changes the call state. If a long-running call exists on the primary site before the secondary site is started, and the call state remains unmodified, that call's data is not replicated to the secondary site. Should a failure occur before a long-running call state has been replicated, the call is lost during failover.

When planning for the capacity of a Oracle WebLogic Server SIP Container installation, keep in mind that, after a failover, a given site must be able to support all of the calls from the failed site as well as from its own geographic location. Essentially this means that all sites that are involved in a geographically-redundant configuration will operate at less than maximum capacity until a failover occurs.

4.6.3 Steps for Configuring Geographic Persistence

In order to use the Oracle WebLogic Server SIP Container geographic persistence features, you must perform certain configuration tasks on both the primary "home" site and on the secondary replication site.

Table 4–2 Steps for Configuring Geographic Persistence

Steps for Primary "Home" Site	Steps for Secondary "Replication" Site:
1. Install Oracle WebLogic Server SIP Container software and create replicated domain.	1. Install Oracle WebLogic Server SIP Container software and create replicated domain.
2. Enable RDBMS storage for long-lived call states (recommended).	2. Enable RDBMS storage for long-lived call states (recommended).
3. Configure persistence options to: define the unique regional site ID; identify the secondary site's URL; and enable replication hints.	3. Configure JMS Servers and modules required for replicating data.
	4. Configure persistence options to define the unique regional site ID.

Note: In most production deployments, two sites will perform replication services for each other, so you will generally configure each installation as both a primary and secondary site.

Oracle WebLogic Server SIP Container provides domain templates to automate the configuration of most of the resources described in [Table 4–2](#). See [Section 4.6.4, "Using the Configuration Wizard Templates for Geographic Persistence"](#) for information about using the templates.

If you have an existing Oracle WebLogic Server SIP Container domain and want to use geographic persistence, follow the instructions in [Section 4.6.5, "Manually Configuring Geographical Redundancy"](#) to create the resources.

4.6.4 Using the Configuration Wizard Templates for Geographic Persistence

Oracle WebLogic Server SIP Container provides two Configuration Wizard templates for using geographic persistence features:

- `WLSS_HOME/common/templates/domains/geo1domain.jar` configures a primary site having a site ID of 1. The domain replicates data to the engine tier servers created in `geo2domain.jar`.
- `WLSS_HOME/common/templates/domains/geo2domain.jar` configures a secondary site that replicates call state data from the domain created with `geo1domain.jar`. This installation has site ID of 2.

The server port numbers in both domain templates are unique, so you can test geographic persistence features on a single machine if necessary. Follow the instructions in the sections that follow to install and configure each domain.

4.6.4.1 Installing and Configuring the Primary Site

Follow these steps to create a new primary domain from the template:

1. Start the Configuration Wizard application (config.sh).
2. Accept the default selection, *Create a new WebLogic domain*, and click **Next**.
3. Select Base this domain on an existing template, and click Browse to display the Select a Template dialog.
4. Select the template named `geo1domain.jar`, and click **OK**.
5. Click **Next**.
6. Enter the username and password for the Administrator of the new domain, and click **Next**.
7. Select a JDK to use, and click **Next**.
8. Select No to keep the settings defined in the source template file, and click **Next**.
9. Click **Create** to create the domain.

The template creates a new domain with two engine tier servers in a cluster, two SIP data tier servers in a cluster, and an Administration Server (AdminServer). The engine tier cluster includes the following resources and configuration:

- A JDBC datasource, `wlss.callstate.datasource`, required for storing long-lived call state data. If you want to use this functionality, edit the datasource to include your RDBMS connection information as described in [Section 4.4.3.1, "Modify the JDBC Datasource Connection Information"](#).
 - A persistence configuration (shown in the SipServer node, Configuration > Persistence tab of the Administration Console) that defines:
 - Default handling of persistence hints for both RDBMS and geographic persistence.
 - A Geo Site ID of 1.
 - A Geo Remote T3 URL of `t3://localhost:8011,localhost:8061`, which identifies the engine tier servers in the "geo2" domain as the replication site for geographic redundancy.
10. Click Done to exit the configuration wizard.
 11. Follow the steps under [Section 4.6.4.2, "Installing the Secondary Site"](#) to create the domain that performs the replication.

4.6.4.2 Installing the Secondary Site

Follow these steps to use a template to create a secondary site from replicating call state data from the "geo1" domain:

1. Start the Configuration Wizard application (config.sh).
2. Accept the default selection, *Create a new WebLogic domain*, and click **Next**.
3. Select Base this domain on an existing template, and click Browse to display the Select a Template dialog.
4. Select the template named `geo2domain.jar`, and click **OK**.
5. Click **Next**.
6. Enter the username and password for the Administrator of the new domain, and click **Next**.

7. Select a JDK to use, and click **Next**.
8. Select No to keep the settings defined in the source template file, and click **Next**.
9. Click **Create** to create the domain.

The template creates a new domain with two engine tier servers in a cluster, two SIP data tier servers in a cluster, and an Administration Server (AdminServer). The engine tier cluster includes the following resources and configuration:

- A JDBC datasource, `wlss.callstate.datasource`, required for storing long-lived call state data. If you want to use this functionality, edit the datasource to include your RDBMS connection information as described in [Section 4.4.3.1, "Modify the JDBC Datasource Connection Information"](#).
- A persistence configuration (shown in the SipServer node, Configuration > Persistence tab of the Administration Console) that defines:
 - Default handling of persistence hints for both RDBMS and geographical redundancy.
 - A Geo Site ID of 2.
- A JMS system module, `SystemModule-Callstate`, that includes:
 - `ConnectionFactory-Callstate`, a connection factory required for backing up call state data from a primary site.
 - `DistributedQueue-Callstate`, a uniform distributed queue required for backing up call state data from a primary site.

The JMS system module is targeted to the site's engine tier cluster

- Two JMS Servers, `JMSServer-1` and `JMSServer-2`, are deployed to `engine1-site2` and `engine2-site2`, respectively.
10. Click **Done** to exit the configuration wizard.

4.6.5 Manually Configuring Geographical Redundancy

If you have an existing replicated Oracle WebLogic Server SIP Container installation, or pair of installations, you must create by hand the JMS and JDBC resources required for enabling geographical redundancy. You must also configure each site to perform replication. These basic steps for enabling geographical redundancy are:

1. **Configure JDBC Resources.** Oracle recommends configuring both the primary and secondary sites to store long-lived call state data in an RDBMS.
2. **Configure Persistence Options.** Persistence options must be configured on both the primary and secondary sites to enable engine tier hints to write to an RDBMS or to replicate data to a geographically-redundant installation.
3. **Configure JMS Resources.** A secondary site must have available JMS Servers and specific JMS module resources in order to replicate call state data from another site.

The sections that follow describe each step in detail.

4.6.5.1 Configuring JDBC Resources (Primary and Secondary Sites)

Follow the instructions in [Section 4.4, "Storing Long-Lived Call State Data In A RDBMS"](#) to configure the JDBC resources required for storing long-lived call states in an RDBMS.

4.6.5.2 Configuring Persistence Options (Primary and Secondary Sites)

Both the primary and secondary sites must configure the correct persistence settings in order to enable replication for geographical redundancy. Follow these steps to configure persistence:

1. Use your browser to access the URL `http://address:port/console` where *address* is the Administration Server's listen address and *port* is the listen port.
2. Click **Lock & Edit** to obtain a configuration lock.
3. Select the SipServer node in the left pane. The right pane of the console provides two levels of tabbed pages that are used for configuring and monitoring Oracle WebLogic Server SIP Container.
4. Select the **Configuration > Persistence** tab in the right pane.
5. Configure the Persistence attributes as follows:
 - **Default Handling:** Select "all" to persist long-lived call state data to an RDBMS and to replicate data to an external site for geographical redundancy (recommended). If your installation does not store call state data in an RDBMS, select "geo" instead of "all."
 - **Geo Site ID:** Enter a unique number from 1 to 9 to distinguish this site from all other configured sites. Note that the site ID of 0 is reserved to indicate call states that are local to the site in question (call states not replicated from another site).
 - **Geo Remote T3 URL:** For primary sites (or for secondary sites that replicate their own data to another site), enter the T3 URL or URLs of the engine tier servers that will replicate this site's call state data. If the secondary engine tier cluster uses a cluster address, you can enter a single T3 URL, such as `t3://mycluster:7001`. If the secondary engine tier cluster does not use a cluster address, enter the URLs for each individual engine tier server separated by a comma, such as `t3://engine1-east-coast:7001,t3://engine2-east-coast:7002,t3://engine3-east-coast:7001,t4://engine4-east-coast:7002`.
6. Click **Save** to save your configuration changes.
7. Click **Activate Changes** to apply your changes to the engine tier servers.

4.6.5.3 Configuring JMS Resources (Secondary Site Only)

Any site that replicates call state data from another site must configure certain required JMS resources. The resources are not required for sites that do not replicate data from another site.

Follow these steps to configure JMS resources:

1. Use your browser to access the URL `http://address:port/console` where *address* is the Administration Server's listen address and *port* is the listen port.
2. Click **Lock & Edit** to obtain a configuration lock.
3. Select the **Services > Messaging > JMS Servers** tab in the left pane.
4. Click **New** in the right pane.
5. Enter a unique name for the JMS Server or accept the default name. Click **Next** to continue.
6. In the Target list, select the name of a single engine tier server node in the installation. Click **Finish** to create the new Server.

7. Repeat Steps 3-6 for to create a dedicated JMS Server for each engine tier server node in your installation.
8. Select the **Services > Messaging > JMS Modules** node in the left pane.
9. Click **New** in the right pane.
10. Fill in the fields of the Create JMS System Module page as follows:
 - **Name:** Enter a name for the new module, or accept the default name.
 - **Descriptor File Name:** Enter the prefix a configuration file name in which to store the JMS module configuration (for example, `systemmodule-callstate`).
11. Click **Next** to continue.
12. Select the name of the engine tier cluster, and choose the option **All servers in the cluster**.
13. Click **Next** to continue.
14. Select **Would you like to add resources to this JMS system module** and click **Finish** to create the module.
15. Click **New** to add a new resource to the module.
16. Select the **Connection Factory** option and click **Next**.
17. Fill in the fields of the Create a new JMS System Module Resource as follows:
 - **Name:** Enter a descriptive name for the resource, such as `ConnectionFactory-Callstate`.
 - **JNDI Name:** Enter the name `wlss.callstate.backup.site.connection.factory`.
18. Click **Next** to continue.
19. Click **Finish** to save the new resource.
20. Select the name of the connection factory resource you just created.
21. Select the **Configuration > Load Balance** tab in the right pane.
22. De-select the **Server Affinity Enabled** option, and click **Save**.
23. Re-select the **Services > Messaging > JMS Modules** node in the left pane.
24. Select the name of the JMS module you created in the right pane.
25. Click **New** to create another JMS resource.
26. Select the **Distributed Queue** option and click **Next**.
27. Fill in the **Name** field of the Create a new JMS System Module Resource by entering a descriptive name for the resource, such as `DistributedQueue-Callstate`.
28. **JNDI Name:** Enter the name Fill in the fields of the Create a new JMS System Module Resource as follows:
 - **Name:** Enter a descriptive name for the resource, such as `ConnectionFactory-Callstate`.
 - **JNDI Name:** Enter the name `wlss.callstate.backup.site.queue`.
29. Click **Next** to continue.
30. Click **Finish** to save the new resource.

31. Click **Save** to save your configuration changes.
32. Click **Activate Changes** to apply your changes to the engine tier servers.

4.6.6 Understanding Geo-Redundant Replication Behavior

This section provides more detail into how multiple sites replicate call state data. Administrators can use this information to better understand the mechanics of geo-redundant replication and to better troubleshoot any problems that may occur in such a configuration. Note, however, that the internal workings of replication across Oracle WebLogic Server SIP Container installations is subject to change in future releases of the product.

4.6.6.1 Call State Replication Process

When a call is initiated on a primary Oracle WebLogic Server SIP Container site, call setup and processing occurs normally. When a SIP dialog boundary is reached, the call is replicated (in-memory) to the site's SIP data tier, and becomes eligible for replication to a secondary site. Oracle WebLogic Server SIP Container may choose to aggregate multiple call states for replication in order to optimize network usage.

A single replica in the SIP data tier then places the call state data to be replicated on a JMS queue configured on the replica site. Data is transmitted to one of the available engines (specified in the `geo-remote-t3-url` element in `sipserver.xml`) in a round-robin fashion. Engines at the secondary site monitor their local queue for new messages.

Upon receiving a message, an engine on the secondary site persists the call state data and assigns it the site ID value of the primary site. The site ID distinguishes replicated call state data on the secondary site from any other call state data actively managed by the secondary site. Timers in replicated call state data remain dormant on the secondary site, so that timer processing does not become a bottleneck to performance.

4.6.6.2 Call State Processing After Failover

To perform a failover, the Administrator must change a global load balancer policy to begin routing calls from the primary, failed site to the secondary site. After this process is completed, the secondary site begins processing requests for the backed-up call state data. When a request is made for data that has been replicated from the failed site, the engine retrieves the data and activates the call state, taking ownership for the call. The activation process involves:

- Setting the site ID associated with the call to zero (making it appear local).
- Activating all dormant timers present in the call state.

By default, call states are activated only for individual calls, and only after those calls are requested on the backup site. `SipServerRuntimeMBean` includes a method, `activateBackup(byte site)`, that can be used to force a site to take over all call state data that it has replicated from another site. The Administrator can execute this method using a WLST configuration script. Alternatively, an application deployed on the server can detect when a request for replicated site data occurs, and then execute the method. [Example 4-7](#) shows sample code from a JSP that activates a secondary site, changing ownership of all call state data replicated from site 1. Similar code could be used within a deployed Servlet. Note that either a JSP or Servlet must run as a privileged user in order to execute the `activateBackup` method.

In order to detect whether a particular call state request, Servlets can use the `WlssSipApplicationSession.getGeoSiteId()` method to examine the site ID

associated with a call. Any non-zero value for the site ID indicates that the Servlet is working with call state data that was replicated from another site.

Example 4–7 Activating a Secondary Site Using JMX

```
<%
  byte site = 1;

  InitialContext ctx = new InitialContext();
  MBeanServer server = (MBeanServer) ctx.lookup("java:comp/env/jmx/runtime");
  Set set = server.queryMBeans(new ObjectName("*:*",Type=SipServerRuntime"),
null);
  if (set.size() == 0) {
    throw new IllegalStateException("No MBeans Found!!!");
  }

  ObjectInstance oi = (ObjectInstance) set.iterator().next();
  SipServerRuntimeMBean bean = (SipServerRuntimeMBean)
    MBeanServerInvocationHandler.newProxyInstance(server,
      oi.getObjectInstance());

  bean.activateBackup(site);
%>
```

Note that after a failover, the load balancer must route all calls having the same callId to the newly-activated site. Even if the original, failed site is restored to service, the load balancer must not partition calls between the two geographical sites.

4.6.7 Removing Backup Call States

You may also choose to stop replicating call states to a remote site in order to perform maintenance on the remote site or to change the backup site entirely. Replication can be stopped by setting the **Site Handling** attribute to "none" on the primary site as described in [Section 4.6.5.2, "Configuring Persistence Options \(Primary and Secondary Sites\)"](#).

After disabling geographical replication on the primary site, you also may want to remove backup call states on the secondary site. `SipServerRuntimeMBean` includes a method, `deleteBackup(byte site)`, that can be used to force a site to remove all call state data that it has replicated from another site. The Administrator can execute this method using a WLST configuration script or via an application deployed on the secondary site. The steps for executing this method are similar to those for using the `activateBackup` method, described in [Section 4.6.6.2, "Call State Processing After Failover"](#).

4.6.8 Monitoring Replication Across Regional Sites

The `ReplicaRuntimeMBean` includes two new methods to retrieve data about geographically-redundant replication:

- `getBackupStoreOutboundStatistics()` provides information about the number of calls queued to a secondary site's JMS queue.
- `getBackupStoreInboundStatistics()` provides information about the call state data that a secondary site replicates from another site.

See *Oracle Fusion Middleware Communication Services Java API Reference* for more information about `ReplicaRuntimeMBean`.

4.6.9 Troubleshooting Geographical Replication

In addition to using the `ReplicaRuntimeMBean` methods described in [Section 4.6.8, "Monitoring Replication Across Regional Sites"](#), Administrators should monitor any SNMP traps that indicate failed database writes on a secondary site installation.

Administrators must also ensure that all sites participating in geographically-redundant configurations use unique site IDs.

4.7 Caching SIP Data in the Engine Tier

As described in [Chapter 7, "Oracle WebLogic Server SIP Container Base Platform Topologies"](#), in the default Oracle WebLogic Server SIP Container configuration the engine tier cluster is stateless. A separate SIP data tier cluster manages call state data in one or more partitions, and engine tier servers fetch and write data in the SIP data tier as necessary. Engines can write call state data to multiple replicas in each partition to provide automatic failover should a SIP data tier replica going offline.

Oracle WebLogic Server SIP Container also provides the option for engine tier servers to cache a portion of the call state data locally, as well as in the SIP data tier. When a local cache is used, an engine tier server first checks its local cache for existing call state data. If the cache contains the required data, and the local copy of the data is up-to-date (compared to the SIP data tier copy), the engine locks the call state in the SIP data tier but reads directly from its cache. This improves response time performance for the request, because the engine does not have to retrieve the call state data from a SIP data tier server.

The engine tier cache stores only the call state data that has been most recently used by engine tier servers. Call state data is moved into an engine's local cache as necessary in order to respond to client requests or to refresh out-of-date data. If the cache is full when a new call state must be written to the cache, the least-recently accessed call state entry is first removed from the cache. The size of the engine tier cache is not configurable.

Using a local cache is most beneficial when a SIP-aware load balancer manages requests to the engine tier cluster. With a SIP-aware load balancer, all of the requests for an established call are directed to the same engine tier server, which improves the effectiveness of the cache. If you do not use a SIP-aware load balancer, the effectiveness of the cache is limited, because subsequent requests for the same call may be distributed to different engine tier servers (having different cache contents).

4.7.1 Configuring Engine Tier Caching

Engine tier caching is enabled by default. To disable partial caching of call state data in the engine tier, specify the `engine-call-state-cache-enabled` element in `sipserver.xml`:

```
<engine-call-state-cache-enabled>>false</engine-call-state-cache-enabled>
```

When enabled, the cache size is fixed at a maximum of 250 call states. The size of the engine tier cache is not configurable.

4.7.2 Monitoring and Tuning Cache Performance

`SipPerformanceRuntime` monitors the behavior of the engine tier cache. [Table 4-3](#) describes the MBean attributes.

Table 4–3 SipPerformanceRuntime Attribute Summary


Attribute	Description
cacheRequests	Tracks the total number of requests for session data items.
cacheHits	The server increments this attribute each time a request for session data results in a version of that data being found in the engine tier server's local cache. Note that this counter is incremented even if the cached data is out-of-date and needs to be updated with data from the SIP data tier.
cacheValidHits	This attribute is incremented each time a request for session data is fully satisfied by a cached version of the data.

When enabled, the size of the cache is fixed at 250 call states. Because the cache consumes memory, you may need to modify the JVM settings used to run engine tier servers to meet your performance goals. Cached call states are maintained in the tenured store of the garbage collector. Try reducing the fixed "NewSize" value when the cache is enabled (for example, `-XX:MaxNewSize=32m -XX:NewSize=32m`). Note that the actual value depends on the call state size used by applications, as well as the size of the applications themselves.

4.8 Monitoring and Troubleshooting SIP Data Tier Servers

A runtime MBean, (`ReplicaRuntimeMBean`), provides valuable information about the current state and configuration of the SIP data tier. See *Oracle Fusion Middleware Communication Services Java API Reference* for a description of the attributes provided in this MBean.

Many of these attributes can be viewed using the SIP Servers Monitoring > Data Tier Information tab in the Administration Console, as shown in [Figure 4–6](#).

Figure 4–6 SIP Data Tier Monitoring in the Administration Console

[Example 4–8](#) shows a simple WLST session that queries the current attributes of a single Managed Server instance in a SIP data tier partition. [Table 4–1](#) describes the MBean services in more detail.

Example 4–8 Displaying ReplicaRuntimeMBean Attributes

```

connect('weblogic','weblogic','t3://datahost1:7001')
custom()
cd('com.bea')
cd('com.bea:ServerRuntime=replica1,Name=replica1,Type=ReplicaRuntime')
ls()
-rw- BackupStoreInboundStatistics          null
-rw- BackupStoreOutboundStatistics        null
-rw- BytesReceived                        0
-rw- BytesSent                            0
-rw- CurrentViewId                        2
-rw- DataItemCount                        0
-rw- DataItemsToRecover                   0
-rw- DatabaseStoreStatistics              null
-rw- HighKeyCount                         0
-rw- HighTotalBytes                       0
-rw- KeyCount                             0
-rw- Name                                  replica1
-rw- Parent                               com.bea:Name=replica1,Type=S
erverRuntime
-rw- PartitionId                          0
-rw- PartitionName                        part-1
-rw- ReplicaId                            0
-rw- ReplicaName                           replica1
-rw- ReplicaServersInCurrentView           java.lang.String[replica1,
replica2]
-rw- ReplicasInCurrentView                 [I@75378c
-rw- State                                  ONLINE
-rw- TimerQueueSize                        0
-rw- TotalBytes                            0
-rw- Type                                   ReplicaRuntime

```

Table 4–4 ReplicaRuntimeMBean Method and Attribute Summary

Method/Attribute	Description
dumpState()	Records the entire state of the selected SIP data tier server instance to the Oracle WebLogic Server SIP Container log file. You may want to use the dumpState() method to provide additional diagnostic information to a Technical Support representative in the event of a problem.
BackupStoreInboundStatistics	Provides statistics about call state data replicated from a remote geographical site.
BackupStoreOutboundStatistics	Provides statistics about call state data replicated to a remote geographical site.
BytesReceived	The total number of bytes received by this SIP data tier server. Bytes are received as servers in the engine tier provide call state data to be stored.
BytesSent	The total number of bytes sent from this SIP data tier server. Bytes are sent to engine tier servers when requested to provide the stored call state.

Table 4–4 (Cont.) ReplicaRuntimeMBean Method and Attribute Summary

Method/Attribute	Description
CurrentViewId	The current view ID. Each time the layout of the SIP data tier changes, the view ID is incremented. For example, as multiple servers in a SIP data tier cluster are started for the first time, the view ID is incremented when each server begins participating in the SIP data tier. Similarly, the view is incremented if a server is removed from the SIP data tier, either intentionally or due to a failure.
DataItemCount	The total number of stored call state keys for which this server has data. This attribute may be lower than the KeyCount attribute if the server is currently recovering data.
DataItemsToRecover	The total number of call state keys that must still be recovered from other replicas in the partition. A SIP data tier server may recover keys when it has been taken offline for maintenance and is then restarted to join the partition.
HighKeyCount	The highest total number of call state keys that have been managed by this server since the server was started.
HighTotalBytes	The highest total number of bytes occupied by call state data that this server has managed since the server was started.
KeyCount	The number of call data keys that are stored on the replica.
PartitionId	The numerical partition ID (from 0 to 7) of this server's partition.
PartitionName	The name of this server's partition.
ReplicaId	The numerical replica ID (from 0 to 2) of this server's replica.
ReplicaName	The name of this server's replica.
ReplicaServersInCurrentView	The names of other Oracle WebLogic Server SIP Container instances that are participating in the partition.
State	The current state of the replica. SIP data tier servers can have one of three different statuses: <ul style="list-style-type: none"> ■ ONLINE—indicates that the server is available for managing call state transactions. ■ OFFLINE—indicates that the server is shut down or unavailable. ■ ONLINE_LOCK_AUTHORITY_ONLY—indicates that the server was rebooted and is currently being updated (from other replicas) with the current call state data. A recovering server cannot yet process call state transactions, because it does not maintain a full copy of the call state managed by the partition.

Table 4–4 (Cont.) ReplicaRuntimeMBean Method and Attribute Summary

Method/Attribute	Description
TimerQueueSize	<p>The current number of timers queued on the SIP data tier server. This generally corresponds to the KeyCount value, but may be less if new call states are being added but their associated timers have not yet been queued.</p> <p>Note: Engine tier servers periodically check with SIP data tier instances to determine if timers associated with a call have expired. In order for SIP timers to function properly, all engine tier servers must actively synchronize their system clocks to a common time source. Oracle recommends using a Network Time Protocol (NTP) client or daemon on each engine tier instance and synchronizing to a selected NTP server. See Section 2.5, "Configuring Timer Processing".</p>
TotalBytes	<p>The total number of bytes consumed by the call state managed in this server.</p>

Monitoring and Troubleshooting

The following sections describe how to configure use the Oracle WebLogic Server SIP Container "echo server" process to improve SIP data tier failover performance when a server becomes physically disconnected from the network:

- [Section 5.1, "Avoiding and Recovering from Server Failures"](#)
- [Section 5.2, "Overview of Failover Detection"](#)
- [Section 5.3, "Improving Failover Performance for Physical Network Failures"](#)
- [Section 5.4, "Configuring SNMP"](#)
- [Section 5.5, "Understanding and Responding to SNMP Traps"](#)
- [Section 5.6, "Using the WebLogic Diagnostics Framework \(WLDF\)"](#)
- [Section 5.7, "Logging SIP Requests and Responses"](#)
- [Section 5.8, "Tuning JVM Garbage Collection for Production Deployments"](#)
- [Section 5.9, "Avoiding JVM Delays Caused By Random Number Generation"](#)

5.1 Avoiding and Recovering from Server Failures

A variety of events can lead to the failure of a server instance. Often one failure condition leads to another. Loss of power, hardware malfunction, operating system crashes, network partitions, or unexpected application behavior may each contribute to the failure of a server instance.

Oracle WebLogic Server SIP Container uses a highly clustered architecture as the basis for minimizing the impact of failure events. However, even in a clustered environment it is important to prepare for a sound recovery process in the event that an individual server or server machine fails.

The following sections summarize Oracle WebLogic Server SIP Container failure prevention and recovery features, and describe the configuration artifacts that are required in order to restore different portions of a Oracle WebLogic Server SIP Container domain

5.1.1 Failure Prevention and Automatic Recovery Features

Oracle WebLogic Server SIP Container, and the underlying WebLogic Server platform, provide many features that protect against server failures. In a production system, all available features should be used in order to ensure uninterrupted service.

5.1.1.1 Overload Protection

Oracle WebLogic Server SIP Container detects increases in system load that could affect the performance and stability of deployed SIP Servlets, and automatically throttles message processing at predefined load thresholds.

Using overload protection helps you avoid failures that could result from unanticipated levels of application traffic or resource utilization.

Oracle WebLogic Server SIP Container attempts to avoid failure when certain conditions occur:

- The rate at which SIP sessions are created reaches a configured value, or
- The size of the SIP timer and SIP request-processing execute queues reaches a configured length.

The underlying WebLogic Server platform also detects increases in system load that can affect deployed application performance and stability. WebLogic Server allows administrators to configure failure prevention actions that occur automatically at predefined load thresholds. Automatic overload protection helps you avoid failures that result from unanticipated levels of application traffic or resource utilization as indicated by:

- A workload manager's capacity being exceeded
- The HTTP session count increasing to a predefined threshold value
- Impending out of memory conditions

5.1.1.2 Redundancy and Failover for Clustered Services

You can increase the reliability and availability of your applications by using multiple engine tier servers in a dedicated cluster, as well as multiple SIP data tier servers (replicas) in a dedicated SIP data tier cluster. Because engine tier clusters maintain no stateful information about SIP dialogs (calls), the failure of an engine tier server does not result in any data loss or dropped calls. Multiple replicas in a SIP data tier partition store redundant copies of call state information, and automatically failover to one another should a replica fail.

5.1.1.3 Automatic Restart for Failed Server Instances

WebLogic Server self-health monitoring features improve the reliability and availability of server instances in a domain. Selected subsystems within each server instance monitor their health status based on criteria specific to the subsystem. (For example, the JMS subsystem monitors the condition of the JMS thread pool while the core server subsystem monitors default and user-defined execute queue statistics.) If an individual subsystem determines that it can no longer operate in a consistent and reliable manner, it registers its health state as "failed" with the host server.

Each WebLogic Server instance, in turn, checks the health state of its registered subsystems to determine its overall viability. If one or more of its critical subsystems have reached the FAILED state, the server instance marks its own health state FAILED to indicate that it cannot reliably host an application.

When used in combination with Node Manager, server self-health monitoring enables you to automatically reboot servers that have failed. This improves the overall reliability of a domain, and requires no direct intervention from an administrator.

5.1.1.4 Managed Server Independence Mode

Managed Servers maintain a local copy of the domain configuration. When a Managed Server starts, it contacts its Administration Server to retrieve any changes to the domain configuration that were made since the Managed Server was last shut down. If a Managed Server cannot connect to the Administration Server during startup, it can use its locally-cached configuration information—this is the configuration that was current at the time of the Managed Server's most recent shutdown. A Managed Server that starts up without contacting its Administration Server to check for configuration updates is running in *Managed Server Independence (MSI)* mode. By default, MSI mode is enabled.

5.1.1.5 Automatic Migration of Failed Managed Servers

When using Linux or UNIX operating systems, you can use WebLogic Server's server migration feature to automatically start a candidate (backup) server if a Network tier server's machine fails or becomes partitioned from the network. The server migration feature uses node manager, in conjunction with the `wlsifconfig.sh` script, to automatically boot candidate servers using a floating IP address. Candidate servers are booted only if the primary server hosting a Network tier instance becomes unreachable. See "Whole Server Migration" in *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server* documentation for more information about using the server migration feature.

5.1.1.6 Geographic Redundancy for Regional Site Failures

In addition to server-level redundancy and failover capabilities, Oracle WebLogic Server SIP Container enables you to configure peer sites to protect against catastrophic failures, such as power outages, that can affect an entire domain. This enables you to failover from one geographical site to another, avoiding complete service outages.

5.1.2 Directory and File Backups for Failure Recovery

Recovery from the failure of a server instance requires access to the domain's configuration data. By default, the Administration Server stores a domain's primary configuration data in a file called `domain_name/config/config.xml`, where `domain_name` is the root directory of the domain. The primary configuration file may reference additional configuration files for specific WebLogic Server services, such as JDBC and JMS, and for Oracle WebLogic Server SIP Container services, such as SIP container properties and SIP data tier configuration. The configuration for specific services are stored in additional XML files in subdirectories of the `domain_name/config` directory, such as `domain_name/config/jms`, `domain_name/config/jdbc`, and `domain_name/config/custom` for Oracle WebLogic Server SIP Container configuration files.

The Administration Server can automatically archive multiple versions of the domain configuration (the entire `domain-name/config` directory). The configuration archives can be used for system restoration in cases where accidental configuration changes need to be reversed. For example, if an administrator accidentally removes a configured resource, the prior configuration can be restored by using the last automated backup.

The Administration Server stores only a finite number of automated backups locally in `domain_name/config`. For this reason, automated domain backups are limited in their ability to guard against data corruption, such as a failed hard disk. Automated backups also do not preserve certain configuration data that are required for full domain restoration, such as LDAP repository data and server start-up scripts. Oracle

recommends that you also maintain multiple backup copies of the configuration and security offline, in a source control system.

This section describes file backups that Oracle WebLogic Server SIP Container performs automatically, as well as manual backup procedures that an administrator should perform periodically.

5.1.2.1 Enabling Automatic Configuration Backups

Follow these steps to enable automatic domain configuration backups on the Administration Server for your domain:

1. Access the Administration Console for your domain.
2. In the left pane of the Administration Console, select the name of the domain.
3. In the right pane, click the **Configuration > General** tab.
4. Select **Advanced** to display advanced options.
5. Select **Configuration Archive Enabled**.
6. In the Archive Configuration Count box, enter the maximum number of configuration file revisions to save.
7. Click **Save**.

When you enable configuration archiving, the Administration Server automatically creates a configuration JAR file archive. The JAR file contains a complete copy of the previous configuration (the complete contents of the `domain-name\config` directory). JAR file archive files are stored in the `domain-name\configArchive` directory. The files use the naming convention `config-number.jar`, where `number` is the sequential number of the archive.

When you save a change to a domain's configuration, the Administration Server saves the previous configuration in `domain-name\configArchive\config.xml#n`. Each time the Administration Server saves a file in the `configArchive` directory, it increments the value of the `#n` suffix, up to a configurable number of copies—5 by default. Thereafter, each time you change the domain configuration:

- The archived files are rotated so that the newest file has a suffix with the highest number,
- The previous archived files are renamed with a lower number, and
- The oldest file is deleted.

Keep in mind that configuration archives are stored locally within the domain directory, and they may be overwritten according to the maximum number of revisions you selected. For these reasons, you must also create your own off-line archives of the domain configuration, as described in [Section 5.1.2.2, "Storing the Domain Configuration Offline"](#).

5.1.2.2 Storing the Domain Configuration Offline

Although automatic backups protect against accidental configuration changes, they do not protect against data loss caused by a failure of the hard disk that stores the domain configuration, or accidental deletion of the domain directory. To protect against these failures, you must also store a complete copy of the domain configuration offline, preferably in a source control system.

Oracle recommends storing a copy of the domain configuration at regular intervals. For example, back up a new revision of the configuration when:

- you first deploy the production system
- you add or remove deployed applications
- the configuration is tuned for performance
- any other permanent change is made.

The domain configuration backup should contain the complete contents of the `domain_name/config` directory. For example:

```
cd ~/user_projects/domains/mydomain
tar cvf domain-backup-06-17-2007.jar config
```

Store the new archive in a source control system, preserving earlier versions should you need to restore the domain configuration to an earlier point in time.

5.1.2.3 Backing Up Server Start Scripts

In a Oracle WebLogic Server SIP Container deployment, the start scripts used to boot engine and SIP data tier servers are generally customized to include domain-specific configuration information such as:

- JVM Garbage Collection parameters required to achieve throughput targets for SIP message processing (see [Section 5.8, "Tuning JVM Garbage Collection for Production Deployments"](#)). Different parameters (and therefore, different start scripts) are generally used to boot engine and SIP data tier servers.
- Configuration parameters and startup information for the Oracle WebLogic Server SIP Container heartbeat mechanism. If you use the heartbeat mechanism, engine tier server start scripts should include startup options to enable and configure the heartbeat mechanism. SIP data tier server start scripts should include startup options to enable heartbeats and start the `WlsSEchoServer` process.

Backup each distinct start script used to boot engine tier, SIP data tier, or diameter relay servers in your domain.

5.1.2.4 Backing Up Logging Servlet Applications

If you use Oracle WebLogic Server SIP Container logging Servlets (see [Section 5.7, "Logging SIP Requests and Responses"](#)) to perform regular logging or auditing of SIP messages, backup the complete application source files so that you can easily redeploy the applications should the staging server fail or the original deployment directory becomes corrupted.

5.1.2.5 Backing Up Security Data

The WebLogic Security service stores its configuration data `config.xml` file, and also in an LDAP repository and other files.

5.1.2.5.1 Backing Up SerializedSystemIni.dat and Security Certificates All servers create a file named `SerializedSystemIni.dat` and place it in the server's root directory. This file contains encrypted security data that must be present to boot the server. You must back up this file.

If you configured a server to use SSL, also back up the security certificates and keys. The location of these files is user-configurable.

5.1.2.5.2 Backing Up the WebLogic LDAP Repository The default Authentication, Authorization, Role Mapper, and Credential Mapper providers that are installed with Oracle WebLogic Server SIP Container store their data in an LDAP server. Each Oracle

WebLogic Server SIP Container contains an embedded LDAP server. The Administration Server contains the master LDAP server, which is replicated on all Managed Servers. If any of your security realms use these installed providers, you should maintain an up-to-date backup of the following directory tree:

```
domain_name\adminServer\ldap
```

where `domain_name` is the domain's root directory and `adminServer` is the directory in which the Administration Server stores runtime and security data.

Each Oracle WebLogic Server SIP Container has an LDAP directory, but you only need to back up the LDAP data on the Administration Server—the master LDAP server replicates the LDAP data from each Managed Server when updates to security data are made. WebLogic security providers cannot modify security data while the domain's Administration Server is unavailable. The LDAP repositories on Managed Servers are replicas and cannot be modified.

The `ldap/ldapfiles` subdirectory contains the data files for the LDAP server. The files in this directory contain user, group, group membership, policies, and role information. Other subdirectories under the `ldap` directory contain LDAP server message logs and data about replicated LDAP servers.

Do not update the configuration of a security provider while a backup of LDAP data is in progress. If a change is made—for instance, if an administrator adds a user—while you are backing up the `ldap` directory tree, the backups in the `ldapfiles` subdirectory could become inconsistent. If this does occur, consistent, but potentially out-of-date, LDAP backups are available.

Once a day, a server suspends write operations and creates its own backup of the LDAP data. It archives this backup in a ZIP file below the `ldap\backup` directory and then resumes write operations. This backup is guaranteed to be consistent, but it might not contain the latest security data.

5.1.2.6 Backing Up Additional Operating System Configuration Files

Certain files maintained at the operating system level are also critical in helping you recover from system failures. Consider backing up the following information as necessary for your system:

- Load Balancer configuration scripts. For example, any automated scripts used to configure load balancer pools and virtual IP addresses for the engine tier cluster, as well as NAT configuration settings.
- NTP client configuration scripts used to synchronize the system clocks of engine and SIP data tier servers.
- Host configuration files for each Oracle WebLogic Server SIP Container machine (host names, virtual and real IP addresses for multi-homed machines, IP routing table information).

5.1.3 Restarting a Failed Administration Server

When you restart a failed Administration Server, no special steps are required. Start the Administration Server as you normally would.

If the Administration Server shuts down while Managed Servers continue to run, you do not need to restart the Managed Servers that are already running in order to recover management of the domain. The procedure for recovering management of an active domain depends upon whether you can restart the Administration Server on the same machine it was running on when the domain was started.

5.1.3.1 Restarting an Administration Server on the Same Machine

If you restart the WebLogic Administration Server while Managed Servers continue to run, by default the Administration Server can discover the presence of the running Managed Servers.

Note: Make sure that the startup command or startup script does not include `-Dweblogic.management.discover=false`, which disables an Administration Server from discovering its running Managed Servers.

The root directory for the domain contains a file, `running-managed-servers.xml`, which contains a list of the Managed Servers in the domain and describes whether they are running or not. When the Administration Server restarts, it checks this file to determine which Managed Servers were under its control before it stopped running.

When a Managed Server is gracefully or forcefully shut down, its status in `running-managed-servers.xml` is updated to "not-running". When an Administration Server restarts, it does not try to discover Managed Servers with the "not-running" status. A Managed Server that stops running because of a system crash, or that was stopped by killing the JVM or the command prompt (shell) in which it was running, will still have the status "running" in `running-managed-servers.xml`. The Administration Server will attempt to discover them, and will throw an exception when it determines that the Managed Server is no longer running.

Restarting the Administration Server does not cause Managed Servers to update the configuration of static attributes. *Static attributes* are those that a server refers to only during its startup process. Servers instances must be restarted to take account of changes to static configuration attributes. Discovery of the Managed Servers only enables the Administration Server to monitor the Managed Servers or make runtime changes in attributes that can be configured while a server is running (dynamic attributes).

5.1.3.2 Restarting an Administration Server on Another Machine

If a machine crash prevents you from restarting the Administration Server on the same machine, you can recover management of the running Managed Servers as follows:

1. Install the Oracle WebLogic Server SIP Container software on the new administration machine (if this has not already been done).
2. Make your application files available to the new Administration Server by copying them from backups or by using a shared disk. Your application files should be available in the same relative location on the new file system as on the file system of the original Administration Server.
3. Make your configuration and security data available to the new administration machine by copying them from backups or by using a shared disk. For more information, refer to [Section 5.1.2.2, "Storing the Domain Configuration Offline"](#) and [Section 5.1.2.5, "Backing Up Security Data"](#).
4. Restart the Administration Server on the new machine.

Make sure that the startup command or startup script does not include `-Dweblogic.management.discover=false`, which disables an Administration Server from discovering its running Managed Servers.

When the Administration Server starts, it communicates with the Managed Servers and informs them that the Administration Server is now running on a different IP address.

5.1.4 Restarting Failed Managed Servers

If the machine on which the failed Managed Server runs can contact the Administration Server for the domain, simply restart the Managed Server manually or automatically using Node Manager. Note that you must configure Node Manager and the Managed Server to support automated restarts.

If the Managed Server cannot connect to the Administration Server during startup, it can retrieve its configuration by reading locally-cached configuration data. A Managed Server that starts in this way is running in Managed Server Independence (MSI) mode.

To start up a Managed Server in MSI mode:

1. Ensure that the following files are available in the Managed Server's root directory:
 - `msi-config.xml`
 - `SerializedSystemIni.dat`
 - `boot.properties`

If these files are not in the Managed Server's root directory:

- a. Copy the `config.xml` and `SerializedSystemIni.dat` file from the Administration Server's root directory (or from a backup) to the Managed Server's root directory.
- b. Rename the configuration file to `msi-config.xml`. When you start the server, it will use the copied configuration files.

Note: Alternatively, use the `-Dweblogic.RootDirectory=path` startup option to specify a root directory that already contains these files.

2. Start the Managed Server at the command line or using a script.

The Managed Server will run in MSI mode until it is contacted by its Administration Server. For information about restarting the Administration Server in this scenario, see [Section 5.1.3, "Restarting a Failed Administration Server"](#).

5.2 Overview of Failover Detection

In a production system, engine tier servers continually access SIP data tier replicas in order to retrieve and write call state data. The Oracle WebLogic Server SIP Container architecture depends on engine tier nodes to detect when a SIP data tier server has failed or become disconnected. When an engine cannot access or write call state data because a replica is unavailable, the engine connects to another replica in the same partition and reports the offline server. The replica updates the current view of the SIP data tier to account for the offline server, and other engines are then notified of the updated view as they access and retrieve call state data.

By default, an engine tier server uses its RMI connection to the replica to determine if the replica has failed or become disconnected. The algorithms used to determine a

failure of an RMI connection are reliable, but ultimately they depend on the TCP protocol's retransmission timers to diagnose a disconnection (for example, if the network cable to the replica is removed). Because the TCP retransmission timer generally lasts a full minute or longer, Oracle WebLogic Server SIP Container provides an alternate method of detecting failures that can diagnose a disconnected replica in a matter of a few seconds.

5.2.1 WlssEchoServer Failure Detection

`WlssEchoServer` is a separate process that you can run on the same server hardware as a SIP data tier replica. The purpose of `WlssEchoServer` is to provide a simple UDP echo service to engine tier nodes to be used for determining when a SIP data tier server goes offline, for example in the event that the network cable is disconnected. The algorithm for detecting failures with `WlssEchoServer` is as follows:

1. For all normal traffic, engine tier servers communicate with SIP data tier replicas using TCP. TCP is used as the basic transport between the engine tier and SIP data tier regardless of whether or not `WlssEchoServer` is used.
2. Engine tier servers send a periodic heartbeat message to each configured `WlssEchoServer` over UDP. During normal operation, `WlssEchoServer` responds to the heartbeats so that the connection between the engine node and replica is verified.
3. Should there be a complete failure of the SIP data tier stack, or the network cable is disconnected, the heartbeat messages are not returned to the engine node. In this case, the engine node can mark the replica as being offline *without* having to wait for the normal TCP connection timeout.
4. After identifying the offline server, the engine node reports the failure to an available SIP data tier replica, and the SIP data tier view is updated as described in the previous section.

Also, should a SIP data tier server notice that its local `WlssEchoServer` process has died, it automatically shuts down. This behavior ensures even quicker failover because avoids the time it takes engine nodes to notice and report the failure as described in [Section 5.2, "Overview of Failover Detection"](#).

You can configure the heartbeat mechanism on engine tier servers to increase the performance of failover detection as necessary. You can also configure the listen port and log file that `WlssEchoServer` uses on SIP data tier servers.

5.2.2 Forced Shutdown for Failed Replicas

If any engine tier server cannot communicate with a particular replica, the engine access another, available replica in the SIP data tier to report the offline server. The replica updates its view of the affected partition to remove the offline server. The updated view is then distributed to all engine tier servers that later access the partition. Propagating the view in this manner helps to ensure that engine servers do not attempt to access the offline replica.

The replica that updates the view also issues a one-time request to the offline replica to ask it to shut down. This is done to try to shut-down running replica servers that cannot be accessed by one or more engine servers due to a network outage. If an active replica can reach the replica marked as "offline," the offline replica shuts down.

5.3 Improving Failover Performance for Physical Network Failures

Note: Using `WlssEchoServer` is not required in all Oracle WebLogic Server SIP Container installations. Enable the echo server only when your system requires detection of a network or replica failure faster than the configured TCP timeout interval.

Observe the following requirements and restrictions when using `WlssEchoServer` to detect replica failures:

- If you use the heartbeat mechanism to detect failures, you must ensure that the `WlssEchoServer` process is always running on each replica server machine. If the `WlssEchoServer` process fails or is stopped, the replica will be treated as being "offline" even if the server process is unaffected.
- Note that `WlssEchoServer` listens on all IP addresses available on the server machine.
- `WlssEchoServer` requires a dedicated port number to listen for heartbeat messages.

5.3.1 Starting `WlssEchoServer` on SIP Data Tier Server Machines

`WlssEchoServer` is a Java program that you can start directly from a shell or command prompt. The basic syntax for starting `WlssEchoServer` is:

```
java -classpath WLSS_HOME/server/lib/wlss/wlssechosvr.jar options
com.bea.wcp.util.WlssEchoServer
```

Where `WLSS_HOME` is the path to the Oracle WebLogic Server SIP Container installation and `options` may include one of the options described in [Table 5–1](#).

Table 5–1 *WlssEchoServer Options*

Option	Description
<code>-Dwlss.ha.echoserver.ipaddress</code>	Specifies the IP address on which the <code>WlssEchoServer</code> instance listens for heartbeat messages. If you do not specify an IP address, the instance listens on any available IP address (0.0.0.0).
<code>-Dwlss.ha.echoserver.port</code>	Specifies the port number used to listen for heartbeat messages. Ensure that the port number you specify is not used by any other process on the server machine. By default <code>WlssEchoServer</code> uses port 6734.
<code>-Dwlss.ha.echoserver.logfile</code>	Specifies the log file location and name. By default, log messages are written to <code>./echo_servertime.log</code> where <code>time</code> is the time expressed in milliseconds.

Oracle recommends that you include the command to start `WlssEchoServer` in the same script you use to start each Oracle WebLogic Server SIP Container SIP data tier instance. If you use the `startManagedWebLogic.sh` script to start an engine or SIP data tier server instance, add a command to start `WlssEchoServer` before the final command used to start the server. For example, change the lines:

```
"$JAVA_HOME/bin/java" ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS} \
-Dweblogic.Name=${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
```

to read:

```
"$JAVA_HOME/bin/java" -classpath WLSS_HOME/server/lib/wlss/wlssechosvr.jar \
-Dwlss.ha.echoserver.ipaddress=192.168.1.4 \
-Dwlss.ha.echoserver.port=6734 com.bea.wcp.util.WlssEchoServer &
"$JAVA_HOME/bin/java" ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS} \
-Dweblogic.Name=${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
```

5.3.2 Enabling and Configuring the Heartbeat Mechanism on Servers

To enable the `WlssEchoServer` heartbeat mechanism, you must include the `-Dreplica.host.monitor.enabled JVM` argument in the command you use to start all engine and SIP data tier servers. Oracle recommends adding this option directly to the script used to start Managed Servers in your system. For example, in the `startManagedWebLogic.sh` script, change the line:

```
# JAVA_OPTIONS="-Dweblogic.attribute=value -Djava.attribute=value"
```

to read:

```
JAVA_OPTIONS="-Dreplica.host.monitor.enabled=true"
```

Several additional JVM options configure the functioning of the heartbeat mechanism. [Table 5–1](#) describes the options used to configure failure detection.

Table 5–2 *WlssEchoServer Options*

Option	Description
<code>-Dreplica.host.monitor.enabled</code>	This system property is required on both engine and SIP data tier servers to enable the heartbeat mechanism.
<code>-Dwlss.ha.heartbeat.interval</code>	Specifies the number of milliseconds between heartbeat messages. By default heartbeats are sent every 1,000 milliseconds.
<code>-Dwlss.ha.heartbeat.count</code>	Specifies the number of consecutive, missed heartbeats that are permitted before a replica is determined to be offline. By default, a replica is marked offline if the <code>WlssEchoServer</code> process on the server fails to respond to 3 heartbeat messages.
<code>-Dwlss.ha.heartbeat.SoTimeout</code>	Specifies the UDP socket timeout value.

5.4 Configuring SNMP

Oracle WebLogic Server SIP Container includes a dedicated SNMP MIB to monitor activity on engine tier and SIP data tier server instances. The Oracle WebLogic Server SIP Container MIB is available on both Managed Servers and the Administration Server of a domain. However, Oracle WebLogic Server SIP Container engine and SIP data tier traps are generated only by the Managed Server instances that make up each tier. If your Administration Server is not a target for the `sipserver` custom resource,

it will generate only WebLogic Server SNMP traps (for example, when a server in a cluster fails). Administrators should monitor both WebLogic Server and Oracle WebLogic Server SIP Container traps to evaluate the behavior of the entire domain.

Note: Oracle WebLogic Server SIP Container MIB objects are read-only. You cannot modify a Oracle WebLogic Server SIP Container configuration using SNMP.

5.4.1 Browsing the MIB

The Oracle WebLogic Server SIP Container MIB file is installed in `WLSS_HOME/server/lib/wlss/BEA-WLSS-MIB.asn1`. Use an available SNMP management tool or MIB browser to view the contents of this file. See also [Section 5.5.2, "Trap Descriptions"](#) for a description of common SNMP traps.

5.4.2 Steps for Configuring SNMP

To enable SNMP monitoring for the entire Oracle WebLogic Server SIP Container domain, follow these steps:

1. Login to the Administration Console for the Oracle WebLogic Server SIP Container domain.
2. In the left pane, select the **Diagnostics > SNMP** node.
3. In the Server SNMP Agents table, click the **New** button to create a new agent.

Note: Ensure that you create a new Server SNMP agent, rather than a Domain-Scoped agent.

4. Enter a unique name for the new SNMP agent (for example, "engine1snmp") and click **OK**.
5. Select the newly-created SNMP agent from the Server SNMP Agents table.
6. On the **Configuration > General** tab:
 - a. Select the Enabled check box to enable the agent.
 - b. Enter an unused port number in the SNMP UDP Port field.

Note: If you run multiple Managed Server instances on the same machine, each server instance must use a dedicated SNMP agent with a unique SNMP port number.

- c. Click **Save**.
7. Repeat the above steps to generate a unique SNMP agent for each server in your deployment (SIP data tier server, engine tier server, and Administration Server).

5.5 Understanding and Responding to SNMP Traps

The following sections describe the Oracle WebLogic Server SIP Container SNMP traps in more detail. Recovery procedures for responding to individual traps are also included where applicable.

5.5.1 Files for Troubleshooting

The following Oracle WebLogic Server SIP Container log and configuration files are frequently helpful for troubleshooting problems, and may be required by your technical support contact:

- `DOMAIN_DIR/config/config.xml`
- `DOMAIN_DIR/config/custom/sipserver.xml`
- `DOMAIN_DIR/servername/*.log` (server and message logs)
- `sip.xml` (in the `/WEB-INF` subdirectory of the application)
- `web.xml` (in the `/WEB-INF` subdirectory of the application)

General information that can help the technical support team includes:

- The specific versions of:
 - Oracle WebLogic Server SIP Container
 - Java SDK
 - Operating System
- Thread dumps for hung Oracle WebLogic Server SIP Container processes
- Network analyzer logs

5.5.2 Trap Descriptions

[Table 5–3](#) lists the Oracle WebLogic Server SIP Container SNMP traps and indicates whether the trap is generated by servers in the engine tier or SIP data tier. Each trap is described in the sections that follow.

Table 5–3 Oracle WebLogic Server SIP Container SNMP Traps

Server Node in which Trap is Generated	Trap Name
Engine Tier Servers	Section 5.5.2.1, "connectionLostToPeer"
	Section 5.5.2.2, "connectionReestablishedToPeer"
	Section 5.5.2.4, "overloadControlActivated, overloadControlDeactivated"
	Section 5.5.2.9, "sipAppDeployed"
	Section 5.5.2.10, "sipAppUndeployed"
Engine and SIP Data Tier Servers, if servers are members of a cluster	Section 5.5.2.11, "sipAppFailedToDeploy"
	Section 5.5.2.8, "serverStopped"
SIP Data Tier Servers	Section 5.5.2.3, "dataTierServerStopped"
	Section 5.5.2.5, "replicaAddedToPartition"
	Section 5.5.2.6, "replicaRemovedEnginesRegistration"
	Section 5.5.2.7, "replicaRemovedFromPartition"

5.5.2.1 connectionLostToPeer

This trap is generated by an engine tier server instance when it loses its connection to a replica in the SIP data tier. It may indicate a network connection problem between the

engine and SIP data tiers, or may be generated with additional traps if a SIP data tier server fails.

5.5.2.1.1 Recovery Procedure If this trap occurs in isolation from other traps indicating a server failure, it generally indicates a network failure. Verify or repair the network connection between the affected engine tier server and the SIP data tier server.

If the trap is accompanied by additional traps indicating a SIP data tier server failure (for example, `dataTierServerStopped`), follow the recovery procedures for the associated traps.

5.5.2.2 `connectionReestablishedToPeer`

This trap is generated by an engine tier server instance when it successfully reconnects to a SIP data tier server after a prior failure (after a `connectionLostToPeer` trap was generated). Repeated instances of this trap may indicate an intermittent network failure between the engine and SIP data tiers.

5.5.2.2.1 Recovery Procedure See [Section 5.5.2.1, "connectionLostToPeer"](#).

5.5.2.3 `dataTierServerStopped`

Oracle WebLogic Server SIP Container SIP data tier nodes generate this alarm when an unrecoverable error occurs in a WebLogic Server instance that is part of the SIP data tier. Note that this trap may be generated by the server that is shutting down, by another replica in the same partition, or in some cases by both servers (network outages can sometimes trigger both servers to generate the same trap).

5.5.2.3.1 Recovery Procedure See the Recovery Procedure for [Section 5.5.2.8, "serverStopped"](#).

5.5.2.4 `overloadControlActivated, overloadControlDeactivated`

Oracle WebLogic Server SIP Container engine tier nodes use a configurable throttling mechanism that helps you control the number of new SIP requests that are processed. After a configured overload condition is observed, Oracle WebLogic Server SIP Container destroys new SIP requests by responding with "503 Service Unavailable" to the caller. The servers continues to destroy new requests until the overload condition is resolved according to a configured threshold control value. This alarm is generated when the throttling mechanism is activated. The throttling behavior should eventually return the server to a non-overloaded state, and further action may be unnecessary.

5.5.2.4.1 Recovery Procedure Follow this recovery procedure:

1. Check other servers to see if they are nearly overloaded.
2. Check to see if the load balancer is correctly balancing load across the application servers, or if it is overloading one or more servers. If additional servers are nearly overloaded, Notify Tier 4 support immediately.
3. If the issue is limited to one server, notify Tier 4 support within one hour.

5.5.2.4.2 Additional Overload Information If you set the queue length as an incoming call overload control, you can monitor the length of the queue using the Administration Console. If you specify a session rate control, you cannot monitor the session rate using the Administration Console. (The Administration Console only displays the current number of SIP sessions, not the rate of new sessions generated.)

5.5.2.5 replicaAddedToPartition

Oracle WebLogic Server SIP Container SIP data tier nodes generate this alarm when a server instance is added to a partition in the SIP data tier.

5.5.2.5.1 Recovery Procedure This trap is generated during normal startup procedures when SIP data tier servers are booted.

5.5.2.6 replicaRemovedEnginesRegistration

SIP data tier nodes generate this alarm if an engine server client that was not registered (or was removed from the list of registered engines) attempts to communicate with the SIP data tier. This trap is generally followed by a `serverStopped` trap indicating that the engine tier server was shut down to preserve SIP data tier consistency.

5.5.2.6.1 Recovery Procedure Restart the engine tier server. Repeated occurrences of this trap may indicate a network problem between the engine tier server and one or more replicas.

5.5.2.7 replicaRemovedFromPartition

Oracle WebLogic Server SIP Container SIP data tier nodes generate this alarm when a server is removed from the SIP data tier, either as a result of a normal shutdown operation or because of a failure. There must be at least one replica remaining in a partition to generate this trap; if a partition has only a single replica and that replica fails, the trap cannot be generated. In addition, because engine tier nodes determine when a replica has failed, an engine tier node must be running in order for this trap to be generated.

5.5.2.7.1 Recovery Procedure If this trap is generated as a result of a server instance failure, additional traps will be generated to indicate the exception. See the recovery procedures for traps generated in addition to `replicaRemovedFromPartition`.

5.5.2.8 serverStopped

This trap indicates that the WebLogic Server instance is now down. This trap applies to both engine tier and SIP data tier server instances, but only when the servers are members of a named WebLogic Server cluster. If this trap is received spontaneously and not as a result of a controlled shutdown, follow the steps below.

5.5.2.8.1 Recovery Procedure Follow this recovery procedure:

1. Use the following command to identify the hung process:

```
ps -ef | grep java
```

There should be only one PID for each WebLogic Server instance running on the machine.

2. After identifying the affected PID, use the following command to kill the process:

```
kill -3 [pid]
```

3. This command generates the actual thread dump. If the process is not immediately killed, repeat the command several times, spaced 5-10 seconds apart, to help diagnose potential deadlock problems, until the process is killed.
4. Attempt to restart Oracle WebLogic Server SIP Container immediately.
5. Make a backup copy of all SIP logs on the affected server to aid in troubleshooting. The location of the logs varies based on the server configuration.

6. Copy each log to assist Tier 4 support with troubleshooting the problem.

Note: Oracle WebLogic Server SIP Container logs are truncated according to your system configuration. Make backup logs immediately to avoid losing critical troubleshooting information.

7. Notify Tier 4 support and include the log files with the trouble ticket.
8. Monitor the server closely over next 24 hours. If the source of the problem cannot be identified in the log files, there may be a hardware or network issue that will reappear over time.

5.5.2.8.2 Additional Shutdown Information The Administration Console generates SNMP messages for managed WebLogic Server instances only until the ServerShutDown message is received. Afterwards, no additional messages are generated.

5.5.2.9 sipAppDeployed

Oracle WebLogic Server SIP Container engine tier nodes generate this alarm when a SIP Servlet is deployed to the container.

5.5.2.9.1 Recovery Procedure This trap is generated during normal deployment operations and does not indicate an exception.

5.5.2.10 sipAppUndeployed

Oracle WebLogic Server SIP Container engine tier nodes generate this alarm when a SIP application shuts down, or if a SIP application is undeployed. This generally occurs when Oracle WebLogic Server SIP Container is shutdown while active requests still exist.

5.5.2.10.1 Recovery Procedure During normal shutdown procedures this alarm should be filtered out and should not reach operations. If the alarm occurs during the course of normal operations, it indicates that someone has shutdown the application or server unexpectedly, or there is a problem with the application. Notify Tier 4 support immediately.

5.5.2.11 sipAppFailedToDeploy

Oracle WebLogic Server SIP Container engine tier nodes generate this trap when an application deploys successfully as a Web Application but fails to deploy as a SIP application.

5.5.2.11.1 Recovery Procedure The typical failure is caused by an invalid `sip.xml` configuration file and should occur only during software installation or upgrade procedures. When it occurs, undeploy the application, validate the `sip.xml` file, and retry the deployment.

Note: This alarm should never occur during normal operations. If it does, contact Tier 4 support immediately.

5.6 Using the WebLogic Diagnostics Framework (WLDF)

The WebLogic Diagnostic Framework (WLDF) consists of a number of components that work together to collect, archive, and access diagnostic information about a

WebLogic Server instance and its applications. Oracle WebLogic Server SIP Container version integrates with several components of the WLDF in order to monitor and diagnose the operation of engine and SIP data tier nodes, as well as deployed SIP Servlets:

- **Data Collectors**—Oracle WebLogic Server SIP Container integrates with the Harvester service to collect information from runtime MBeans, and with the Logger service to archive SIP requests and responses.
- **Watches and Notifications**—Administrators can use the Watches and Notifications component to create complex rules, based on Oracle WebLogic Server SIP Container runtime MBean attributes, that trigger automatic notifications using JMS, JMX, SNMP, SMTP, and so forth.
- **Image Capture**—Oracle WebLogic Server SIP Container instances can collect certain diagnostic data and write the data to an image file when requested by an Administrator. This data can then be used to diagnose problems in a running server.
- **Instrumentation**—Oracle WebLogic Server SIP Container instruments the server and application code with monitors to help you configure diagnostic actions that are performed on SIP messages (requests and responses) that match certain criteria.

The sections that follow provide more details about how Oracle WebLogic Server SIP Container integrates with each of the above WLDF components.

5.6.1 Data Collection and Logging

Oracle WebLogic Server SIP Container uses the WLDF Harvester service to collect data from the attributes of these runtime MBeans:

- `ReplicaRuntimeMBean`
- `SipApplicationRuntimeMBean`
- `SipServerRuntimeMBean`

You can add charts and graphs of this data to your own custom views using the WLDF console extension. To do so, first enable the WLDF console extension by copying the JAR file into the `console-ext` subdirectory of your domain directory:

```
cp ~/bea/wlserver_10.3/server/lib/console-ext/diagnostics-console-extension.jar
~/bea/user_projects/domains/mydomain/console-ext
```

When accessing the WLDF console extension, the Oracle WebLogic Server SIP Container runtime MBean attributes are available in the Metrics tab of the extension.

Oracle WebLogic Server SIP Container also uses the WLDF Logger service to archive SIP and Diameter messages to independent, dedicated log files (by default, `domain_home/logs/server_name/sipMessages.log`). You can configure the name and location of the log file, as well as log rotation policies, using the Configuration > Message Debug tab in the SIP Server Administration Console extension. Note that a server restart is necessary in order to initiate independent logging and log rotation.

5.6.2 Watches and Notifications

The data collected from Oracle WebLogic Server SIP Container runtime MBeans can be used to create automated monitors, or "watches," that observe a server's diagnostic state. One or more notifications can then be configured for use by a watch, in order to

generate a message using SMTP, SNMP, JMX, or JMS when your configured watch conditions and rules occur.

To use watches and notifications, you select the Diagnostics > Diagnostic Modules node in the left pane of the Administration Console and create a new module with the watch rules and notifications required for monitoring your servers. The watch rules can use the metrics collected from Oracle WebLogic Server SIP Container runtime MBeans, messages written to the log file, or events generated by the diagnostic framework.

5.6.3 Image Capture

Oracle WebLogic Server SIP Container adds its own image capture information to the diagnostic image generated by the WLDF. You can generate diagnostic images either on demand, or automatically by configuring watch rules.

The information contained in diagnostic images is intended for use by Oracle technical support personnel when troubleshooting a potential server problem and includes:

- SIP data tier partition and replica configuration
- Call state and timer statistics
- Work manager statistics

5.6.4 Instrumentation

The WLDF instrumentation system creates diagnostic monitors and inserts them into Oracle WebLogic Server SIP Container or application code at specific points in the flow of execution. Oracle WebLogic Server SIP Container integrates with the instrumentation service to provide a built-in DyeInjection monitor. When enabled, this monitor injects dye flags into the diagnostic context when certain SIP messages enter or exist the system. Dye flags are injected based on the monitor's configuration properties, and on certain request attributes.

Oracle WebLogic Server SIP Container adds the dye flags described in [Table 5–4](#) below, as well as the WebLogic Server dye flags USER and ADDR. See *Oracle Fusion Middleware Configuring and Using the Diagnostics Framework for Oracle WebLogic Server* for more information.

Table 5–4 Oracle WebLogic Server SIP Container DyeInjection Flags

Dye Flag	Description
PROTOCOL_SIP	Set in the diagnostic context of all SIP protocol messages.
SIP_REQ	Set in the diagnostic context for all SIP requests that match the value of the property SIP_REQ.
SIP_RES	Set if the SIP response matches the value of property SIP_RES.
SIP_REQURI	Set if the SIP request's reqURI matches the value of property SIP_REQURI.
SIP_ANY_HEADER	Set if the SIP request contains a header that matches the value of the property SIP_ANY_HEADER.

Table 5–4 (Cont.) Oracle WebLogic Server SIP Container DyeInjection Flags

Dye Flag	Description
SIP_RES	This flag is set in the diagnostic context for all SIP responses that match the value of the property SIP_RES.
SIP_REQURI	This flag is set if a SIP request's request URI matches the value of property SIP_REQURI.
SIP_ANY_HEADER	This flag is set if a SIP request contains a header matching the value of the property SIP_ANY_HEADER. The value of SIP_ANY_HEADER is specified using the format <i>messageType.headerName=headerValue</i> where <i>headerValue</i> is either a value or regular expression. For example, you can specify the property as SIP_ANY_HEADER=request.Contact=sip:sipp@localhost:5061 or SIP_ANY_HEADER=response.Contact=sip:findme@172.17.30.50:5060.

Dye flags can be applied to both incoming and outbound SIP messages. The flags are useful for dye filtering, and can be used by delegating monitors to trigger further diagnostic actions.

Oracle WebLogic Server SIP Container provides several delegating monitors that can be applied at the application and server scope, and which may examine dye flags set by the `DyeInjection` monitor. The delegating monitors are described in [Table 5–4](#).

Table 5–5 Oracle WebLogic Server SIP Container Diagnostic Monitors

Monitor Name	Monitor Type	Scope	Pointcuts
occas/Sip_Servlet_Before_Service	Before	Application	At entry of <code>SipServlet.do*</code> or <code>SipServlet.service</code> methods of all implementing subclasses.
occas/Sip_Servlet_After_Service	After	Application	At exit of <code>SipServlet.do*</code> or <code>SipServlet.service</code> methods of all implementing subclasses.
occas/Sip_Servlet_Around_Service	Around	Application	At entry and exit of <code>SipServlet.do*</code> or <code>SipServlet.service</code> methods of all implementing subclasses.
occas/Sip_Servlet_Before_Session	Before	Application	At entry of <code>getAttribute</code> , <code>set</code> , <code>remove</code> , and <code>invalidate</code> methods for both <code>SipSession</code> and <code>SipApplicationSession</code> .
occas/Sip_Servlet_After_Session	After	Application	At exit of <code>getAttribute</code> , <code>set</code> , <code>remove</code> , and <code>invalidate</code> methods for both <code>SipSession</code> and <code>SipApplicationSession</code> .
occas/Sip_Servlet_Around_Session	Around	Application	At entry and exit of <code>getAttribute</code> , <code>set</code> , <code>remove</code> , and <code>invalidate</code> methods for both <code>SipSession</code> and <code>SipApplicationSession</code> .

Table 5–5 (Cont.) Oracle WebLogic Server SIP Container Diagnostic Monitors

Monitor Name	Monitor Type	Scope	Pointcuts
occas/SipSessionDebug	Around	Application	<p>This is a built-in, application-scoped monitor having fixed pointcuts and a fixed debug action. Before and after a pointcut, the monitor performs the <code>SipSessionDebug</code> diagnostic action, which calculates the size of the SIP session after serializing the underlying object.</p> <p>The pointcuts for this monitor are as follows:</p> <ol style="list-style-type: none"> 1. Before and after calls to <code>getSession</code> and <code>getApplicationSession</code> of the <code>SipServletMessage</code> class hierarchy. 2. Before and after calls to <code>getAttribute</code>, <code>setAttribute</code>, and <code>removeAttribute</code> methods in the <code>SipSession</code> and <code>SipApplicationSession</code> classes. <p>Note: The <code>occas/SessionDebugAction-Before</code> event is not triggered for the <code>req.getSession()</code> or <code>req.getApplicationSession()</code> joinpoints. Only the <code>occas/SessionDebugAction-After</code> is triggered, because the Session is made available for inspection only after the joinpoints have executed.</p> <p>Note: If you compile your application using Apache Ant, you must enable the <code>debug</code> attribute to embed necessary debug information into the generated class files.</p>
occas/Sip_Servlet_Before_Message_Send_Internal	Before	Server	At entry of Oracle WebLogic Server SIP Container code that writes messages to the wire.
occas/Sip_Servlet_After_Message_Send_Internal	After	Server	At exit of Oracle WebLogic Server SIP Container code that writes messages to the wire.
occas/Sip_Servlet_Around_Message_Send_Internal	Around	Server	At entry and exit of Oracle WebLogic Server SIP Container code that writes messages to the wire.

5.6.4.1 Configuring Server-Scoped Monitors

To use the server-scoped monitors, you must create a new diagnostic module and create and configure one or more monitors in the module. For the built-in `DyeInjection` monitor, you then add monitor properties to define the specific dye flags. For delegating monitors such as `occas/Sip_Servlet_Before_Message_Send_Internal`, you add monitor properties to define diagnostic actions.

Follow these steps to configure the Oracle WebLogic Server SIP Container `DyeInjection` monitor, a delegate monitor, and enable dye filtering:

1. Access the Administration Console for your domain.
2. Select the `Diagnostics > Diagnostic Modules` node in the left pane of the console.
3. Click `New` to create a new Diagnostic Module. Give the module a descriptive name, such as "instrumentationModule," and click `OK`.

4. Select the new "instrumentationModule" from the list of modules in the table.
5. Select the Targets tab.
6. Select a server on which to target the module and click Save.
7. Return to the Diagnostics > Diagnostic Modules node and select instrumentationModule from the list of modules.
8. Select the Configuration > Instrumentation tab.
9. Select Enabled to enable instrumentation at the server level, then click Save.
10. Add the DyeInjection monitor to the module:
 - a. Click Add/Remove.
 - b. Select the name of a monitor from the Available list (for example, DyeInjection), and use the arrows to move it to the Chosen list.
 - c. Click OK.
 - d. Select the newly-created monitor from the list of available monitors.
 - e. Ensure that the monitor is enabled, and edit the Properties field to add any required properties. For the DyeInjection monitor, sample properties include:

```
SIP_RES=180
SIP_REQ=INVITE
SIP_ANY_HEADER=request.Contact=sip:sipp@localhost:5061
```
 - f. Click Save
11. Add one or more delegate monitors to the module:
 - a. Return to the Configuration > Instrumentation tab for the new module.
 - b. Click Add/Remove.
 - c. Select the name of a delegate monitor from the Available list (for example, occas/Sip_Servlet_Before_Message_Send_Internal), and use the arrows to move it to the Chosen list.
 - d. Click OK.
 - e. Select the newly-created monitor from the list of available monitors.
 - f. Ensure that the monitor is enabled, then select one or more Actions from the available list, and use the arrows to move the actions to the Chosen list. For the occas/Sip_Servlet_Before_Message_Send_Internal monitor, sample actions include DisplayArgumentsAction, StackDumpAction, ThreadDumpAction, and TraceAction.
 - g. Select the check box to EnableDyeFiltering.
 - h. Select one or more Dye Masks, such as SIP_REQ, from the Available list and use the arrows to move them to the Chosen list.
 - i. Click Save

Note: You can repeat the above steps to create additional delegate monitors.

5.6.4.2 Configuring Application-Scoped Monitors

You configure application-scoped monitors in an XML configuration file named `weblogic-diagnostics.xml`. You must store the `weblogic-diagnostics.xml` file in the SIP module's or enterprise application's `META-INF` directory.

The XML file enables instrumentation at the application level, defines point cuts, and also defines delegate monitor dye masks and actions. [Example 5–1](#) shows a sample configuration file that uses the `occas/Sip_Servlet_Before_Service` monitor.

Example 5–1 Sample `weblogic-diagnostics.xml` File

```
<wldf-resource xmlns="http://www.bea.com/ns/weblogic/90/diagnostics">
  <instrumentation>
    <enabled>true</enabled>
    <include>demo.ProxyServlet</include>
    <wldf-instrumentation-monitor>
      <name>occas/Sip_Servlet_Before_Service</name>
      <enabled>true</enabled>
      <dye-mask>SIP_ANY_HEADER</dye-mask>
      <dye-filtering-enabled>true</dye-filtering-enabled>
      <action>DisplayArgumentsAction</action>
    </wldf-instrumentation-monitor>
  </instrumentation>
</wldf-resource>
```

In this example, if an incoming request's diagnostic context contains the `SIP_ANY_HEADER` dye flag, then the `occas/Sip_Servlet_Before_Service` monitor is triggered and the `DisplayArgumentsAction` is executed.

5.7 Logging SIP Requests and Responses

Oracle WebLogic Server SIP Container enables you to perform Protocol Data Unit (PDU) logging for the SIP requests and responses it processes. Logged SIP messages are placed either in the domain-wide log file for Oracle WebLogic Server SIP Container, or in the log files for individual Managed Server instances. Because SIP messages share the same log files as Oracle WebLogic Server SIP Container instances, you can use advanced server logging features such as log rotation, domain log filtering, and maximum log size configuration when managing logged SIP messages.

Administrators configure SIP PDU logging by defining one or more SIP Servlets using the `com.bea.wcp.sip.engine.tracing.listener.TraceMessageListenerImpl` class. Logging criteria are then configured either as parameters to the defined servlet, or in separate XML files packaged with the application.

As SIP requests are processed or SIP responses generated, the logging Servlet compares the message with the filtering patterns defined in a standalone XML configuration file or Servlet parameter. SIP requests and responses that match the specified pattern are written to the log file along with the name of the logging servlet, the configured logging level, and other details. To avoid unnecessary pattern matching, the Servlet marks new SIP Sessions when an initial pattern is matched and then logs subsequent requests and responses for that session automatically.

Logging criteria are defined either directly in `sip.xml` as parameters to a logging Servlet, or in external XML configuration files. See [Section 5.7.3, "Specifying the Criteria for Logging Messages"](#).

Note: Engineers can implement PDU logging functionality in their Servlets either by creating a delegate with the `TraceMessageListenerFactory` in the Servlet's `init()` method, or by using the tracing class in deployed Java applications. Using the delegate enables you to perform custom logging or manipulate incoming SIP messages using the default trace message listener implementation. See [Section 5.7.7, "Adding Tracing Functionality to SIP Servlet Code"](#) for an example of using the factory in a Servlet's `init()` method.

5.7.1 Defining Logging Servlets in sip.xml

Logging Servlets for SIP messages are created by defining Servlets having the implementation class `com.bea.wcp.sip.engine.tracing.listener.TraceMessageListenerImpl`. The definition for a sample `msgTraceLogger` is shown in [Example 5-2](#).

Example 5-2 Sample Logging Servlet

```
<servlet>
  <servlet-name>msgTraceLogger</servlet-name>

  <servlet-class>com.bea.wcp.sip.engine.tracing.listener.TraceMessageListenerImpl</servlet-class>
  <init-param>
    <param-name>domain</param-name>
    <param-value>>true</param-value>
  </init-param>
  <init-param>
    <param-name>level</param-name>
    <param-value>full</param-value>
  </init-param>
  <load-on-startup/>
</servlet>
```

5.7.2 Configuring the Logging Level and Destination

Logging attributes such as the level of logging detail and the destination log file for SIP messages are passed as initialization parameters to the logging Servlet. [Table 5-2](#) lists the parameters and parameter values that you can specify as `init-param` entries. [Example 5-2](#) shows the sample `init-param` entries for a Servlet that logs full SIP message information to the domain log file.

5.7.3 Specifying the Criteria for Logging Messages

The criteria for selecting SIP messages to log can be defined either in XML files that are packaged with the logging Servlet's application, or as initialization parameters in the Servlet's `sip.xml` deployment descriptor. The sections that follow describe each method.

5.7.3.1 Using XML Documents to Specify Logging Criteria

If you do not specify logging criteria as an initialization parameter to the logging Servlet, the Servlet looks for logging criteria in a pair of XML descriptor files in the top level of the logging application. These descriptor files, named `request-pattern.xml` and `response-pattern.xml`, define patterns that Oracle

WebLogic Server SIP Container uses for selecting SIP requests and responses to place in the log file.

Note: By default Oracle WebLogic Server SIP Container logs both requests and responses. If you do not want to log responses, you must define a `response-pattern.xml` file with empty matching criteria.

A typical pattern definition defines a condition for matching a particular value in a SIP message header. For example, the sample `response-pattern.xml` used by the `msgTraceLogger` Servlet matches all MESSAGE requests. The contents of this descriptor are shown in

Example 5-3 Sample response-pattern.xml for msgTraceLogger Servlet

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE pattern
  PUBLIC "Registration//Organization//Type Label//Definition Language"
  "trace-pattern.dtd">
<pattern>
  <equal>
    <var>response.method</var>
    <value>MESSAGE</value>
  </equal>
</pattern>
```

Additional operators and conditions for matching SIP messages are described in [Section 5.7.6, "trace-pattern.dtd Reference"](#). Most conditions, such as the `equal` condition shown in [Example 5-3](#), require a variable (`var` element) that identifies the portion of the SIP message to evaluate. [Table 5-2](#) lists some common variables and sample values. For additional variable names and examples, see *Section 16: Mapping Requests to Servlets* in the SIP Servlet API 1.1 specification (<http://jcp.org/en/jsr/detail?id=289>); Oracle WebLogic Server SIP Container enables mapping of both request and response variables to logging Servlets.

Table 5-6 Pattern-matching Variables and Sample Values

Variable	Sample Values
request.method, response.method	MESSAGE, INVITE, ACK, BYE, CANCEL
request.uri.user, response.uri.user	guest, admin, joe
request.to.host, response.to.host	server.mydomain.com

Both `request-pattern.xml` and `response-pattern.xml` use the same Document Type Definition (DTD). See [Section 5.7.6, "trace-pattern.dtd Reference"](#) for more information.

5.7.3.2 Using Servlet Parameters to Specify Logging Criteria

Pattern-matching criteria can also be specified as initialization parameters to the logging Servlet, rather than as separate XML documents. The parameter names used to specify matching criteria are `request-pattern-string` and `response-pattern-string`. They are defined along with the logging level and destination as described in [Section 5.7.2, "Configuring the Logging Level and Destination"](#).

The value of each pattern-matching parameter must consist of a valid XML document that adheres to the DTD for standalone pattern definition documents (see [Section 5.7.3.1, "Using XML Documents to Specify Logging Criteria"](#)). Because the XML documents that define the patterns and values must not be parsed as part of the `sip.xml` descriptor, you must enclose the contents within the CDATA tag. [Example 5-4](#) shows the full `sip.xml` entry for the sample logging Servlet, `invTraceLogger`. The final two `init-param` elements specify that the Servlet log only INVITE request methods and OPTIONS response methods.

Example 5-4 Logging Criteria Specified as `init-param` Elements

```
<servlet>
  <servlet-name>invTraceLogger</servlet-name>

  <servlet-class>com.bea.wcp.sip.engine.tracing.listener.TraceMessageListenerImpl</s
  ervlet-class>
  <init-param>
    <param-name>domain</param-name>
    <param-value>>true</param-value>
  </init-param>
  <init-param>
    <param-name>level</param-name>
    <param-value>full</param-value>
  </init-param>
  <init-param>
    <param-name>request-pattern-string</param-name>
    <param-value>
      <![CDATA[
        <?xml version="1.0" encoding="UTF-8"?>
        <!DOCTYPE pattern
          PUBLIC "Registration//Organization//Type Label//Definition
Language"
          "trace-pattern.dtd">
        <pattern>
          <equal>
            <var>request.method</var>
            <value>INVITE</value>
          </equal>
        </pattern>
      ]]>
    </param-value>
  </init-param>
  <init-param>
    <param-name>response-pattern-string</param-name>
    <param-value>
      <![CDATA[
        <?xml version="1.0" encoding="UTF-8"?>
        <!DOCTYPE pattern
          PUBLIC "Registration//Organization//Type Label//Definition
Language"
          "trace-pattern.dtd">
        <pattern>
          <equal>
            <var>response.method</var>
            <value>OPTIONS</value>
          </equal>
        </pattern>
      ]]>
    </param-value>
  </init-param>
```

```
<load-on-startup/>
</servlet>
```

5.7.4 Specifying Content Types for Unencrypted Logging

By default Oracle WebLogic Server SIP Container 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, Oracle WebLogic Server SIP Container 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, Oracle WebLogic Server SIP Container 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, Oracle WebLogic Server SIP Container 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.

[Example 5-5](#) shows a sample message-debug configuration that logs String content for three additional Content-Type values, in addition to text/* and application/sdp content.

Example 5-5 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>
```

5.7.5 Enabling Log Rotation and Viewing Log Files

The Oracle WebLogic Server SIP Container logging infrastructure enables you to automatically write to a new log file when the existing log file reaches a specified size. You can also view log contents using the Administration Console or configure additional server-level events that are written to the log.

5.7.6 trace-pattern.dtd Reference

trace-pattern.dtd defines the required contents of the request-pattern.xml and response-pattern.xml, documents, as well as the values for the request-pattern-string and response-pattern-string Servlet init-param variables.

Example 5-6 trace-pattern.dtd

```

<!--
The different types of conditions supported.
- >

<!ENTITY % condition "and | or | not |
                    equal | contains | exists | subdomain-of">

<!--
A pattern is a condition: a predicate over the set of SIP requests.
- >

<!ELEMENT pattern (%condition;)>

<!--
An "and" condition is true if and only if all its constituent conditions
are true.
- >

<!ELEMENT and (%condition;)+>

<!--
An "or" condition is true if at least one of its constituent conditions
is true.
- >

<!ELEMENT or (%condition;)+>

<!--
Negates the value of the contained condition.
- >

<!ELEMENT not (%condition;)>

<!--
True if the value of the variable equals the specified literal value.
- >

<!ELEMENT equal (var, value)>

<!--
True if the value of the variable contains the specified literal value.
- >

<!ELEMENT contains (var, value)>

<!--
True if the specified variable exists.
- >

<!ELEMENT exists (var)>

<!--
- >

<!ELEMENT subdomain-of (var, value)>

<!--
Specifies a variable. Example:
  <var>request.uri.user</var>

```

```
- >

<!ELEMENT var (#PCDATA)>

<!--
Specifies a literal string value that is used to specify rules.
- >

<!ELEMENT value (#PCDATA)>

<!--
Specifies whether the "equal" test is case sensitive or not.
- >

<!ATTLIST equal ignore-case (true|false) "false">

<!--
Specifies whether the "contains" test is case sensitive or not.
- >

<!ATTLIST contains ignore-case (true|false) "false">

<!--
The ID mechanism is to allow tools to easily make tool-specific
references to the elements of the deployment descriptor. This allows
tools that produce additional deployment information (i.e information
beyond the standard deployment descriptor information) to store the
non-standard information in a separate file, and easily refer from
these tools-specific files to the information in the standard sip-app
deployment descriptor.
- >

<!ATTLIST pattern id ID #IMPLIED>
<!ATTLIST and id ID #IMPLIED>
<!ATTLIST or id ID #IMPLIED>
<!ATTLIST not id ID #IMPLIED>
<!ATTLIST equal id ID #IMPLIED>
<!ATTLIST contains id ID #IMPLIED>
<!ATTLIST exists id ID #IMPLIED>
<!ATTLIST subdomain-of id ID #IMPLIED>
<!ATTLIST var id ID #IMPLIED>
<!ATTLIST value id ID #IMPLIED>
```

5.7.7 Adding Tracing Functionality to SIP Servlet Code

Tracing functionality can be added to your own Servlets or to Java code by using the `TraceMessageListenerFactory`. `TraceMessageListenerFactory` enables clients to reuse the default trace message listener implementation behaviors by creating an instance and then delegating to it. The factory implementation instance can be found in the servlet context for SIP Servlets by looking up the value of the `TraceMessageListenerFactory.TRACE_MESSAGE_LISTENER_FACTORY` attribute.

Note: Instances created by the factory are not registered with Oracle WebLogic Server SIP Container to receive callbacks upon SIP message arrival and departure.

To implement tracing in a Servlet, you use the factory class to create a delegate in the Servlet's `init()` method as shown in [Example 5-7](#).

Example 5-7 Using the `TraceMessageListenerFactory`

```
public final class TraceMessageListenerImpl extends SipServlet implements
MessageListener {
    private MessageListener delegate;

    public void init() throws ServletException {
        ServletContext sc = (ServletContext) getServletContext();
        TraceMessageListenerFactory factory = (TraceMessageListenerFactory)
sc.getAttribute(TraceMessageListenerFactory.TRACE_MESSAGE_LISTENER_FACTORY);
        delegate = factory.createTraceMessageListener(getServletConfig());
    }
    public final void onRequest(SipServletRequest req, boolean incoming) {
        delegate.onRequest(req, incoming);
    }
    public final void onResponse(SipServletResponse resp, boolean incoming) {
        delegate.onResponse(resp, incoming);
    }
}
```

5.7.8 Order of Startup for Listeners and Logging Servlets

If you deploy both listeners and logging servlets, the listener classes are loaded first, followed by the Servlets. Logging Servlets are deployed in order according to the load order specified in their Web Application deployment descriptor.

5.8 Tuning JVM Garbage Collection for Production Deployments

Production installations of Oracle WebLogic Server SIP Container generally require extremely small response times (under 50 milliseconds) for clients at all times, even under peak server loads. A key factor in maintaining brief response times is the proper selection and tuning of the JVM's Garbage Collection (GC) algorithm for Oracle WebLogic Server SIP Container instances in the engine tier.

Whereas certain tuning strategies are designed to yield the lowest average garbage collection times or to minimize the frequency of full GCs, those strategies can sometimes result in one or more very long periods of garbage collection (often several seconds long) that are offset by shorter GC intervals. With a production SIP Server installation, all long GC intervals must be avoided in order to maintain response time goals.

The sections that follow describe GC tuning strategies for JRockit and Sun's JVM that generally result in best response time performance.

Note: For more information on JRockit, see *Oracle Fusion Middleware Introduction to Oracle WebLogic Server*.

5.8.1 Modifying JVM Parameters in Server Start Scripts

If you use custom startup scripts to start Oracle WebLogic Server SIP Container engines and replicas, simply edit those scripts to include the recommended JVM options described in the sections that follow.

The Configuration Wizard also installs default startup scripts when you configure a new domain. These scripts are installed in the `MIDDLEWARE_HOME/user_projects/domains/domain_name/bin` directory by default, and include:

- `startWebLogic.cmd`, `startWebLogic.sh`—These scripts start the Administration Server for the domain.
- `startManagedWebLogic.cmd`, `startManagedWebLogic.sh`—These scripts start managed engines and replicas in the domain.

If you use the Oracle-installed scripts to start engines and replicas, you can override JVM memory arguments by first setting the `USER_MEM_ARGS` environment variable in your command shell.

Note: Setting the `USER_MEM_ARGS` environment variable overrides all default JVM memory arguments specified in the Oracle-installed scripts. Always set `USER_MEM_ARGS` to the full list of JVM memory arguments you intend to use. For example, when using the Sun JVM, always add `-XX:MaxPermSize=128m` to the `USER_MEM_ARGS` value, even if you only intend to change the default heap space (`-Xms`, `-Xmx`) parameters.

5.8.2 Tuning Garbage Collection with JRockit

JRockit provides several monitoring tools that you can use to analyze the JVM heap at any given moment, including:

- JRockit Runtime Analyzer—provides a view into the runtime behavior of garbage collection and pause times.
- JRockit Stack Dumps—reveals applications' thread activity to help you troubleshoot and/or improve performance.

Use these and other tools in a controlled environment to determine the effects of JVM settings before you use the settings in a production deployment.

The following sections describe suggested starting JVM options for use with the JRockit. If you use JRockit with the deterministic garbage collector (recommended), use the options described in [Section 5.8.3, "Using Oracle JRockit Real Time \(Deterministic Garbage Collection\)"](#).

5.8.3 Using Oracle JRockit Real Time (Deterministic Garbage Collection)

Very short response times are most easily achieved by using JRockit Real Time, which implements a deterministic garbage collector.

Oracle recommends using the following JVM arguments for engine tier servers in replicated cluster configurations:

```
-Xms1024m -Xmx1024m -XgcPrio:deterministic -XpauseTarget=30ms -XXtlasize:min=8k  
-XXnosystemgc
```

Note: The above settings are configured by default in the `$WLSS_HOME/common/bin/wlssCommenv.sh` file when you use the Configuration Wizard to create a new domain with the JRockit JVM.

You may need to increase the `-XpauseTarget` value for allocation-intensive applications. The value can be decreased for smaller applications under light loads.

Adjust the heap size according to the amount of live data used by deployed applications. As a starting point, set the heap size from 2 to 3 times the amount required by your applications. A value closer to 3 times the required amount generally yields the best performance.

For replica servers, increase the available memory:

```
-Xms3072m -Xmx3072m -XgcPrio:deterministic -XpauseTarget=30ms -XXtlasize:min=8k
-XXnosystemgc
```

These settings fix the heap size and enable the dynamic garbage collector with deterministic garbage collection. `-XpauseTarget` sets the maximum pause time and `-XXtlasize=3k` sets the thread-local area size. `-XXnosystemgc` prevents `System.gc()` application calls from forcing garbage collection.

5.8.4 Using Oracle JRockit without Deterministic Garbage Collection

When using Oracle's JRockit JVM without deterministic garbage collection (not recommended for production deployments), the best response time performance is obtained by using the generational concurrent garbage collector.

The full list of example startup options for an engine tier server are:

```
-Xms1024m -Xmx1024m -Xgc:gencon -XXnosystemgc -XXtlasize:min=3k -XXkeeparearatio=0
-Xns:48m
```

Note: Fine tune the heap size according to the amount of live data used by deployed applications.

The full list of example startup options for a replica server are:

```
-Xms3072m -Xmx3072m -Xgc:gencon -XXnosystemgc -XXtlasize:min=3k -XXkeeparearatio=0
-Xns:48m
```

5.8.5 Tuning Garbage Collection with Sun JDK

When using Sun's JDK, the goal in tuning garbage collection performance is to reduce the time required to perform a full garbage collection cycle. You should not attempt to tune the JVM to minimize the frequency of full garbage collections, because this generally results in an eventual forced garbage collection cycle that may take up to several full seconds to complete.

The simplest and most reliable way to achieve short garbage collection times over the lifetime of a production server is to use a fixed heap size with the default collector and the parallel young generation collector, restricting the new generation size to at most one third of the overall heap.

The following example JVM settings are recommended for most engine tier servers:

```
-server -Xmx1024m -XX:MaxPermSize=128m -XX:+UseParNewGC -XX:+UseConcMarkSweepGC
```

```
-XX:+UseTLAB -XX:+CMSIncrementalMode -XX:+CMSIncrementalPacing
-XX:CMSIncrementalDutyCycleMin=0 -XX:CMSIncrementalDutyCycle=10
-XX:MaxTenuringThreshold=0 -XX:SurvivorRatio=256
-XX:CMSInitiatingOccupancyFraction=60 -XX:+DisableExplicitGC
```

For replica servers, use the example settings:

```
-server -Xmx3072m -XX:MaxPermSize=128m -XX:+UseParNewGC -XX:+UseConcMarkSweepGC
-XX:+UseTLAB -XX:+CMSIncrementalMode -XX:+CMSIncrementalPacing
-XX:CMSIncrementalDutyCycleMin=0 -XX:CMSIncrementalDutyCycle=10
-XX:MaxTenuringThreshold=0 -XX:SurvivorRatio=256
-XX:CMSInitiatingOccupancyFraction=60 -XX:+DisableExplicitGC
```

The above options have the following effect:

- `-XX:+UseTLAB`—Uses thread-local object allocation blocks. This improves concurrency by reducing contention on the shared heap lock.
- `-XX:+UseParNewGC`—Uses a parallel version of the young generation copying collector alongside the concurrent mark-and-sweep collector. This minimizes pauses by using all available CPUs in parallel. The collector is compatible with both the default collector and the Concurrent Mark and Sweep (CMS) collector.
- `-Xms`, `-Xmx`—Places boundaries on the heap size to increase the predictability of garbage collection. The heap size is limited in replica servers so that even Full GCs do not trigger SIP retransmissions. `-Xms` sets the starting size to prevent pauses caused by heap expansion.
- `-XX:MaxTenuringThreshold=0`—Makes the full `NewSize` available to every `NewGC` cycle, and reduces the pause time by not evaluating tenured objects. Technically, this setting promotes all live objects to the older generation, rather than copying them.
- `-XX:SurvivorRatio=128`—Specifies a high survivor ratio, which goes along with the zero tenuring threshold to ensure that little space is reserved for absent survivors.

5.9 Avoiding JVM Delays Caused By Random Number Generation

The library used for random number generation in Sun's JVM relies on `/dev/random` by default for UNIX platforms. This can potentially block the Oracle WebLogic Server SIP Container process because on some operating systems `/dev/random` waits for a certain amount of "noise" to be generated on the host machine before returning a result. Although `/dev/random` is more secure, Oracle recommends using `/dev/urandom` if the default JVM configuration delays Oracle WebLogic Server SIP Container startup.

To determine if your operating system exhibits this behavior, try displaying a portion of the file from a shell prompt:

```
head -n 1 /dev/random
```

1. Open the `$JAVA_HOME/jre/lib/security/java.security` file in a text editor.
2. Change the line:

```
securerandom.source=file:/dev/random
```

to read:

```
securerandom.source=file:/dev/urandom
```

3. Save your change and exit the text editor.

Configuring Diameter Client Nodes and Relay Agents

The following sections describe how to configure individual servers to act as Diameter nodes or relays in a Oracle WebLogic Server SIP Container domain:

- [Section 6.1, "Overview of Diameter Protocol Configuration"](#)
- [Section 6.2, "Steps for Configuring Diameter Client Nodes and Relay Agents"](#)
- [Section 6.3, "Installing the Diameter Domain"](#)
- [Section 6.4, "Enabling the Diameter Console Extension"](#)
- [Section 6.5, "Creating TCP, TLS, and SCTP Network Channels for the Diameter Protocol"](#)
- [Section 6.6, "Configuring Diameter Nodes"](#)
- [Section 6.7, "Example Domain Configuration"](#)
- [Section 6.8, "Troubleshooting Diameter Configurations"](#)

6.1 Overview of Diameter Protocol Configuration

A typical Oracle WebLogic Server SIP Container domain includes support for the Diameter base protocol and one or more IMS Diameter interface applications (Sh, Ro, Rf) deployed to engine tier servers that act as Diameter client nodes. SIP Servlets deployed on the engines can use the available Diameter applications to initiate requests for user profile data, accounting, and credit control, or to subscribe to and receive notification of profile data changes.

One or more server instances may be also be configured as Diameter relay agents, which route Diameter messages from the client nodes to a configured Home Subscriber Server (HSS) or other nodes in the network, but do not modify the messages. Oracle recommends configuring one or more servers to act as relay agents in a domain. The relays simplify the configuration of Diameter client nodes, and reduce the number of network connections to the HSS. Using at least two relays ensures that a route can be established to an HSS even if one relay agent fails.

Note: In order to support multiple HSSs, the 3GPP defines the Dh interface to look up the correct HSS. Oracle WebLogic Server SIP Container does not provide a Dh interface application, and can be configured only with a single HSS.

Note that relay agent servers do not function as either engine or SIP data tier instances—they should not host applications, store call state data, maintain SIP timers, or even use SIP protocol network resources (sip or sips network channels).

Oracle WebLogic Server SIP Container also provides simulator applications for the Sh and Ro protocols. You can use the simulator applications for testing while developing Sh and Ro clients. The simulator applications are not intended for deployment to a production system.

6.2 Steps for Configuring Diameter Client Nodes and Relay Agents

To configure Diameter support in a Oracle WebLogic Server SIP Container domain, follow these steps:

1. Install the Oracle WebLogic Server SIP Container Diameter Domain. Install the Diameter domain, which contains a sample configuration and template applications configured for different Diameter node types. You may use the Diameter domain as a template for your own domain, or to better understand how different Diameter node types are configured.
2. Enable the Diameter console extension. If you are working with the sample Diameter domain, the Diameter console extension is already enabled. If you are starting with a basic Oracle WebLogic Server SIP Container domain, edit the `config.xml` file to enable the extension.
3. Create Diameter network channels. Create the network channels necessary to support Diameter over TCP, TLS, or SCTP transports on engine tier servers and relays.
4. Create and configure the Diameter nodes. Configure the Diameter protocol client applications on engine tier servers with the host name, peers, and routes to relay agents or other network elements, such as an HSS. You can also configure Diameter nodes that operate in standalone mode, without a Oracle WebLogic Server SIP Container instance.

The sections that follow describe each step in detail. See also the [Section 6.7, "Example Domain Configuration"](#).

6.3 Installing the Diameter Domain

The Configuration Wizard includes a Diameter domain template that creates a domain having four Oracle WebLogic Server SIP Container instances:

- An Administration Server (AdminServer)
- A Diameter Sh client node (hssclient)
- A Diameter relay node (relay)
- An HSS simulator (hss)

You can use the installed Diameter domain as the basis for creating your own domain. Or, you can use the customized Diameter Web Applications as templates for configuring existing Oracle WebLogic Server SIP Container instances to function as HSS client or relay agent nodes. The configuration instructions in the sections that follow assume that you have access to the Diameter domain configuration. Follow these steps to install the domain:

1. Change to the `WLS_HOME\common\bin` directory, where `WLS_HOME` is the directory in which you installed the Oracle WebLogic Server 10g Release 3 portion

- Oracle WebLogic Server SIP Container (for example, `c:\bea\wlserver_10.3\common\bin`).
2. Execute the `config.cmd` or `config.sh` script to launch the Configuration Wizard.
 3. Select Create a new WebLogic Domain and click **Next**.
 4. Select Base this domain on an existing template, and click **Browse**.
 5. Select the `diameterdomain.jar` template and click **OK**.
 6. Click **Next**.
 7. Enter a username and password for the Administrator of the new domain, and click **Next**.
 8. Select the startup mode and JDKs to use with the new domain, and click **Next**.
 9. Select No to accept the default template options, and click **Next**.
 10. Click Create to create the new domain using the default domain name and domain location (`/user_projects/domains/diameter`).
 11. Click Done.

Table 6–3 describes the server configuration installed with the Diameter domain.

Table 6–1 Key Configuration Elements of the Diameter Domain

Server Name	Network Channel Configuration	Diameter Applications	Notes
AdminServer	n/a	n/a	The Administration Server provides no SIP or Diameter protocol functionality.
hssclient	diameter (TCP over port 3868) sip (UDP/TCP over port 5060)	WlssShApplication	The <code>hssclient</code> engine functions as a Diameter Sh client node. The server contains network channels supporting both SIP and Diameter traffic. The Diameter node configuration deploys <code>WlssShApplication</code> (<code>com.bea.wcp.diameter.sh.WlssShApplication</code>) to provide IMS Sh interface functionality for deployed SIP Servlets.
relay	diameter (TCP over port 3869)	RelayApplication	The <code>relay</code> engine functions as a Diameter Sh relay node. The server contains a network channel to support both Diameter traffic. The server does not contain a channel to support SIP traffic, as a relay performs no SIP message processing. The Diameter node configuration deploys <code>RelayApplication</code> (<code>com.bea.wcp.diameter.relay.RelayApplication</code>) to provide relay services. The node configuration also defines a realm-based route for relaying messages from the <code>hssclient</code> engine.
hss	diameter (TCP over port 3870)	HssSimulator	The <code>hss</code> engine's Diameter node configuration deploys only the <code>HssSimulator</code> application (<code>com.bea.wcp.diameter.sh.HssSimulator</code>). The server is configured with a Diameter network channel.

6.4 Enabling the Diameter Console Extension

Oracle WebLogic Server SIP Container provides a console extension to help you create and configure Diameter nodes. The actual configuration generated by the extension is stored in a `diameter.xml` configuration file, stored in the `config/custom` subdirectory of the domain directory.

The sample Diameter domain already enables the Diameter console extension. If you are working with a domain that does not enable the extension, edit the `config.xml` file for the domain to specify the custom resource for the extension. [Example 6-1](#) highlights the `config.xml` entries necessary to enable the console extension. Use a text editor to add the highlighted lines in the correct location in `config.xml`.

Example 6-1 `config.xml` Entries for Enabling the Diameter Console Extension

```
<custom-resource>
  <name>sipserver</name>
  <target>hssclient</target>
  <descriptor-file-name>custom/sipserver.xml</descriptor-file-name>

  <resource-class>com.bea.wcp.sip.management.descriptor.resource.SipServerResource</
resource-class>

  <descriptor-bean-class>com.bea.wcp.sip.management.descriptor.beans.SipServerBean</
descriptor-bean-class>
</custom-resource>
  <custom-resource>
    <name>diameter</name>
    <target>hssclient,relay,hss</target>
    <deployment-order>200</deployment-order>
    <descriptor-file-name>custom/diameter.xml</descriptor-file-name>
    <resource-class>com.bea.wcp.diameter.DiameterResource</resource-class>

  <descriptor-bean-class>com.bea.wcp.diameter.management.descriptor.beans.DiameterBe
an</descriptor-bean-class>
</custom-resource>
  <custom-resource>
    <name>ProfileService</name>
    <target>hssclient</target>
    <deployment-order>300</deployment-order>
    <descriptor-file-name>custom/profile.xml</descriptor-file-name>

  <resource-class>com.bea.wcp.profile.descriptor.resource.ProfileServiceResource</re
source-class>

  <descriptor-bean-class>com.bea.wcp.profile.descriptor.beans.ProfileServiceBean</de
scriptor-bean-class>
</custom-resource>
  <admin-server-name>AdminServer</admin-server-name>
</domain>
```

6.5 Creating TCP, TLS, and SCTP Network Channels for the Diameter Protocol

The Oracle WebLogic Server SIP Container Diameter implementation supports the Diameter protocol over the TCP, TLS, and SCTP transport protocols. (SCTP transport is provided with certain restrictions as described in [Section 6.5.2, "Configuring and Using SCTP for Diameter Messaging"](#).)

To enable incoming Diameter connections on a server, you must configure a dedicated network channel of the appropriate protocol type:

- "diameter" channels use TCP transport
- "diameters" channels use TCP/TLS transport
- "diameter-sctp" channels use TCP/SCTP transport.

Servers that use a TCP/TLS channel for Diameter (diameters channels) must also enable two-way SSL. Oracle WebLogic Server SIP Container may automatically upgrade Diameter TCP connections to use TLS as described in the Diameter specification (RFC 3558).

To configure a TCP or TCP/TLS channel for use with the Diameter provider, follow these steps:

1. Access the Administration Console for the Oracle WebLogic Server SIP Container domain.
2. Click **Lock & Edit** to obtain a configuration lock.
3. In the left pane, select the name of the server to configure.
4. In the right pane, select **Protocols > Channels** to display the configured channels.
5. Click **New** to configure a new channel.
6. Fill in the fields of the Identity Properties page as follows:
 - **Name:** Enter an administrative name for this channel, such as "Diameter TCP/TLS Channel."
 - **Protocol:** Select "diameter" to support the TCP transport, "diameters" to support both TCP and TLS transports, or "diameter-sctp" to support TCP transport.

Note: If a server configures at least one TLS channel, the server operates in TLS mode and will reject peer connections from nodes that do not support TLS (as indicated in their capabilities exchange).

7. Click **Next** to continue.
8. Fill in the fields of the Network Channel Addressing page as follows:
 - **Listen Address:** Enter the IP address or DNS name for this channel. On a multi-homed machine, enter the exact IP address of the interface you want to configure, or a DNS name that maps to the exact IP address.
 - **Listen Port:** Enter the port number used to communication via this channel. Diameter nodes conventionally use port 3868 for incoming connections.
 - **External Listen Port:** Re-enter the Listen Port value.
9. Click **Next** to continue.
10. Chose attributes in the Network Channel Properties page as follows:
 - **Enabled:** Select this attribute to ensure that the new channel accepts network traffic.
 - **Tunneling Enabled:** Un-check this attribute for Diameter channels.
 - **HTTP Enabled for this Protocol:** Un-check this attribute for Diameter channels.

- **Outbound Enabled:** Select this attribute to ensure that the node can initiate Diameter messages using the channel.
11. Click **Next** to continue.
 12. For "diameters" channels, select the following two attributes:
 - **Two Way SSL Enabled:** Two-way SSL is required for TLS transport.
 - **Client Certificate Enforced:** Select this attribute to honor available client certificates for secure communication.
 13. Click **Finish** to create the new channel.
 14. Select the name of the newly-created channel in the Network Channel table.
 15. Display the advanced configuration items for the newly-created channel by clicking the Advanced link.
 16. Change the **Idle Connection Timeout** value from the default (65 seconds) to a larger value that will ensure the Diameter connection remains consistently available.

Note: If you do not change the default value, the Diameter connection will be dropped and recreated every 65 seconds with idle traffic.

17. Click **Save**.
18. Click **Activate Changes**.

The servers installed with the Diameter domain template include network channel configurations for Diameter over TCP transport. Note that the relays server includes only a diameter channel and *not* a sip or sips channel. Relay agents should not host SIP Servlets or other applications, therefore no SIP transports should be configured on relay server nodes.

6.5.1 Configuring Two-Way SSL for Diameter TLS Channels

Diameter channels that use TLS (diameters channels) require that you enable two-way SSL, which is disabled by default. Select Two Way Enabled SSL in the steps above to enable two-way SSL.

6.5.2 Configuring and Using SCTP for Diameter Messaging

SCTP is a reliable, message-based transport protocol that is designed for use in telephony networks. SCTP provides several benefits over TCP:

- SCTP preserves the internal structure of messages when transmitting data to an endpoint, whereas TCP transmits raw bytes that must be received in order.
- SCTP supports multihoming, where each endpoint may have multiple IP addresses. The SCTP protocol can transparently failover to another IP address should a connection fail.
- SCTP provides multistreaming capabilities, where multiple streams in a connection transmit data independently of one another.

Oracle WebLogic Server SIP Container supports SCTP for Diameter network traffic, with several limitations:

- Only 1 stream per connection is currently supported.

- SCTP can be used only for Diameter network traffic; SIP traffic cannot use a configured SCTP channel.
- TLS is not supported over SCTP.

In addition, Oracle WebLogic Server SIP Container only supports SCTP channels on the following software platforms:

- Red Hat Enterprise Linux 4.0 AS, ES, WS (Kernel 2.6.9, GCC 3.4 or higher) on 32- or 64-bit hardware
- Sun Solaris 10 on SPARC

SCTP channels can operate on either IPv4 or IPv6 networks. [Section 6.5](#) describes how to create a new SCTP channel. To enable multihoming capabilities for an existing SCTP channel, specify the IPv4 address 0 . 0 . 0 . 0 as the listen address for the channel (or use the :: address for IPv6 networks).

6.6 Configuring Diameter Nodes

The Diameter node configuration for Oracle WebLogic Server SIP Container engines is stored in the `diameter.xml` configuration file (`domain_home/config/custom/diameter.xml`). If you want to provide diameter services (client, server, or relay functionality) on an engine tier server, you must create a new node configuration and target the configuration to an existing engine server instance.

Diameter node configurations are divided into several categories:

- General configuration defines the host identity and realm for the node, as well as basic connection information and default routing behavior.
- Application configuration defines the Diameter application(s) that run on the node, as well as any optional configuration parameters passed to those applications.
- Peer configuration defines the other Diameter nodes with which this node operates.
- Routes configuration defines realm-based routes that the node can use when resolving messages.

The sections that follow describe how to configure each aspect of a Diameter node.

6.6.1 Creating a New Node Configuration (General Node Configuration)

Follow these steps to create a new Diameter node configuration and target it to an existing Oracle WebLogic Server SIP Container engine tier instance:

1. Log in to the Administration Console for the Oracle WebLogic Server SIP Container domain you want to configure.
2. Click Lock & Edit to obtain a configuration lock.
3. Select the Diameter node in the left pane of the Console.
4. Click New in the right pane to create a new Diameter configuration.
5. Fill in the fields of the Create a New Configuration page as described in [Table 6–3](#), then click Finish.

Table 6–2 Diameter Node General Configuration Properties

Property Name	Description
Name	Enter the an administrative name for this Diameter node configuration.
Host	<p>Enter the host identity of this Diameter node, or leave the field blank to automatically assign the host name of the target engine tier server as the Diameter node's host identity. Note that the host identity may or may not match the DNS name.</p> <p>When configuring Diameter support for multiple Sh client nodes, it is best to omit the <code>host</code> element from the <code>diameter.xml</code> file. This enables you to deploy the same Diameter Web Application to all servers in the engine tier cluster, and the host name is dynamically obtained for each server instance.</p>
Realm	<p>Enter the realm name for which this node has responsibility, or leave the field blank to use the domain name portion of the target engine tier server's fully-qualified host name (for example, <code>host@oracle.com</code>).</p> <p>You can run multiple Diameter nodes on a single host using different realms and listen port numbers.</p> <p>Note: An HSS, Application Server, and relay agents must all agree on a realm name or names. The realm name for the HSS and Application Server need not match.</p>
Address	<p>Enter the listen address for this Diameter node, using either the DNS name or IP address, or leave the field blank to use the host identity as the listen address.</p> <p>Note: The host identity may or may not match the DNS name of the Diameter node. Oracle recommends configuring the Address property with an explicit DNS name or IP address to avoid configuration errors.</p>
TLS	Select this option if the Diameter node us configured with support for TLS (diameters network channels). This field is used to advertise TLS capabilities when the node is interrogated by another Diameter node.
Debug	Select this option if you want to enable debug message output. Debug messages are disabled by default.
Message Debug	Select this option if you want to enable tracing for Diameter messages processed by this node. Message tracing is disabled by default.
Dynamic Peers Allowed	Select this option to allow dynamic discovery of Diameter peer nodes. Dynamic peer support is disabled by default. Oracle recommends enabling dynamic peers only when using the TLS transport, because no access control mechanism is available to restrict hosts from becoming peers.
Peer Retry Delay	Enter the amount of time, in seconds, this node waits before retrying a request to a Diameter peer. The default value is 30 seconds.
Request Timeout	Enter the amount of time, in milliseconds, this node waits for an answer message before timing out.
Watchdog Timeout	Enter the number of seconds this node uses for the value of the Diameter Tw watchdog timer interval.
Targets	Enter one or more target engine tier server names. The Diameter node configuration only applies to servers listed in this field.

Table 6–2 (Cont.) Diameter Node General Configuration Properties

Property Name	Description
Default Route Action	Specify an action type that describes the role of this Diameter node when using a default route. The value of this element can be one of the following: <ul style="list-style-type: none"> ▪ none ▪ local ▪ relay ▪ proxy ▪ redirect
Default Route Servers	Specifies one or more target servers for the default route. Any server you include in this element must also be defined as a peer to this Diameter node, or dynamic peer support must be enabled.

6. Click **Activate Changes** to apply the configuration to target servers.

After creating a general node configuration, the configuration name appears in the list of Diameter nodes. You can select the node to configure Diameter applications, peers, and routes, as described in the sections that follow.

6.6.2 Configuring Diameter Applications

Each Diameter node can deploy one or more applications. You configure Diameter applications in the Administration Console using the **Configuration > Applications** page for a selected Diameter node. Follow these steps:

1. Log in to the Administration Console for the Oracle WebLogic Server SIP Container domain you want to configure.
2. Click **Lock & Edit** to obtain a configuration lock.
3. Select the Diameter node in the left pane of the Console.
4. Select the name of a Diameter node configuration in the right pane of the Console.
5. Select the **Configuration > Applications** tab.
6. Click **New** to configure a new Diameter application, or select an existing application configuration from the table.
7. Fill in the application properties as follows:
 - **Application Name:** Enter a name for the application configuration.
 - **Class Name:** Enter the classname of the application to deploy on this node.
 - **Parameters:** Enter optional parameters to pass to the application upon startup.
8. Click **Finish** to create the new application configuration.
9. Click **Activate Changes** to apply the configuration to the Diameter node.

Oracle WebLogic Server SIP Container includes several Diameter applications to support clients using the Sh, Rf, and Ro interfaces, Diameter relays, and simulators for the Sh and Ro interfaces. The sections that follow provide more information about configuring these Oracle WebLogic Server SIP Container Diameter applications.

You can also use the base Diameter API included in Oracle WebLogic Server SIP Container to create and deploy your own Diameter applications.

6.6.2.1 Configuring the Sh Client Application

The Sh client application is implemented as a provider to the Profile Service API. The application transparently generates and responds to the Diameter command codes defined in the Sh application specification. The 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.

The Diameter nodes on which you deploy the Sh client application should be configured with:

- The host names of any relay agents configured in the domain, defined as Diameter peer nodes. If no relay agents are used, all engine tier servers must be added to the list of peers, or dynamic peers must be enabled.
- One or more routes to access relay agent nodes (or the HSS) in the domain.

To configure the Sh client application, you specify the `com.bea.wcp.diameter.sh.WlssShApplication` class. `WlssShApplication` accepts the following parameters:

- `destination.host` configures a static route to the specified host. Include a `destination.host` param definition only if servers communicate directly to an HSS (static routing), without using a relay agent. Omit the `destination.host` param completely when routing through relay agents.
- `destination.realm`—configures a static route to the specified realm. Specify the realm name of relay agent servers or the HSS, depending on whether or not the domain uses relay agents.

[Example 6–2](#) shows a sample node configuration for an Sh client node that uses a relay.

Example 6–2 Sample Diameter Node Configuration with Sh Client Application

```
<?xml version='1.0' encoding='utf-8'?>
<diameter xmlns="http://www.bea.com/ns/wlcp/diameter/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">
  <configuration>
    <name>hssclient</name>
    <target>hssclient</target>
    <host>hssclient</host>
    <realm>oracle.com</realm>
    <!-- Omit the host and realm elements to dynamically assign the host name
         and domain name of individual engine tier servers. - >
    <message-debug-enabled>true</message-debug-enabled>
    <application>
      <name>WlssShApplication</name>
      <class-name>com.bea.wcp.diameter.sh.WlssShApplication</class-name>
      <param>
        <!-- Include a destination.host param definition only if servers will
             communicate directly to an HSS (static routing), without using
             a relay agent. Omit the destination.host param completely when
             routing through relay agents. - >
        <!-- Specify the realm name of relay agent servers or the HSS,
             depending on whether or not the domain uses relay agents. - >
        <name>destination.realm</name>
        <value>hss.com</value>
      </param>
    </application>
  </configuration>
</diameter>
```

```

</application>
<peer>
<!-- Include peer entries for each relay agent server used in the domain.
      If no relay agents are used, include a peer entry for the HSS
      itself, as well as for all other Sh client nodes (all other engine
      tier servers in the domain).
      Alternately, use the allow-dynamic-peers functionality in
      combination with TLS transport to allow peers to be recognized
      automatically. - >
<host>relay</host>
<address>localhost</address>
<!-- The address element can specify either a DNS name or IP address,
      whereas the host element must specify a diameter host identity.
      The diameter host identity may or may not match the DNS name. - >
<port>3869</port>
</peer>
<!-- Enter a default route to a selected relay agent. If the domain does
      not use a relay agent, specify a default route to relay messages
      directly to the HSS. - >
<default-route>
  <action>relay</action>
  <server>relay</server>
</default-route>
</configuration>
</diameter>

```

6.6.2.2 Configuring the Rf Client Application

The Oracle WebLogic Server SIP Container Rf client application enables SIP Servlets to issue offline charging messages using the IMS Rf interface. To configure the Rf application, specify the class `com.bea.wcp.diameter.charging.RfApplication`. The Rf application accepts the following parameters:

- `cdf.host` specifies the host name of the Charging Data Function (CDF).
- `cdf.realm` specifies the realm of the CDF.

6.6.2.3 Configuring the Ro Client Application

The Oracle WebLogic Server SIP Container Ro client application enables SIP Servlets to issue online charging messages using the IMS Ro interface. To configure the Rf application, specify the class `com.bea.wcp.diameter.charging.RoApplication`. The Ro application accepts the following parameters:

- `ocs.host` specifies the host identity of the Online Charging Function (OCF). The OCF you specify host must also be configured as the peer for the Diameter node on which the Ro application is deployed.
- `ocs.realm` can be used instead of `ocs.host` for realm-based routing when using more than one OCF host. The corresponding realm definition must also exist in the Diameter node's configuration.

6.6.2.4 Configuring a Diameter Relay Agent

Relay agents are not required in a Diameter configuration, but Oracle recommends using at least two relay agent servers to limit the number of direct connections to the HSS, and to provide multiple routes to the HSS in the event of a failure.

Note: You must ensure that relay servers *do not* also act as Oracle WebLogic Server SIP Container engine tier servers or SIP data tier servers. This means that the servers should not be configured with "sip" or "sips" network channels.

Relay agent nodes route Sh messages between client nodes and the HSS, but they do not modify the messages except as defined in the Diameter Sh specification. Relays always route responses from the HSS back the client node that initiated the message, or the message the response is dropped if that node is unavailable.

To configure a Diameter relay agent, simply configure the node to deploy an application with the class `com.bea.wcp.diameter.relay.RelayApplication`.

The node on which you deploy the relay application should also configure:

- All other nodes as peers to the relay node.
- A default route that specifies the relay action.

[Example 6-3](#) shows the sample `diameter.xml` configuration for a relay agent node.

Example 6-3 Diameter Relay Node Configuration

```
<?xml version='1.0' encoding='utf-8'?>
<diameter xmlns="http://www.bea.com/ns/wlcp/diameter/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">
  <configuration>
<name>relay</name>
    <target>relay</target>
    <host>relay</host>
    <realm>oracle.com</realm>
    <message-debug-enabled>true</message-debug-enabled>
    <application>
      <name>RelayApplication</name>
      <class-name>com.bea.wcp.diameter.relay.RelayApplication</class-name>
    </application>
    <!-- Define peer connection information for each Diameter node, or use
         the allow-dynamic-peers functionality in combination with TLS
         transport to allow peers to be recognized automatically. - >
    <peer>
      <host>hssclient</host>
      <address>localhost</address>
      <port>3868</port>
    </peer>
    <peer>
      <host>hss</host>
      <address>localhost</address>
      <port>3870</port>
    </peer>
    <route>
      <realm>oracle.com</realm>
      <application-id>16777217</application-id>
      <action>relay</action>
      <server>hssclient</server>
    </route>
    <!-- Enter a default route for this agent to relay messages
         to the HSS. - >
    <default-route>
```

```

    <action>relay</action>
    <server>hss</server>
  </default-route>
</configuration>
</diameter>

```

6.6.2.5 Configuring the Sh and Rf Simulator Applications

Oracle WebLogic Server SIP Container contains two simulator applications that you can use in development or testing environments to evaluate Diameter client applications. To configure a simulator application, you simply deploy the corresponding class to a configured Diameter node:

- `com.bea.wcp.diameter.sh.HssSimulator` simulates an HSS in your domain for testing Sh client applications.
- `com.bea.wcp.diameter.rf.RfSimulator` simulates an CDF host for testing Rf client applications

Note: These simulators are provided for testing or development purposes only, and is not meant as a substitute for a production HSS or CDF.

Diameter nodes that deploy simulator applications can be targeted to running engine tier servers, or they may be started as standalone Diameter nodes. When started in standalone mode, simulator applications accept the command-line options described in [Table 6-3](#).

Table 6-3 *Command-Line Options for Simulator Applications*

Option	Description
<code>-r, -realm realm_name</code>	Specifies the realm name of the Diameter node.
<code>-h, -host host_name</code>	Specifies the host identity of the node.
<code>-a, -address address</code>	Specifies the listen address for this node.
<code>-p, -port port_number</code>	Specifies the listen port number for this node.
<code>-d, -debug</code>	Enables debug output.
<code>-m, -mdebug</code>	Enables Diameter message tracing.

6.6.2.6 Enabling Profile Service (using an Sh backend)

As noted earlier, Sh, Ro, and Rf applications can be configured and used separately, but Sh can take advantage of the Profile Service API. To do so:

1. Configure `ShApplication` in `diameter.xml` (see [Example 6-5](#) for more information).
2. Add a `profile.xml` file to `domain_home/config/custom/profile.xml`. You can either install the Diameter domain as a template and modify the file, or you can manually create `profile.xml` as shown in [Example 6-4](#).

Example 6-4 *profile.xml sample*

```

<profile-service xmlns="http://www.bea.com/ns/wlcp/wlss/profile/300">
  <mapping>

```

```

<map-by>prefix</map-by>
<map-by-prefix>
  <provider-prefix-set>
    <name>sh</name>
    <prefix>sh</prefix>
  </provider-prefix-set>
</map-by-prefix>
</mapping>
<provider>
  <name>sh</name>
  <provider-class>com.bea.wcp.profile.ShProviderCached</provider-class>
</provider>
</profile-service>

```

6.6.3 Configuring Peer Nodes

A Diameter node should define peer connection information for each other Diameter node in the realm, or enable dynamic peers in combination with TLS transport to allow peers to be recognized automatically. You configure Diameter peer nodes in the Administration Console using the Configuration > Peers page for a selected Diameter node. Follow these steps:

1. Log in to the Administration Console for the Oracle WebLogic Server SIP Container domain you want to configure.
2. Click **Lock & Edit** to obtain a configuration lock.
3. Select the Diameter node in the left pane of the Console.
4. Select the name of a Diameter node configuration in the right pane of the Console.
5. Select the **Configuration > Peers** tab.
6. Click **New** to define a new peer entry.
7. Fill in the fields of the Create a New Peer page as follows:
 - **Host:** Enter the peer node's host identity.
 - **Address:** Enter the peer node's address (DNS name or IP address).
 - **Port Number:** Enter the listen port number of the peer node.
 - **Protocol:** Select the protocol used to communicate with the peer (TCP or SCTP).

Note: Oracle WebLogic Server SIP Container attempts to connect to the peer using *only* the protocol you specify (TCP or SCTP). The other protocol is not used, even if a connection fails using the selected protocol. TCP is used as by default if you do not specify a protocol.

- **Watchdog:** Indicate whether the peer supports the Diameter Tw watchdog timer interval.
8. Click **Finish** to create the new peer entry.
 9. Click **Activate Changes** to apply the configuration.

6.6.4 Configuring Routes

Certain Diameter nodes, such as relays, should configure realm-based routes for use when resolving Diameter messages. You configure Diameter routes in the

Administration Console using the Configuration > Routes page for a selected Diameter node. Follow these steps:

1. Log in to the Administration Console for the Oracle WebLogic Server SIP Container domain you want to configure.
2. Click **Lock & Edit** to obtain a configuration lock.
3. Select the Diameter node in the left pane of the Console.
4. Select the name of a Diameter node configuration in the right pane of the Console.
5. Select the **Configuration > Routes** tab.
6. Click **New** to configure a new Route.
7. Fill in the fields of the Create a New Route page as follows:
 - **Name:** Enter an administrative name for the route.
 - **Realm:** Enter the target realm for this route.
 - **Application ID:** Enter the target Diameter application ID for this route.
 - **Action:** Select an action that this node performs when using the configured route. The action type may be one of: none, local, relay, proxy, or redirect.
 - **Server Names:** Enter the names of target servers that will use the route.
8. Click **Finish** to create the new route entry.
9. Click **Activate Changes** to apply the configuration.

See [Example 6-3](#) for an example `diameter.xml` node configuration containing a route entry.

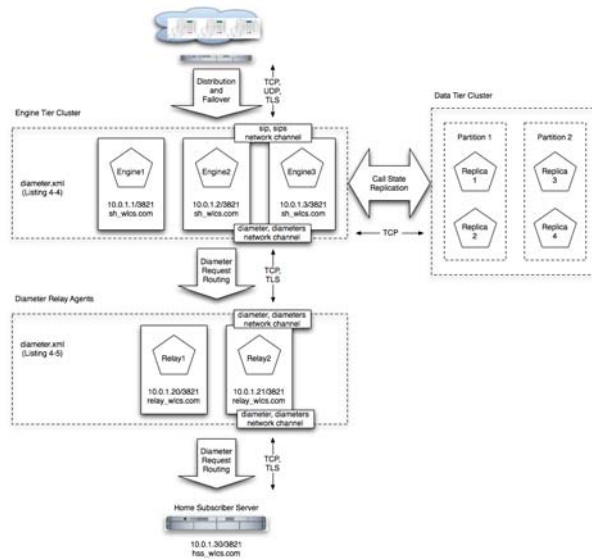
6.7 Example Domain Configuration

This section describes a sample Oracle WebLogic Server SIP Container configuration that provides basic Diameter Sh protocol capabilities. The layout of the sample domain includes the following:

- Three engine tier servers which host SIP applications and also deploy the Diameter Sh application for accessing user profiles.
- Four SIP data tier servers arranged into two partitions with two replicas each.
- Two servers that act as Diameter relay agents and forward diameter requests to an HSS.

[Figure 6-1](#) shows the individual servers in the sample configuration.

Figure 6–1 Sample Diameter Domain



Example 6–5 shows the contents of the `diameter.xml` file used to configure engine tier servers (Sh Clients) in the sample domain. Example 6–6 shows the `diameter.xml` file used to configure the relay agents.

Example 6–5 diameter.xml Configuration for Sample Engine Tier Cluster (Sh Clients)

```
<?xml version='1.0' encoding='utf-8'?>
<diameter xmlns="http://www.bea.com/ns/wlcp/diameter/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">
<configuration>
  <name>clientnodes</name>
  <target>Engine1</target>
  <target>Engine2</target>
  <target>Engine3</target>
  <realm>sh_wlss.com</realm>
  <application>
    <name>WlssShApplication</name>
    <class-name>com.bea.wcp.diameter.sh.WlssShApplication</class-name>
    <param>
      <name>destination.realm</name>
      <value>relay_wlss.com</value>
    </param>
  </application>
  <peer>
    <host>Relay1</host>
    <address>10.0.1.20</address>
    <port>3821</port>
  </peer>
  <peer>
    <host>Relay2</host>
    <address>10.0.1.21</address>
    <port>3821</port>
  </peer>
  <default-route>
    <action>relay</action>
    <server>Relay1</server>
  </default-route>
</configuration>
</diameter>
```



```

    </default-route>
    <route>
      <action>relay</action>
      <server>Relay2</server>
    </route>
  </configuration>
</diameter>

```

Example 6-6 diameter.xml Configuration for Sample Relay Agents

```

<?xml version='1.0' encoding='utf-8'?>
<diameter xmlns="http://www.bea.com/ns/wlcp/diameter/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">
  <configuration>
    <name>relaynodes</name>
    <target>Relay1</target>
    <target>Relay2</target>
    <realm>relay_wlss.com</realm>
    <application>
      <name>RelayApplication</name>
      <class-name>com.bea.wcp.diameter.relay.RelayApplication</class-name>
    </application>
    <peer>
      <host>Engine1</host>
      <address>10.0.1.1</address>
      <port>3821</port>
    </peer>
    <peer>
      <host>Engine2</host>
      <address>10.0.1.2</address>
      <port>3821</port>
    </peer>
    <peer>
      <host>Engine3</host>
      <address>10.0.1.3</address>
      <port>3821</port>
    </peer>
    <peer>
      <host>Relay1</host>
      <address>10.0.1.20</address>
      <port>3821</port>
    </peer>
    <peer>
      <host>Relay2</host>
      <address>10.0.1.21</address>
      <port>3821</port>
    </peer>
    <peer>
      <host>hss</host>
      <address>hssserver</address>
      <port>3870</port>
    </peer>
    <default-route>
      <action>relay</action>
      <server>hss</server>
    </default-route>
  </configuration>
</diameter>

```

6.8 Troubleshooting Diameter Configurations

SIP Servlets deployed on Oracle WebLogic Server SIP Container use the available Diameter applications to initiate requests for user profile data, accounting, and credit control, or to subscribe to and receive notification of profile data changes. If a SIP Servlet performing these requests generates an error similar to:

```
Failed to dispatch Sip message to servlet ServletName  
java.lang.IllegalArgumentException: No registered provider for protocol: Protocol  
The message may indicate that you have not properly configured the associated  
Diameter application for the protocol. See Section 6.6.2, "Configuring Diameter  
Applications" for more information.
```

If you experience problems connecting to a Diameter peer node, verify that you have configured the correct protocol for communicating with the peer in [Section 6.6.3, "Configuring Peer Nodes"](#). Keep in mind that Oracle WebLogic Server SIP Container tries only the protocol you specify for the peer configuration (or TCP if you do not specify a protocol).

Oracle WebLogic Server SIP Container Base Platform Topologies

The following sections provide an overview of the Oracle WebLogic Server SIP Container architecture and deployment topologies:

- [Section 7.1, "Goals of the Oracle WebLogic Server SIP Container Base Platform"](#)
- [Section 7.2, "Load Balancer"](#)
- [Section 7.3, "Engine Tier"](#)
- [Section 7.4, "SIP Data tier"](#)
- [Section 7.5, "Geographically-Redundant Installations"](#)
- [Section 7.6, "Example Hardware Configurations"](#)
- [Section 7.7, "Alternate Configurations"](#)

7.1 Goals of the Oracle WebLogic Server SIP Container Base Platform


Oracle WebLogic Server SIP Container is designed to provide a highly scalable, highly available, performant server for deploying SIP applications. OWLSC is used for deploying any JSR 289 compliant SIP applications.

The Oracle WebLogic Server SIP Container architecture is simple to manage and easily adaptable to make use of available hardware. The basic architecture consists of these components:

- [Section 7.2, "Load Balancer"](#)
- [Section 7.3, "Engine Tier"](#)
- [Section 7.4, "SIP Data tier"](#)

[Figure 7-1](#) shows the components of a basic Oracle WebLogic Server SIP Container installation. The sections that follow describe each component of the architecture in more detail.

Figure 7–1 Oracle WebLogic Server SIP Container Architecture



7.2 Load Balancer

Although it is not provided as part of the Oracle WebLogic Server SIP Container product, a load balancer (or multiple load balancers) is an essential component of any production Oracle WebLogic Server SIP Container installation. The primary goal of a load balancer is to provide a single public address that distributes incoming SIP requests to available servers in the Oracle WebLogic Server SIP Container engine tier. Distribution of requests ensures that Oracle WebLogic Server SIP Container engines are fully utilized.

Most load balancers have configurable policies to ensure that client requests are distributed according to the capacity and availability of individual machines, or according to any other load policies required by your installation. Some load balancers provide additional features for managing SIP network traffic, such as support for routing policies based on source IP address, port number, or other fields available in SIP message headers. Many load balancer products also provide additional fault tolerance features for telephony networks, and can be configured to consistently route SIP requests for a given call to the same engine server on which the call was initiated.

In a Oracle WebLogic Server SIP Container installation, the load balancer is also essential for performing maintenance activities such as upgrading individual servers (Oracle WebLogic Server SIP Container software or hardware) or upgrading applications without disrupting existing SIP clients. The Administrator modifies load balancer policies to move client traffic off of one or more servers, and then performs the required upgrades on the unused server instances. Afterwards, the Administrator modifies the load balancer policies to allow client traffic to resume on the upgraded servers.

Oracle provides detailed information for setting up load balancers with the Oracle WebLogic Server SIP Container engine tier for basic load distribution. See [Section 3.2](#),

["Configuring Load Balancer Addresses"](#) to configure a load balancer used with Oracle WebLogic Server SIP Container.

7.3 Engine Tier

The engine tier is a cluster of Oracle WebLogic Server SIP Container instances that hosts the SIP Servlets and other applications that provide features to SIP clients. The engine tier is a stateless cluster of servers, and it stores no permanent or transient information about the state of SIP dialogs. Instead, all stateful information about SIP dialogs is stored and retrieved from SIP Data Tier ([Section 7.4](#)), which also provides replication and failover services for SIP session data.

Engine tier servers can optionally cache a portion of the session data managed by the SIP data tier. Caching is most useful in configurations that use a SIP-aware load balancer. See [Section 4.7, "Caching SIP Data in the Engine Tier"](#).

The primary goal of the engine tier is to provide maximum throughput and low response time to SIP clients. As the number of calls, or the average duration of calls to your system increases, you can easily add additional server instances to the engine tier to manage the additional load.

Note that although the engine tier consists of multiple Oracle WebLogic Server SIP Container instances, you manage the engine tier as a single, logical entity; SIP Servlets are deployed uniformly to all server instances (by targeting the cluster itself) and the load balancer need not maintain an affinity between SIP clients and servers in the engine tier.

Note: Oracle WebLogic Server SIP Container start scripts use default values for many JVM parameters that affect performance. For example, JVM garbage collection and heap size parameters may be omitted, or may use values that are acceptable only for evaluation or development purposes. In a production system, you must rigorously profile your applications with different heap size and garbage collection settings in order to realize adequate performance. See [Section 5.8, "Tuning JVM Garbage Collection for Production Deployments"](#) for suggestions about maximizing JVM performance in a production domain.

Because the engine tier relies on SIP data tier servers in order to retrieve call state data, Oracle recommends using dual, Gigabit Ethernet Network Interface Cards (NICs) on engine and SIP data tier machines to provide redundant network connections.

7.4 SIP Data tier

The SIP data tier is a cluster of Oracle WebLogic Server SIP Container instances that provides a high-performance, highly-available, in-memory database for storing and retrieving the session state data for SIP Servlets. The goals of the SIP data tier are as follows:

- To provide reliable, performant storage for session data required by SIP applications in the Oracle WebLogic Server SIP Container engine tier.
- To enable administrators to easily scale hardware and software resources as necessary to accommodate the session state for all concurrent calls.

Within the SIP data tier, session data is managed in one or more "partitions" where each partition manages a fixed portion of the concurrent call state. For example, in a system that uses two partitions, the first partition manages one half of the concurrent call state (sessions A through M) while the second partition manages another half of the concurrent call states (sessions N through Z). With three partitions, each partition manages a third of the call state, and so on. Additional partitions can be added as necessary to manage a large number of concurrent calls. A simple hashing algorithm is used to ensure that each call state is uniquely assigned to only one SIP data tier partition.

Within each partition, multiple servers can be added to provide redundancy and failover should other servers in the partition fail. When multiple servers participate in the same partition, the servers are referred to as "replicas" because each server maintains a duplicate copy of the partition's call state. For example, if a two-partition system has two servers in the first partition, each server manages a replica of call states A through M. If one or more servers in a partition fails or is disconnected from the network, any available replica can automatically provide call state data to the engine tier. The SIP data tier can have a maximum of three replicas, providing two levels of redundancy.

See [Chapter 4, "Configuring SIP Data Tier Partitions and Replicas"](#) for more information about configuring the SIP data tier for high availability.

Note: Because the engine tier relies on SIP data tier servers in order to retrieve call state data, Oracle recommends using dual Network Interface Cards (NICs) on engine and SIP data tier machines to provide redundant network connections.

7.4.1 Example of Writing and Retrieving Call State Data

When an initial SIP message is received, Oracle WebLogic Server SIP Container uses Servlet mapping rules to direct the message to the appropriate SIP Servlet deployed in the engine tier. The engine tier maintains no stateful information about SIP dialogs, but instead persists the call state to the engine tier at SIP transaction boundaries. A hashing algorithm is applied to the call state to select a single SIP data tier partition in which to store the call state data. The engine tier server then "writes" the call state to each replica within that partition and locks the call state. For example, if the SIP data tier is configured to use two servers within each partition, the engine tier opens a connection to both replicas in the partition, and writes and locks the call state on each replica.

In a default configuration, the replicas maintain the call state information only in memory (available RAM). Call state data can also be configured for longer-term storage in an RDBMS, and it may also be persisted to an off-site Oracle WebLogic Server SIP Container installation for geographic redundancy.

When subsequent SIP messages are generated for the SIP dialog, the engine tier must first retrieve the call state data from the SIP data tier. The hashing algorithm is again applied to determine the partition that stores the call state data. The engine tier then asks each replica in the partition to unlock and retrieve the call state data, after which a Servlet on the engine tier can update the call state data.

7.4.2 RDBMS Storage for Long-Lived Call State Data

Oracle WebLogic Server SIP Container enables you to store long-lived call state data in an Oracle RDBMS in order to conserve RAM. The SIP data tier persists a call state's

data to the RDBMS after the call dialog has been established, and retrieves or deletes the persisted call state data as necessary to modify or remove the call state. See [Section 4.4, "Storing Long-Lived Call State Data In A RDBMS"](#).

7.5 Geographically-Redundant Installations

Oracle WebLogic Server SIP Container can be installed in a geographically-redundant configuration for customers who have multiple, regional data centers, and require continuing operation even after a catastrophic site failure. The geographically-redundant configuration enables multiple Oracle WebLogic Server SIP Container installations (complete with engine and SIP data tier clusters) to replicate call state transactions between one another. If the a particular site's installation were to suffer a critical failure, the administrator could choose to redirect all network traffic to the secondary, replicated site to minimize lost calls. See [Section 4.6, "Using Geographically-Redundant SIP Data Tiers."](#)


7.6 Example Hardware Configurations

Oracle WebLogic Server SIP Container's flexible architecture enables you to configure engine and SIP data tier servers in a variety of ways to support high throughput and/or provide high availability.

7.7 Alternate Configurations

Not all Oracle WebLogic Server SIP Container requirements require the performance and reliability provided by multiple servers in the engine and SIP data tier. On a development machine, for example, it is generally more convenient to deploy and test applications on a single server, rather than a cluster of servers.

Oracle WebLogic Server SIP Container enables you to combine engine and SIP data tier services on a single server instance when replicating call states is unnecessary. In a combined-tier configuration, the same Oracle WebLogic Server SIP Container instance provides SIP Servlet container functionality and also manages the call state for applications hosted on the server. Although the combined-tier configuration is most commonly used for development and testing purposes, it may also be used in a production environment if replication is not required for call state data. [Figure 7-2](#) shows an example deployment of multiple combined-tier servers in a production environment.

Figure 7-2 Single-Server Configurations with SIP-Aware Load Balancer

Because each server in a combined-tier server deployment manages only the call state for the applications it hosts, the load balancer must be fully "SIP aware." This means that the load balancer actively routes multiple requests for the same call to the same Oracle WebLogic Server SIP Container instance. If requests in the same call are not pinned to the same server, the call state cannot be retrieved. Also keep in mind that if a Oracle WebLogic Server SIP Container instance fails in the configuration shown in [Figure 7-2](#), all calls handled by that server are lost.

Upgrading Deployed SIP Applications

The following sections describe how to upgrade deployed SIP Servlets and converged SIP/HTTP applications to a newer version of the same application without losing active calls:

- [Section 8.1, "Overview of SIP Application Upgrades"](#)
- [Section 8.2, "Requirements and Restrictions for Upgrading Deployed Applications"](#)
- [Section 8.3, "Steps for Upgrading a Deployed SIP Application"](#)
- [Section 8.4, "Assign a Version Identifier"](#)
- [Section 8.5, "Deploy the Updated Application Version"](#)
- [Section 8.6, "Undeploy the Older Application Version"](#)
- [Section 8.7, "Roll Back the Upgrade Process"](#)
- [Section 8.8, "Accessing the Application Name and Version Identifier"](#)
- [Section 8.9, "Using Administration Mode"](#)

8.1 Overview of SIP Application Upgrades

With Oracle WebLogic Server SIP Container, you can upgrade a deployed SIP application to a newer version without losing existing calls being processed by the application. This type of application upgrade is accomplished by deploying the newer application version alongside the older version. Oracle WebLogic Server SIP Container automatically manages the SIP Servlet mapping so that new requests are directed to the new version. Subsequent messages for older, established dialogs are directed to the older application version until the calls complete. After all of the older dialogs have completed and the earlier version of the application is no longer processing calls, you can safely undeploy it.

Oracle WebLogic Server SIP Container's upgrade feature ensures that no calls are dropped while during the upgrade of a production application. The upgrade process also enables you to revert or rollback the process of upgrading an application. If, for example, you determine that there is a problem with the newer version of the deployed application, you can undeploy the newer version and activate the older version.

Note: When you undeploy an active version of an application, the previous application version remains in administration mode. You must explicitly activate the older version in order to direct new requests to the application.

You can also use the upgrade functionality with a SIP administration channel to deploy a new application version with restricted access for final testing. After performing final testing using the administration channel, you can open the application to general SIP traffic.

Oracle WebLogic Server SIP Container application upgrades provide the same functionality as Oracle WebLogic Server 10g Release 3 application upgrades, with the following exceptions:

- Oracle WebLogic Server SIP Container does not support "graceful" retirement of old application versions. Instead, only timeout-based undeployment is supported using the `-retiretimeout` option to `weblogic.Deployer`.
- If you want to use administration mode with SIP Servlets or converged applications, you must configure a `sips-admin` channel that uses TLS transport.
- Oracle WebLogic Server SIP Container handles application upgrades differently in replicated and non-replicated environments. In replicated environments, the server behaves as if the `save-sessions-enabled` element was set to "true" in the `weblogic.xml` configuration file. This preserves sessions across a redeployment operation.

For non-replicated environments, sessions are destroyed immediately upon redeployment.

8.2 Requirements and Restrictions for Upgrading Deployed Applications

To use the application upgrade functionality of Oracle WebLogic Server SIP Container:

- You must assign version information to your updated application in order to distinguish it from the older application version. Note that only the newer version of a deployed application requires version information; if the currently-deployed application contains no version designation, Oracle WebLogic Server SIP Container automatically treats this application as the "older" version. See [Section 8.4, "Assign a Version Identifier"](#).
- A maximum of two different versions of the same application can be deployed at one time.
- If your application hard-codes the use of an application name (for example, in composed applications where multiple SIP Servlets process a given call), you must replace the application name with calls to a helper method that obtains the base application name. WebLogic Server provides `ApplicationRuntimeMBean` methods for obtaining the base application name and version identifier, as well as determining whether the current application version is active or retiring. See [Section 8.8, "Accessing the Application Name and Version Identifier"](#).
- When applications take part in a composed application (using application composition techniques), Oracle WebLogic Server SIP Container always uses the latest version of an application when only the base name is supplied.
- If you want to deploy an application in administration mode, you must configure a `sips-admin` channel that uses TLS transport.

8.3 Steps for Upgrading a Deployed SIP Application

Follow these steps to upgrade a deployed SIP application to a newer version:

1. Assign a Version Identifier—Package the updated version of the application with a version identifier.
2. Deploy the Updated Application Version—Deploy the updated version of the application alongside the previous version to initiate the upgrade process.
3. Undeploy the Older Application Version—After the older application has finished processing all SIP messages for its established calls, you can safely undeploy that version. This leaves the newly-deployed application version responsible for processing all current and future calls.

Each procedure is described in the sections that follow. You can also roll back the upgrade process if you discover a problem with the newly-deployed application. Applications that are composed of multiple SIP Servlets may also need to use the `ApplicationRuntimeMBean` for accessing the application name and version identifier.

8.4 Assign a Version Identifier

Oracle WebLogic Server SIP Container uses a version identifier—a string value—appended to the application name to distinguish between multiple versions of a given application. The version string can be a maximum of 215 characters long, and must consist of valid characters as identified in [Table 8-1](#).

Table 8-1 Valid and Invalid Characters

Valid ASCII Characters	Invalid Version Constructs
a-z	..
A-Z	.
0-9	
period ("."), underscore ("_"), or hyphen ("-") in combination with other characters	

For deployable SIP Servlet WAR files, you must define the version identifier in the MANIFEST.MF file of the application or specify it on the command line at deployment time.

8.4.1 Defining the Version in the Manifest

Both WAR and EAR deployments must specify a version identifier in the MANIFEST.MF file. [Example 8-1](#) shows an application with the version identifier "v2":

Example 8-1 Version Identifier in Manifest

```
Manifest-Version: 1.0
Created-By: 1.4.1_05-b01 (Sun Microsystems Inc.)
Weblogic-Application-Version: v2
```

If you deploy an application without a version identifier, and later deploy with a version identifier, Oracle WebLogic Server SIP Container recognizes the deployments as separate versions of the same application.

8.5 Deploy the Updated Application Version

To begin the upgrade process, simply deploy the updated application archive using either the Administration Console or `weblogic.Deployer` utility. Use the `-retiretimeout` option to the `weblogic.Deployer` utility if you want to automatically undeploy the older application version after a fixed amount of time. For example:

```
java weblogic.Deployer -name MyApp -version v2 -deploy -retiretimeout 7
```

Oracle WebLogic Server SIP Container examines the version identifier in the manifest file to determine if another version of the application is currently deployed. If two versions are deployed, the server automatically begins routing new requests to the most recently-deployed application. The server allows the other deployed application to complete in-flight calls, directs no new calls to it. This process is referred to as "retiring" the older application, because eventually the older application version will process no SIP messages.

Note that Oracle WebLogic Server SIP Container does not compare the actual version strings of two deployed applications to determine which is the higher version. New calls are always routed to the most recently-deployed version of an application.

Oracle WebLogic Server SIP Container also distinguishes between a deployment that has no version identifier (no version string in the manifest) and a subsequent version that does specify a version identifier. This enables you to easily upgrade applications that were packaged before you began including version information as described in [Section 8.4, "Assign a Version Identifier"](#).

8.6 Undeploy the Older Application Version

After deploying a new version of an existing application, the original deployment process messages only for in-flight calls (calls that were initiated with the original deployment). After those in-flight calls complete, the original deployment no longer processes any SIP messages. In most production environments, you will want to ensure that the original deployment is no longer processing messages before you undeploy the application.

To determine whether a deployed application is processing messages, you can obtain the active session count from the application's `SipApplicationRuntimeMBean` instance. [Example 8-2](#) shows the sample WLST commands for viewing the active session count for the `findme` sample application on the default single-server domain.

Based on the active session count value, you can undeploy the application safely (without losing any in-flight calls) or abruptly (losing the active session counts displayed at the time of undeployment).

Use either the Administration Console or `weblogic.Deployer` utility to undeploy the correct deployment name.

Example 8-2 Sample WLST Session for Examining Session Count

```
connect()
custom()
cd ('examples:Location=myserver,Name=myserver_myserver_findme_
findme,ServerRuntime=myserver,Type=SipApplicationRuntime')
ls()
-rw- ActiveAppSessionCount 0
-rw- ActiveSipSessionCount 0
-rw- AppSessionCount 0
-rw- CachingDisabled true
```

```

-rw-  MBeanInfo                weblogic.management.tools.In
fo@5ae636
-rw-  Name                    myserver_myserver_findme_fin
dme
-rw-  ObjectName              examples:Location=myserver,N
ame=myserver_myserver_findme_findme,ServerRuntime=myserver,Type=SipApplicationRu
ntime
-rw-  Parent                  examples:Location=myserver,N
ame=myserver,Type=ServerRuntime
-rw-  Registered              false
-rw-  SipSessionCount        0
-rw-  Type                    SipApplicationRuntime

-rwx  preDeregister          void :

```

8.7 Roll Back the Upgrade Process

If you deploy a new version of an application and discover a problem with it, you can roll back the upgrade process by:

1. Undeploying the active version of the application.
2. Activating the older version of the application. For example:

```
java weblogic.Deployer -name MyApp -appversion v1 -start
```

Note: When you undeploy an active version of an application, the previous application version remains in administration mode. You must explicitly activate the older version in order to direct new requests to the application.

Alternately, you can simply use the `-start` option to start the older application version, which causes the older version of the application to process new requests and retire the newer version.

8.8 Accessing the Application Name and Version Identifier

If you intend to use Oracle WebLogic Server SIP Container's production upgrade feature, applications that are composed of multiple SIP Servlets should not hard-code the application name. Instead of hard-coding the application name, your application can dynamically access the deployment name or version identifier by using helper methods in `ApplicationRuntimeMBean`.

8.9 Using Administration Mode

You can optionally use the `-adminmode` option with `weblogic.Deployer` to deploy a new version of an application in administration mode. While in administration mode, SIP traffic is accepted only via a configured network channel named `sips-admin` having the TLS transport. If no `sips-admin` channel is configured, or if a request is received using a different channel, the server rejects the request with a 503 message.

To transition the application from administration mode to a generally-available mode, use the `-start` option with `weblogic.Deployer`.

Note: If using TLS is not feasible with your application, you can alternately change the Servlet role mapping rules to allow only 1 user on the newer version of the application. This enables you to deploy the newer version alongside the older version, while restricting access to the newer version.

SIP Servlet Container Configuration Reference

The following sections provide a complete reference to the engine tier configuration file, `sipserver.xml`:

- [Section A.1, "Overview of sipserver.xml"](#)
- [Section A.2, "Editing sipserver.xml"](#)
- [Section A.3, "XML Schema"](#)
- [Section A.4, "Example sipserver.xml File"](#)
- [Section A.5, "XML Element Description"](#)

A.1 Overview of sipserver.xml

The `sipserver.xml` file is an XML document that configures the SIP container features provided by an Oracle WebLogic Server SIP Container instance in the engine tier of a server installation. `sipserver.xml` is stored in the `DOMAIN_DIR/config/custom` subdirectory where `DOMAIN_DIR` is the root directory of the Oracle WebLogic Server SIP Container domain.

A.2 Editing sipserver.xml

You should never move, modify, or delete the `sipserver.xml` file during normal operations.

Oracle recommends using the Administration Console to modify `sipserver.xml` indirectly, rather than editing the file by hand. Using the Administration Console ensures that the `sipserver.xml` document always contains valid XML. See also [Configuring Container Properties Using WLST \(JMX\)](#) in the *Configuration Guide*.

You may need to manually view or edit `sipserver.xml` to troubleshoot problem configurations, repair corrupted files, or to roll out custom configurations to a large number of machines when installing or upgrading Oracle WebLogic Server SIP Container. When you manually edit `sipserver.xml`, you must reboot Oracle WebLogic Server SIP Container instances to apply your changes.

Caution: Always use the SipServer node in the Administration Console or the WLST utility to make changes to a running Oracle WebLogic Server SIP Container deployment.

A.2.1 Steps for Editing sipserver.xml

If you need to modify `sipserver.xml` on a production system, follow these steps:

1. Use a text editor to open the `DOMAIN_DIR/config/custom/sipserver.xml` file, where `DOMAIN_DIR` is the root directory of the Oracle WebLogic Server SIP Container domain.
2. Modify the `sipserver.xml` file as necessary. See [Section A.3, "XML Schema"](#) for a full description of the XML elements.
3. Save your changes and exit the text editor.
4. Reboot or start servers to have your changes take effect:

Caution: Always use the SipServer node in the Administration Console or the WLST utility to make changes to a running Oracle WebLogic Server SIP Container deployment.

5. Test the updated system to validate the configuration.

A.3 XML Schema

The schema file for `sipserver.xml`, `wcp-sipserver.xsd`, is installed inside the `wlss-descriptor-binding.jar` library, located in the `WLSS_HOME/server/lib/wlss` directory.

A.4 Example sipserver.xml File

The following shows a simple example of a `sipserver.xml` file:

```
<?xml version="1.0" encoding="UTF-8"?>
<sip-server xmlns="http://www.bea.com/ns/wlcp/wlss/300">
  <overload>
    <threshold-policy>queue-length</threshold-policy>
    <threshold-value>200</threshold-value>
    <release-value>150</release-value>
  </overload>
</sip-server>
```

A.5 XML Element Description

The following sections describe each element used in the `sipserver.xml` configuration file. Each section describes an XML element that is contained within the main `sip-server` element shown in [Figure A-1](#).

A.5.1 enable-timer-affinity

The `enable-timer-affinity` element determines the way in which engine tier servers process expired timers. By default (when `enable-timer-affinity` is omitted from `sipserver.xml`, or is set to "false"), an engine tier server that polls the SIP data tier for expired timers processes all available expired timers. When `enable-timer-affinity` is set to "true," engine tier servers polling the SIP data tier process only those expired timers that are associated with call states that the engine last modified (or expired timers for call states that have no owner).

A.5.2 overload

The `overload` element enables you to throttle incoming SIP requests according to a configured overload condition. When an overload condition occurs, Oracle WebLogic Server SIP Container destroys new SIP requests by responding with "503 Service Unavailable" until the configured release value is observed, or until the size of the server's capacity constraints is reduced (see [Section A.5.2.3, "Overload Control Based on Capacity Constraints"](#)).

User-configured overload controls are applied only to initial SIP requests; SIP dialogues that are already active when an overload condition occurs may generate additional SIP requests that are not throttled.

To configure an overload control, you define the three elements described in [Table A-1](#).

Table A-1 *Nested overload Elements*

Element	Description
<code>threshold-policy</code>	<p>A String value that identifies the type of measurement used to monitor overload conditions:</p> <ul style="list-style-type: none"> ▪ <code>session-rate</code> measures the rate at which new SIP requests are generated. Oracle WebLogic Server SIP Container determines the session rate by calculating the number of new SIP application connections that were created in the last 5 seconds of operation. See Section A.5.2.2, "Overload Control Based on Session Generation Rate". ▪ <code>queue-length</code> measures the sum of the sizes of the capacity constraint work manager components that processes SIP requests and SIP timers. See Section A.5.2.3, "Overload Control Based on Capacity Constraints". ▪ Note: Execute queues are deprecated and no longer used in Oracle WebLogic Server SIP Container. Capacity constraints are used in place of execute queues. The policy name "queue-length" was kept for backward compatibility. <p>You must use only one of the above policies to define an overload control. See Section A.5.2.1, "Selecting an Appropriate Overload Policy" for more information.</p>

Table A-1 (Cont.) Nested overload Elements

Element	Description
threshold-value	<p>Specifies the measured value that causes Oracle WebLogic Server SIP Container to recognize an overload condition and <i>start</i> throttling new SIP requests:</p> <ul style="list-style-type: none"> When using the <code>session-rate</code> threshold policy, <code>threshold-value</code> specifies the number of new SIP requests per second that trigger an overload condition. See Section A.5.2.2, "Overload Control Based on Session Generation Rate". When using the <code>queue-length</code> threshold policy, <code>threshold-value</code> specifies the size of the combined number of requests in the SIP transport and SIP timer capacity constraint components that triggers an overload condition. See Section A.5.2.3, "Overload Control Based on Capacity Constraints". After the <code>threshold-value</code> is observed, Oracle WebLogic Server SIP Container recognizes an overload condition for a minimum of 512 milliseconds during which time new SIP requests are throttled. If multiple overloads occur over a short period of time, the minimum overload of 512 ms is dynamically increased to avoid repeated overloads. After the minimum overload recognition period expires, the overload condition is terminated only after the configured <code>release-value</code> is observed.
release-value	<p>Specifies the measured value that causes Oracle WebLogic Server SIP Container to end an overload condition and <i>stop</i> throttling new SIP requests:</p> <ul style="list-style-type: none"> When using the <code>session-rate</code> threshold policy, <code>release-value</code> specifies the number of new SIP requests per second that terminates session throttling. See Section A.5.2.2, "Overload Control Based on Session Generation Rate". When using the <code>queue-length</code> threshold policy, <code>release-value</code> specifies the combined number of requests in the capacity constraints that terminates session throttling. See Section A.5.2.3, "Overload Control Based on Capacity Constraints".

A.5.2.1 Selecting an Appropriate Overload Policy

Oracle WebLogic Server SIP Container provides two different policies for throttling SIP requests:

- The `session-rate` policy throttles sessions when the volume new SIP sessions reaches a configured rate (a specified number of sessions per second).
- The `queue-length` policy throttles requests after the sum of the requests in the `wlss.connect` work manager and `wlss.timer.capacity` capacity constraint components reaches a configured size.

Note that you must select only one of the available overload policies. You cannot use both policies simultaneously.

The `session-rate` policy is generally used when a back-end resource having a known maximum throughput (for example, an RDBMS) is used to set up SIP calls. In this case, the `session-rate` policy enables you to tie the Oracle WebLogic Server SIP Container overload policy to the known throughput capabilities of the back-end resource.

With the `queue-length` policy, Oracle WebLogic Server SIP Container monitors both CPU and I/O bottlenecks to diagnose an overload condition. The `queue-length` policy is generally used with CPU-intensive SIP applications in systems that have no predictable upper bound associated with the call rate.

The following sections describe each policy in detail.

A.5.2.2 Overload Control Based on Session Generation Rate

Oracle WebLogic Server SIP Container calculates the session generation rate (sessions per second) by monitoring the number of application sessions created in the last 5 seconds. When the session generation rate exceeds the rate specified in the `threshold-value` element, Oracle WebLogic Server SIP Container throttles initial SIP requests until the session generation rate becomes smaller than the configured `release-value`.

The following example configures Oracle WebLogic Server SIP Container to begin throttling SIP requests when the new sessions are created at a rate higher than 50 sessions per second. Throttling is discontinued when the session rate drops to 40 sessions per second:

```
<overload>
  <threshold-policy>session-rate</threshold-policy>
  <threshold-value>50</threshold-value>
  <release-value>40</release-value>
</overload>
```

A.5.2.3 Overload Control Based on Capacity Constraints

By default, SIP messages are handled by a work manager named `wlss.connect` and SIP timers are processed by a work manager named `wlss.timer`. Each work manager has an associated capacity constraint component that sets the number of requests allotted for SIP message handling and timer processing. Work managers are configured in the `config.xml` file for your Oracle WebLogic Server SIP Container. You can also allocate additional threads to the server at boot time using the startup option

```
-Dweblogic.threadpool.MinPoolSize=number_of_threads.
```

Oracle WebLogic Server SIP Container performs `queue-length` overload control by monitoring the combined lengths of the configured capacity constraints. When the sum of the requests in the two constraints exceeds the length specified in the `threshold-value` element, Oracle WebLogic Server SIP Container throttles initial SIP requests until the total requests are reduced to the configured `release-value`.

[Example A-1](#) shows a sample `overload` configuration from `sipserver.xml`. Here, Oracle WebLogic Server SIP Container begins throttling SIP requests when the combined size of the constraints exceeds 200 requests. Throttling is discontinued when the combined length returns to 200 or fewer simultaneous requests.

Example A-1 Sample overload Definition

```
<overload>
  <threshold-policy>queue-length</threshold-policy>
  <threshold-value>200</threshold-value>
  <release-value>150</release-value>
</overload>
```

A.5.2.4 Two Levels of Overload Protection

User-configured overload controls (defined in `sipserver.xml`) represent the first level of overload protection provided by Oracle WebLogic Server SIP Container. They mark the onset of an overload condition and initiate simple measures to avoid dropped calls (generating 503 responses for new requests).

If the condition that caused the overload persists or worsens, then the work manager component used to perform work in the SIP Servlet container may itself become overloaded. At this point, the server no longer utilizes threads to generate 503 responses, but instead begins to drop messages. In this way, the configured size of the SIP container's work manager components represent the second and final level of overload protection employed by the server.

Always configure overload controls in `sipserver.xml` conservatively, and resolve the circumstances that caused the overload in a timely fashion.

A.5.3 message-debug

The `message-debug` element is used to enable and configure access logging with log rotation for Oracle WebLogic Server SIP Container. This element should be used only in a development environment, because access logging logs *all* SIP requests and responses.

A.5.4 proxy—Setting Up an Outbound Proxy Server

RFC 3261 defines an outbound proxy as "A proxy that receives requests from a client, even though it may not be the server resolved by the Request-URI. Typically, a UA is manually configured with an outbound proxy, or can learn about one through auto-configuration protocols."

In Oracle WebLogic Server SIP Container an outbound proxy server is specified using the `proxy` element in `sipserver.xml`. The `proxy` element defines one or more proxy server URIs. You can change the behavior of the proxy process by setting a proxy policy with the `proxy-policy` tag. [Example A-1](#) describes the possible values for the `proxy` elements.

The default behavior is as if **proxy** policy is in effect. The **proxy** policy means that the request is sent out to the configured outbound proxy and the route headers in the request preserve any routing decision taken by Oracle WebLogic Server SIP Container. This enables the outbound proxy to send the request over to the intended recipient after it has performed its actions on the request. The **proxy** policy comes into effect only for the initial requests. As for the subsequent request the Route Set takes precedence over any policy in a dialog. (If the outbound proxy wants to be in the Route Set it can turn record routing on).

Also if a proxy application written on Oracle WebLogic Server SIP Container wishes to override the configured behavior of outbound proxy traversal, then it can add a special header with name "X-BEA-Proxy-Policy" and value "domain". This header is stripped from the request while sending, but the effect is to ignore the configured outbound proxy. The X-BEA-Proxy-Policy custom header can be used by applications to override the configured policy on a request-by-request basis. The value of the header can be "domain" or "proxy". Note, however, that if the policy is overridden to "proxy," the configuration must still have the outbound proxy URIs in order to route to the outbound proxy.

Table A-2 Nested proxy Elements

Element	Description
routing-policy	An optional element that configures the behavior of the proxy. Valid values are: <ul style="list-style-type: none"> ▪ domain - Proxies messages using the routing rule defined by RFC 3261, ignoring any outbound proxy that is specified. ▪ proxy - Sends the message to the downstream proxy specified in the default proxy URI. If there are multiple proxy specifications they are tried in the order in which they are specified. However, if the transport tries a UDP proxy, the settings for subsequent proxies are ignored.
uri	The TCP or UDP URI of the proxy server. You must specify at least one URI for a proxy element. Place multiple URIs in multiple uri elements within the proxy element.

[Example A-2](#) shows the default proxy configuration for Oracle WebLogic Server SIP Container domains. The request in this case is created in accordance with the SIP routing rules, and finally the request is sent to the outbound proxy "sipoutbound.oracle.com".

Example A-2 Sample proxy Definition

```
<proxy>
  <routing-policy>proxy</routing-policy>
  <uri>sip:sipoutbound.oracle.com:5060</uri>
  <!-- Other proxy uri tags can be added. - >
</proxy>
```

A.5.5 t1-timeout-interval

This element sets the value of the SIP protocol T1 timer, in milliseconds. Timer T1 also specifies the initial values of Timers A, E, and G, which control the retransmit interval for INVITE requests and responses over UDP.

Timer T1 also affects the values of timers F, H, and J, which control retransmit intervals for INVITE responses and requests; these timers are set to a value of 64*T1 milliseconds. See the *SIP: Session Initiation Protocol* for more information about SIP timers.

If `t1-timeout-interval` is not configured, Oracle WebLogic Server SIP Container uses the SIP protocol default value of 500 milliseconds.

A.5.6 t2-timeout-interval

This element sets the value of the SIP protocol T2 timer, in milliseconds. Timer T2 defines the retransmit interval for INVITE responses and non-INVITE requests. See the *SIP: Session Initiation Protocol* for more information about SIP timers.

If `t2-timeout-interval` is not configured, Oracle WebLogic Server SIP Container uses the SIP protocol default value of 4 seconds.

A.5.7 t4-timeout-interval

This element sets the value of the SIP protocol T4 timer, in milliseconds. Timer T4 specifies the maximum length of time that a message remains in the network. Timer T4 also specifies the initial values of Timers I and K, which control the wait times for retransmitting ACKs and responses over UDP.

If `t4-timeout-interval` is not configured, Oracle WebLogic Server SIP Container uses the SIP protocol default value of 5 seconds.

A.5.8 timer-b-timeout-interval

This element sets the value of the SIP protocol Timer B, in milliseconds. Timer B specifies the length of time a client transaction attempts to retry sending a request.

If `timer-b-timeout-interval` is not configured, the Timer B value is derived from timer T1 (64*T1, or 32000 milliseconds by default).

A.5.9 timer-f-timeout-interval

This element sets the value of the SIP protocol Timer F, in milliseconds. Timer F specifies the timeout interval for retransmitting non-INVITE requests.

If `timer-f-timeout-interval` is not configured, the Timer F value is derived from timer T1 (64*T1, or 32000 milliseconds by default).

A.5.10 max-application-session-lifetime

This element sets the maximum amount of time, in minutes, that a SIP application session can exist before Oracle WebLogic Server SIP Container invalidates the session. `max-application-session-lifetime` acts as an upper bound for any timeout value specified using the `session-timeout` element in a `sip.xml` file, or using the `setExpires` API.

A value of -1 (the default) specifies that there is no upper bound to application-configured timeout values.

A.5.11 enable-local-dispatch

`enable-local-dispatch` is a server optimization that helps avoid unnecessary network traffic when sending and forwarding messages. You enable the optimization by setting this element "true." When `enable-local-dispatch` is enabled, if a server instance needs to send or forward a message and the message destination is the engine tier's cluster address or the local server address, then the message is routed internally to the local server instead of being sent via the network. Using this optimization can dramatically improve performance when chained applications process the same request.

You may want to disable this optimization if you feel that routing internal messages could skew the load on servers in the engine tier, and you prefer to route all requests via a configured load balancer.

By default `enable-local-dispatch` is set to "false."

A.5.12 cluster-loadbalancer-map

The `cluster-loadbalancer-map` element is used only when upgrading Oracle WebLogic Server SIP Container software, or when upgrading a production SIP Servlet to a new version. It is not required or used during normal server operations.

During a software upgrade, multiple engine tier clusters are defined to host the older and newer software versions. A `cluster-loadbalancer-map` defines the virtual IP address (defined on your load balancer) that correspond to an engine tier cluster configured for an upgrade. Oracle WebLogic Server SIP Container uses this mapping to ensure that engine tier requests for timers and call state data are received from the correct "version" of the cluster. If a request comes from an incorrect version of the

software, the `cluster-loadbalancer-map` entries are used to forward the request to the correct cluster.

Each `cluster-loadbalancer-map` entry contains the two elements described in [Table A-1](#).

Table A-3 Nested `cluster-loadbalancer-map` Elements

Element	Description
<code>cluster-name</code>	The configured name of an engine tier cluster.
<code>sip-uri</code>	The internal SIP URI that maps to the engine tier cluster. This corresponds to a virtual IP address that you have configured in your load balancer. The internal URI is used to forward requests to the correct cluster version during an upgrade.

[Example A-3](#) shows a sample `cluster-loadbalancer-map` entry used during an upgrade.

Example A-3 Sample `cluster-loadbalancer-map` Entry

```
<cluster-loadbalancer-map>
  <cluster-name>EngineCluster</cluster-name>
  <sip-uri>sip:172.17.0.1:5060</sip-uri>
</cluster-loadbalancer-map>
<cluster-loadbalancer-map>
  <cluster-name>EngineCluster2</cluster-name>
  <sip-uri>sip:172.17.0.2:5060</sip-uri>
</cluster-loadbalancer-map>
```

See [Upgrading Software in the Operations Guide](#) for more information.

A.5.13 default-behavior

This element defines the default behavior of the Oracle WebLogic Server SIP Container instance if the server cannot match an incoming SIP request to a deployed SIP Servlet (or if the matching application has been invalidated or timed out). Valid values are:

- `proxy`—Act as a proxy server.
- `ua`—Act as a User Agent.

`proxy` is used as the default if you do not specify a value.

When acting as a User Agent (UA), Oracle WebLogic Server SIP Container acts in the following way in response to SIP requests:

- ACK requests are discarded without notice.
- CANCEL or BYE requests receive response code 481 - Transaction does not exist.
- All other requests receive response code 500 - Internal server error.

When acting as a proxy requests are automatically forwarded to an outbound proxy (see [Section A.5.4, "proxy—Setting Up an Outbound Proxy Server"](#)) if one is configured. If no proxy is defined, Oracle WebLogic Server SIP Container proxies to a specified Request URI only if the Request URI does not match the IP and port number of a known local address for a SIP Servlet container, or a load balancer address configured for the server. This ensures that the request does not constantly loop to the same servers. When the Request URI matches a local container address or load balancer address, Oracle WebLogic Server SIP Container instead acts as a UA.

A.5.14 default-servlet-name

This element specifies the name of a default SIP Servlet to call if an incoming initial request cannot be matched to a deployed Servlet (using standard `servlet-mapping` definitions in `sip.xml`). The name specified in the `default-servlet-name` element must match the `servlet-name` value of a deployed SIP Servlet. For example:

```
<default-servlet-name>myServlet</default-servlet-name>
```

If the name defined in `default-servlet-name` does not match a deployed Servlet, or no value is supplied (the default configuration), Oracle WebLogic Server SIP Container registers the name `com.bea.wcp.sip.engine.BlankServlet` as the default Servlet. The `BlankServlet` name is also used if a deployed Servlet registered as the `default-servlet-name` is undeployed from the container.

`BlankServlet`'s behavior is configured with the `default-behavior` element. By default the Servlet proxies all unmatched requests. However, if the `default-behavior` element is set to "ua" mode, `BlankServlet` is responsible for returning 481 responses for CANCEL and BYE requests, and 500/416 responses in all other cases. `BlankServlet` does not respond to ACK, and it always invalidates the application session.

A.5.15 retry-after-value

Specifies the number of seconds used in the `Retry-After` header for 5xx responses. This value can also include a parameter or a reason code, such as "Retry-After: 18000;duration=3600" or "Retry-After: 120 (I'm in a meeting)."

If the this value is not configured, Oracle WebLogic Server SIP Container uses the default value of 180 seconds.

A.5.16 sip-security

Oracle WebLogic Server SIP Container enables you to configure one or more trusted hosts for which authentication is not performed. When Oracle WebLogic Server SIP Container receives a SIP message, it calls `getRemoteAddress()` on the SIP Servlet message. If this address matches an address defined in the server's trusted host list, no further authentication is performed for the message.

The `sip-security` element defines one or more trusted hosts, for which authentication is not performed. The `sip-security` element contains one or more `trusted-authentication-host` or `trusted-charging-host` elements, each of which contains a trusted host definition. A trusted host definition can consist of an IP address (with or without wildcard placeholders) or a DNS name. [Example A-4](#) shows a sample `sip-security` configuration.

Example A-4 Sample Trusted Host Configuration

```
<sip-security>
  <trusted-authentication-host>myhost1.mycompany.com</trusted-authentication-host>
  <trusted-authentication-host>172.*</trusted-authentication-host>
</sip-security>
```

A.5.17 route-header

3GPP TS 24.229 Version 7.0.0 (http://www.3gpp.org/ftp/Specs/archive/24_series/24.229/24229-700.zip) requires that IMS Application Servers generating new requests (for example, as a B2BUA) include the S-CSCF route header.

In Oracle WebLogic Server SIP Container, the S-CSCF route header must be statically defined as the value of the `route-header` element in `sipserver.xml`. For example:

```
<route-header>
  <uri>Route: sip:wlssl.bea.com</uri>
</route-header>
```

A.5.18 engine-call-state-cache-enabled

Oracle WebLogic Server SIP Container provides the option for engine tier servers to cache a portion of the call state data locally, as well as in the SIP data tier, to improve performance with SIP-aware load balancers. When a local cache is used, an engine tier server first checks its local cache for existing call state data. If the cache contains the required data, and the local copy of the data is up-to-date (compared to the SIP data tier copy), the engine locks the call state in the SIP data tier but reads directly from its cache.

By default the engine tier cache is enabled. To disable caching, set `engine-call-state-cache-enabled` to `false`:

```
<engine-call-state-cache-enabled>false</engine-call-state-cache-enabled>
```

A.5.19 server-header

Oracle WebLogic Server SIP Container enables you to control when a Server header is inserted into SIP messages. You can use this functionality to limit or eliminate Server headers to reduce the message size for wireless networks, or to increase security.

By default, Oracle WebLogic Server SIP Container inserts no Server header into SIP messages. Set the `server-header` to one of the following string values to configure this behavior:

- `none` (the default) inserts no Server header.
- `request` inserts the Server header only for SIP requests generated by the server.
- `response` inserts the Server header only for SIP responses generated by the server.
- `all` inserts the Server header for all SIP requests and responses.

For example, the following element configures Oracle WebLogic Server SIP Container to insert a Server header for all generated SIP messages:

```
<server-header>all</server-header>
```

See also [Section A.5.20, "server-header-value"](#).

A.5.20 server-header-value

Oracle WebLogic Server SIP Container enables you to control the text that is inserted into the Server header of generated messages. This provides additional control over the size of SIP messages and also enables you to mask the server entity for security purposes. By default, Oracle WebLogic Server SIP Container does not insert a Server header into generated SIP messages (see [Section A.5.19, "server-header"](#)). If Server header insertion is enabled but no `server-header-value` is specified, Oracle WebLogic Server SIP Container inserts the value "WebLogic SIP Server." To configure the header contents, enter a string value. For example:

```
<server-header-value>MyCompany Application Server</server-header-value>
```

A.5.21 persistence

The `persistence` element defines enables or disables writing call state data to an RDBMS and/or to a remote, geographically-redundant Oracle WebLogic Server SIP Container installation. For sites that utilize geographically-redundant replication features, the `persistence` element also defines the site ID and the URL at which to persist call state data.

The `persistence` element contains the sub-elements described in [Table A-1](#).

Table A-4 Nested persistence Elements

Element	Description
<code>default-handling</code>	<p>Determines whether or not Oracle WebLogic Server SIP Container observes persistence hints for RDBMS persistence and/or geographical-redundancy. This element can have one of the following values:</p> <ul style="list-style-type: none"> ▪ all—Specifies that call state data may be persisted to both an RDBMS store and to a geographically-redundant Oracle WebLogic Server SIP Container installation. This is the default behavior. Note that actual replication to either destination also requires that the available resources (JDBC datasource and remote JMS queue) are available. ▪ db—Specifies that long-lived call state data is replicated to an RDBMS if the required JDBC datasource and schema are available. ▪ geo—Specifies that call state data is persisted to a remote, geographically-redundant site if the configured site URL contains the necessary JMS resources. ▪ none—Specifies that only in-memory replication is performed to other replicas in the SIP data tier cluster. Call state data is not persisted in an RDBMS or to an external site.
<code>geo-site-id</code>	Specifies the site ID of this installation. All installations that participate in geographically-redundant replication require a unique site ID.
<code>geo-remote-t3-url</code>	Specifies the remote Oracle WebLogic Server SIP Container installation to which this site replicates call state data. You can specify a single URL corresponding to the engine tier cluster of the remote installation. You can also specify a comma-separated list of addresses corresponding to each engine tier server. The URLs must specify the t3 protocol.

[Example A-5](#) shows a sample configuration that uses RDBMS storage for long-lived call state as well as geographically-redundant replication. Call states are replicated to two engine tier servers in a remote location.

Example A-5 Sample persistence Configuration

```
<persistence>
  <default-handling>all</default-handling>
  <geo-site-id>1</geo-site-id>

  <geo-remote-t3-url>t3://remoteEngine1:7050,t3://remoteEngine2:7051</geo-remote-t3-
url>
</persistence>
```

A.5.22 use-header-form

This element configures the server-wide, default behavior for using or preserving compact headers in SIP messages. You can set this element to one of the following values:

- `compact`—Oracle WebLogic Server SIP Container uses the compact form for all system-generated headers. However, any headers that are copied from an originating message (rather than generated) use their original form.
- `force compact`—Oracle WebLogic Server SIP Container uses the compact form for all headers, converting long headers in existing messages into compact headers as necessary.
- `long`—Oracle WebLogic Server SIP Container uses the long form for all system-generated headers. However, any headers that are copied from an originating message (rather than generated) use their original form.
- `force long`—Oracle WebLogic Server SIP Container uses the long form for all headers, converting compact headers in existing messages into long headers as necessary.

A.5.23 `enable-dns-srv-lookup`

This element enables or disables Oracle WebLogic Server SIP Container DNS lookup capabilities. If you set the element to "true," then the server can use DNS to:

- Discover a proxy server's transport, IP address, and port number when a request is sent to a SIP URI.
- Resolve an IP address and/or port number during response routing, depending on the contents of the Sent-by field.

For proxy discovery, Oracle WebLogic Server SIP Container uses DNS resolution only once per SIP transaction to determine transport, IP, and port number information. All retransmissions, ACKs, or CANCEL requests are delivered to the same address and port using the same transport. For details about how DNS resolution takes place, see *RFC 3263: Session Initiation Protocol (SIP): Locating SIP Servers* (<http://www.ietf.org/rfc/rfc3263.txt>).

When a proxy needs to send a response message, Oracle WebLogic Server SIP Container uses DNS lookup to determine the IP address and/or port number of the destination, depending on the information provided in the sent-by field and via header.

By default, DNS resolution is not used ("false").

Note: Because DNS resolution is performed within the context of SIP message processing, any DNS performance problems result in increased latency performance. Oracle recommends using a caching DNS server in a production environment to minimize potential performance problems.

A.5.24 `connection-reuse-pool`

Oracle WebLogic Server SIP Container includes a connection pooling mechanism that can be used to minimize communication overhead with a Session Border Control (SBC) function or Serving Call Session Control Function (S-CSCF). You can configure multiple, fixed pools of connections to different addresses.

Oracle WebLogic Server SIP Container opens new connections from the connection pool on demand as the server makes requests to a configured address. The server then multiplexes new SIP requests to the address using the already-opened connections, rather than repeatedly terminating and recreating new connections. Opened

connections are re-used in a round-robin fashion. Opened connections remain open until they are explicitly closed by the remote address.

Note that connection re-use pools are not used for incoming requests from a configured address.

To configure a connection re-use pool, you define the four nested elements described in [Table A-1](#).

Table A-5 Nested connection-reuse-pool Elements

Element	Description
pool-name	A String value that identifies the name of this pool. All configured pool-name elements must be unique to the domain.
destination	Specifies the IP address or host name of the destination SBC or S-CSCF. Oracle WebLogic Server SIP Container opens or re-uses connection in this pool only when making requests to the configured address.
destination-port	Specifies the port number of the destination SBC or S-CSCF.
maximum-connections	Specifies the maximum number of opened connections to maintain in this pool.

[Example A-6](#) shows a sample connection-reuse-pool configuration having two pools.

Example A-6 Sample connection-reuse-pool Configuration

```
<connection-reuse-pool>
  <pool-name>SBCPool</pool-name>
  <destination>MySBC</destination>
  <destination-port>7070</destination-port>
  <maximum-connections>10</maximum-connections>
</connection-reuse-pool>
<connection-reuse-pool>
  <pool-name>SCSFPool</pool-name>
  <destination>192.168.1.6</destination>
  <destination-port>7071</destination-port>
  <maximum-connections>10</maximum-connections>
</connection-reuse-pool>
```

A.5.25 globally-routable-uri

This element enables you to specify a Globally-Routable User Agent URI (GRUU) that Oracle WebLogic Server SIP Container automatically inserts into Contact and Route-Set headers when communicating with network elements. The URI specified in this element should be the GRUU for the entire Oracle WebLogic Server SIP Container cluster. (In a single-server domain, use a GRUU for the server itself.)

Note that User Agents (UAs) deployed on Oracle WebLogic Server SIP Container typically obtain GRUUs via a registration request. In this case, the application code is responsible both for requesting and subsequently handling the GRUU. To request a GRUU the UA would include the "+sip.instance" Contact header field parameter in each Contact for which GRUU is required. Upon receiving a GRUU, the UA would use the GRUU as the URI for the contact header field when generating new requests.

A.5.26 domain-alias-name

This element defines one or more domains for which Oracle WebLogic Server SIP Container is responsible. If a message has a destination domain that matches a domain specified with a `domain-alias-name` element, Oracle WebLogic Server SIP Container processes the message locally, rather than forwarding it.

The `sipserver.xml` configuration file can have multiple `main-alias-name` elements. Each element can specify either:

- an individual, fully-qualified domain name, such as `myserver.mycompany.com`, or
- a domain name starting with an initial wildcard character, such as `*.mycompany.com`, used to represent all matching domains. Note that only a single wildcard character is supported, and it must be used as the first element of the domain name.

Note: You can also identify these domain names using the Domain Aliases field in the Configuration > General tab of the SipServer Administration Console extension.

A.5.27 enable-rport

This element determines whether or not Oracle WebLogic Server SIP Container automatically adds an `rport` parameter to `Via` headers when acting as a UAC. By default, the server does not add the `rport` parameter; set the element to "true" to automatically add `rport` to requests generated by the server.

Note: You can also set this parameter to "true" by selecting the Symmetric Response Routing option on the Configuration > General tab of the SipServer Administration console extension.

The `rport` parameter is used for symmetric response routing as described in RFC 3581 (<http://www.ietf.org/rfc/rfc3581.txt>). When a message is received by an RFC 3581-compliant server, such as Oracle WebLogic Server SIP Container, the server responds using the remote UDP port number from which the message was received, rather than the port number specified in the `Via` header. This behavior is frequently used when servers reside behind gateway devices that perform Network Address Translation (NAT). The NAT devices maintain a binding between the internal and external port numbers, and all communication must be initiated via the gateway port.

Note that Oracle WebLogic Server SIP Container is compliant with RFC 3581, and will honor the `rport` parameter even if you set the `enable-rport` element to "false." The `enable-rport` element only specifies whether the server automatically adds `rport` to the requests it generates when acting as a UAC. To disable `rport` handling completely (disable RFC 3581 support), you must start the server with the command-line option, `-Dwlss.udp.uas.rport=false`.

Note: `rport` support as described in RFC 3581 requires that SIP responses include the source port of the original SIP request. Because source port information is frequently treated as sensitive data, Oracle recommends using the TLS transport.

A.5.28 image-dump-level

This element specifies the level of detail to record in Oracle WebLogic Server SIP Container diagnostic image files. You can set this element to one of two values:

- `basic`—Records all diagnostic data except for call state data.
- `full`—Records all diagnostic data including call state data.

Note: Recording call state data in the image file can be time consuming. By default, image dump files are recorded using the `basic` option.

You can also set this parameter using the Configuration > General tab of the SipServer Administration console extension.

A.5.29 stale-session-handling

Oracle WebLogic Server SIP Container uses encoded URIs to identify the call states and application sessions associated with a message. When an application is undeployed or upgraded to a new version, incoming requests may have encoded URIs that specify "stale" or nonexistent call or session IDs. The `stale-session-handling` element enables you to configure the action that Oracle WebLogic Server SIP Container takes when it encounters stale session data in a request. The following actions are possible:

- `drop`—Drops the message without logging an error. This setting is desirable for systems that frequently upgrade applications using Oracle WebLogic Server SIP Container's in-place upgrade feature. Using the `drop` action ensures that messages intended for older, incompatible versions of a deployed application are dropped.
- `error`—Responds with an error, so that a UAC might correct the problem. This is the default action. Messages having a `To:` tag cause a 481 Call/Transaction Does Not Exist error, while those without the tag cause a 404 Not Found error.
- `continue`—Ignores the stale session data and continues processing the request.

Note: When it encounters stale session data, Oracle WebLogic Server SIP Container applies the action specified by `stale-session-handling` before considering the value of the `default-behavior` element. This means that the `default-behavior` is performed only when you have configured `stale-session-handling` to perform the `continue` action.

A.5.30 enable-contact-provisional-response

By default Oracle WebLogic Server SIP Container does not place a `Contact` header in non-reliable provisional (1xx) responses that have a `To` header. If you deploy applications that expect the `Contact` header to be present in such 1xx responses, set this element to `true`:

```
<enable-contact-provisional-response>true</enable-contact-provisional-response>
```

Note that setting this element to `true` does not affect 100 Trying responses.

A.5.31 app-router

The `app-router` stanza contains several elements that configure SIP Servlet v1.1 application router behavior. See [Section A.5.32, "use-custom-app-router"](#), [Section A.5.33, "app-router-config-data"](#), [Section A.5.34, "custom-app-router-jar-file-name"](#), and [Section A.5.35, "default-application-name"](#).

A.5.32 use-custom-app-router

The `use-custom-app-router` element determines whether Oracle WebLogic Server SIP Container uses the default, built-in Application Router (AR), or a custom AR that you specify with the `custom-app-router-jar-file-name` element. The default value, "false," configures the server to use the default AR.

A.5.33 app-router-config-data

The `app-router-config-data` element defines properties to pass to the default or custom Application Router (AR) in the `init` method. All configuration properties must conform to the Java Properties format, and each individual property must be entered on a separate, single line without line breaks or spaces. DAR properties must conform to the detailed property format described in Appendix C of <http://jcp.org/en/jsr/detail?id=289>. [Example A-7](#) shows an example configuration.

Example A-7 Sample app-router-config-data element

```
<?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:wls="http://www.bea.com/ns/weblogic/90/security/wls"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <app-router>
    <use-custom-app-router>false</use-custom-app-router>

    <app-router-config-data>INVITE: ("OriginatingCallWaiting", "DAR:From", "ORIGINATING",
    "", "NO_ROUTE", "0"), ("CallForwarding", "DAR:To", "TERMINATING", "", "NO_ROUTE", "1")
SUBSCRIBE: ("CallForwarding", "DAR:To", "TERMINATING", "", "NO_
ROUTE", "1")</app-router-config-data>
    <custom-app-router-jar-file-name></custom-app-router-jar-file-name>
    <default-application-name></default-application-name>
  </app-router>
</sip-server>
```

You can optionally specify AR initialization properties when starting the Oracle WebLogic Server SIP Container instance by including the `-Djavax.servlet.sip.ar.dar.configuration` Java option. (To specify a property file, rather than a URI, include the prefix `file:///`) If you specify the Java startup option, the container ignores any configuration properties defined in `app-router-config-data`. You can modify the properties in at any time, but the properties are not passed to the AR until the server is restarted with the `-Djavax.servlet.sip.ar.dar.configuration` option omitted.

A.5.34 custom-app-router-jar-file-name

The `custom-app-router-jar-file-name` element specifies the filename of the custom Application Router (AR), packaged as a JAR file, to use. The custom AR implementation must reside in the `$DOMAIN_HOME/approuter` subdirectory.

A.5.35 default-application-name

The `default-application-name` element specifies the name of a default application that the container should call when the custom Application Router (AR) cannot find an application to process an initial request. If no default application is specified, the container returns a 500 error if the AR cannot select an application.

Note: You must first deploy an application before specifying its name as the value of **Default application name**.

SIP Data Tier Configuration Reference

The following sections provide a complete reference to the SIP data tier configuration file, `datatier.xml`:

- [Section B.1, "Overview of datatier.xml"](#)
- [Section B.2, "Editing datatier.xml"](#)
- [Section B.3, "XML Schema"](#)
- [Section B.4, "Example datatier.xml File"](#)
- [Section B.5, "XML Element Description"](#)

B.1 Overview of datatier.xml

The `datatier.xml` configuration file identifies servers that manage the concurrent call state for SIP applications, and defines how those servers are arranged into SIP data tier *partitions*. A *partition* refers to one or more SIP data tier server instances that manage the same portion of the call state. Multiple servers in the same partition are referred to as *replicas* because they all manage a copy of the same portion of the call state.

`datatier.xml` is stored in the `DOMAIN_DIR/config/custom` subdirectory where `DOMAIN_DIR` is the root directory of the Oracle WebLogic Server SIP Container domain.

B.2 Editing datatier.xml

You can edit `datatier.xml` using either the Administration Console or a text editor. Note that changes to the SIP data tier configuration cannot be applied to servers dynamically; you must restart servers in order to change SIP data tier membership or reconfigure partitions.

B.3 XML Schema

This schema file is bundled within the `wcp-sipserver.xsd` library, installed in the `WLSS_HOME/server/lib/wlss` directory.

B.4 Example datatier.xml File

[Example B-1](#) shows the template `datatier.xml` file created using the Configuration Wizard.

Example B-1 Default `datatier.xml` File

```
<st:data-tier xmlns:st="http://www.bea.com/ns/wlcp/wlss/300">
  <st:partition>
    <st:name>partition-0</st:name>
    <st:server-name>replica1</st:server-name>
    <st:server-name>replica2</st:server-name>
  </st:partition>
</st:data-tier>
```

B.5 XML Element Description

`datatier.xml` contains one or more `partition` elements that define servers' membership in a SIP data tier partition. All SIP data tier clusters must have at least one `partition`. Each partition contains the XML elements described in [Table B-1](#).

Table B-1 Nested partition Elements

Element	Description
<code>name</code>	A String value that identifies the name of the partition. Oracle recommends including the number of the partition (starting at 0) in the text of the name for administrative purposes. For example, "partition-0."
<code>server-name</code>	Specifies the name of a Oracle WebLogic Server SIP Container instance that manages call state in this partition. You can define up two three servers per <code>partition</code> element. Multiple servers in the same partition maintain the same call state data, and are referred to as <i>replicas</i> . Oracle recommends including the number of the server (starting with 0) and the number of the partition in the server name for administrative purposes. For example, "replica-0-0."

Diameter Configuration Reference

The following sections provide a complete reference to the Diameter configuration file, `diameter.xml`:

- [Section C.1, "Overview of diameter.xml"](#)
- [Section C.2, "Graphical Representation"](#)
- [Section C.3, "Editing diameter.xml"](#)
- [Section C.4, "XML Schema"](#)
- [Section C.5, "Example diameter.xml File"](#)
- [Section C.6, "XML Element Description"](#)

C.1 Overview of diameter.xml

The `diameter.xml` file configures attributes of a Diameter node, such as:

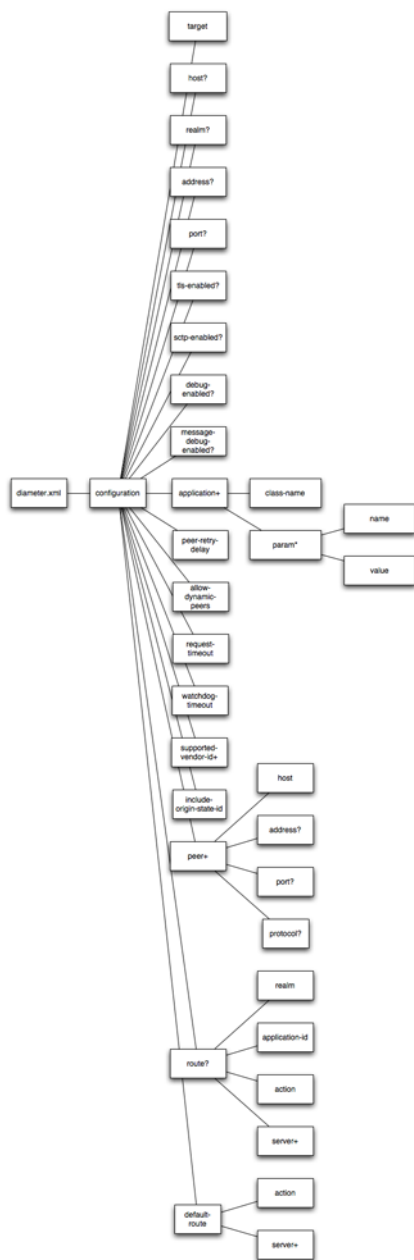
- The host identity of the Diameter node
- The Diameter applications that are deployed on the node
- Connection information for Diameter peer nodes
- Routing information and default routes for handling Diameter messages.

The Diameter protocol implementation reads the configuration file at boot time. `diameter.xml` is stored in the `DOMAIN_DIR/config/custom` subdirectory where `DOMAIN_DIR` is the root directory of the Oracle WebLogic Server SIP Container domain.

C.2 Graphical Representation

[Figure C-1](#) shows the element hierarchy of the `diameter.xml` file.

Figure C-1 Element Hierarchy of diameter.xml



C.3 Editing diameter.xml

You should never move, modify, or delete the `diameter.xml` file during normal operations.

Oracle recommends using the Administration Console to modify `diameter.xml` indirectly, rather than editing the file by hand. Using the Administration Console ensures that the `diameter.xml` document always contains valid XML.

You may need to manually view or edit `diameter.xml` to troubleshoot problem configurations, repair corrupted files, or to roll out custom Diameter node configurations to a large number of machines when installing or upgrading Oracle

WebLogic Server SIP Container. When you manually edit `diameter.xml`, you must reboot Diameter nodes to apply your changes.

Caution: Always use the Diameter node in the Administration Console or the WLST utility, as described in *Configuring Engine Tier Container Properties* in the *Configuration Guide* to make changes to a running Oracle WebLogic Server SIP Container deployment.

C.3.1 Steps for Editing `diameter.xml`

If you need to modify `diameter.xml` on a production system, follow these steps:

1. Use a text editor to open the `DOMAIN_DIR/config/custom/diameter.xml` file, where `DOMAIN_DIR` is the root directory of the Oracle WebLogic Server SIP Container domain.
2. Modify the `diameter.xml` file as necessary. See [Section C.6, "XML Element Description"](#) for a full description of the XML elements.
3. Save your changes and exit the text editor.
4. Reboot or start servers to have your changes take effect:

Caution: Always use the Diameter node in the Administration Console or the WLST utility, as described in *Configuring Engine Tier Container Properties* in the *Configuration Guide*, to make changes to a running Oracle WebLogic Server SIP Container deployment.

5. Test the updated system to validate the configuration.

C.4 XML Schema

The xml schema file (`wcp-diameter.xsd`) is bundled within the `wlssdiameter.jar` library, installed in the `WLSS_HOME/server/lib/wlss` directory.

C.5 Example `diameter.xml` File

See *Configuring Diameter Sh Client Nodes and Relay Agents* in *Configuring Network Resources* for multiple listings of example `diameter.xml` configuration files.

C.6 XML Element Description

The following sections describe each XML element in `diameter.xml`.

C.6.1 configuration

The top level `configuration` element contains the entire diameter node configuration.

C.6.2 target

Specifies one or more target Oracle WebLogic Server SIP Container instances to which the node configuration is applied. The target servers must be defined in the `config.xml` file for your domain.

C.6.3 host

Specifies the host identity for this Diameter node. If no `host` element is specified, the identity is taken from the local server's host name. Note that the host identity may or may not match the DNS name.

Note: When configuring Diameter support for multiple Sh client nodes, it is best to omit the `host` element from the `diameter.xml` file. This enables you to deploy the same Diameter Web Application to all servers in the engine tier cluster, and the host name is dynamically obtained for each server instance.

C.6.4 realm

Specifies the realm name for which this Diameter node has responsibility. You can run multiple Diameter nodes on a single host using different realms and listen port numbers. The HSS, Application Server, and relay agents must all agree on a realm name or names. The realm name for the HSS and Application Server need not match.

If you omit the `realm` element, the realm named is derived using the domain name portion of the host name, if the host name is fully-qualified (for example, `host@oracle.com`).

C.6.5 address

Specifies the listen address for this Diameter node, using either the DNS name or IP address. If you do not specify an address, the node uses the `host` identity as the listen address.

Note: The `host` identity may or may not match the DNS name of the Diameter node. Oracle recommends configuring the `address` element with an explicit DNS name or IP address to avoid configuration errors.

C.6.6 port

Specifies the TCP or TLS listen port for this Diameter node. The default port is 3868.

C.6.7 tls-enabled

This element is used only for standalone node operation to advertise TLS capabilities.

Oracle WebLogic Server SIP Container ignores the `tls-enabled` element for nodes running within a server instance. Instead, TLS transport is reported as enabled if the server instance has configured a Network Channel having TLS support (a `diameters` channel). See *Creating Network Channels for the Diameter Protocol* in *Configuring Network Resources*.

C.6.8 sctp-enabled

This element is used only for standalone node operation to advertise SCTP capabilities.

Oracle WebLogic Server SIP Container ignores the `sctp-enabled` element for nodes running within a server instance. Instead, SCTP transport is reported as enabled if the server instance has configured a Network Channel having SCTP support (a `diameter-sctp` channel). See *Creating Network Channels for the Diameter Protocol in Configuring Network Resources*.

C.6.9 debug-enabled

Specifies a boolean value to enable or disable debug message output. Debug messages are disabled by default.

C.6.10 message-debug-enabled

Specifies a boolean value to enable or disable tracing of Diameter messages. This element is disabled by default.

C.6.11 application

Configures a particular Diameter application to run on the selected node. Oracle WebLogic Server SIP Container includes applications to support nodes that act as Diameter Sh, Ro, and Rf clients, Diameter relay agents, or Home Subscriber Servers (HSS). Note that the HSS application is a simulator that is provided only for development or testing purposes.

C.6.11.1 class-name

Specifies the application class file to load.

C.6.11.2 param*

Specifies one or more optional parameters to pass to the application class.

C.6.11.2.1 name Specifies the name of the application parameter.

C.6.11.2.2 value Specifies the value of the parameter.

C.6.12 peer-retry-delay

Specifies the number of seconds this node waits between retries to Diameter peers. The default value is 30 seconds.

C.6.13 allow-dynamic-peers

Specifies a boolean value that enables or disables dynamic peer configuration. Dynamic peer support is disabled by default. Oracle recommends enabling dynamic peers only when using the TLS transport, because no access control mechanism is available to restrict hosts from becoming peers.

C.6.14 request-timeout

Specifies the number of milliseconds to wait for an answer from a peer before timing out.

C.6.15 watchdog-timeout

Specifies the number of seconds used for the Diameter Tw watchdog timer.

C.6.16 supported-vendor-id+

Specifies one or more vendor IDs to be added to the `Supported-Version-Ids` AVP in the capabilities exchange.

C.6.17 include-origin-state

Specifies whether the node should include the origin state AVP in requests and answers.

C.6.18 peer+

Specifies connection information for an individual Diameter peer. You can choose to configure connection information for individual peer nodes, or allow any node to be dynamically added as a peer. Oracle recommends using dynamic peers only if you are using the TLS transport, because there is no way to filter or restrict hosts from becoming peers when dynamic peers are enabled.

When configuring Sh client nodes, the `peers` element should contain peer definitions for each Diameter relay agent deployed to your system. If your system does not use relay agents, you must include a peer entry for the Home Subscriber Server (HSS) in the system, as well as for all other engine tier nodes that act as Sh client nodes.

When configuring Diameter relay agent nodes, the `peers` element should contain peer entries for all Diameter client nodes that access the peer, as well as the HSS.

C.6.18.1 host

Specifies the host identity for a Diameter peer.

C.6.18.2 address

Specifies the listen address for a Diameter peer. If you do not specify an address, the host identity is used.

C.6.18.3 port

Specifies the TCP or TLS port number for this Diameter peer. The default port is 3868.

C.6.18.4 protocol

Specifies the protocol used by the peer. This element may be one of `tcp` or `sctp`.

C.6.19 route

Defines a realm-based route that this node uses when resolving messages.

When configuring Sh client nodes, you should specify a route to each Diameter relay agent node deployed in the system, as well as a `default-route` to a selected relay. If your system does not use relay agents, simply configure a single `default-route` to the HSS.

When configuring Diameter relay agent nodes, specify a single `default-route` to the HSS.

C.6.19.1 realm

The target realm used by this route.

C.6.19.2 application-id

The target application ID for the route.

C.6.19.3 action

An action type that describes the role of the Diameter node when using this route. The value of this element can be one of the following:

- none
- local
- relay
- proxy
- redirect

C.6.19.4 server+

Specifies one or more target servers for this route. Note that any server specified in the `server` element must also be defined as a `peer` to this Diameter node, or dynamic peer support must be enabled.

C.6.20 default-route

Defines a default route to use when a request cannot be matched to a configured route.

C.6.20.1 action

Specifies the default routing action for the Diameter node. See [Section C.6.19.3, "action"](#).

C.6.20.2 server+

Specifies one or more target servers for the default route. Any server you include in this element must also be defined as a `peer` to this Diameter node, or dynamic peer support must be enabled.

Startup Command Options

[Table D–1](#) provides a reference to the startup configuration options available to Oracle WebLogic Server SIP Container and other Oracle WebLogic Server SIP Container utilities.

Table D–1 Startup Command Options

Application	Startup Option Link
SIP Servlet Application Router	-Djavax.servlet.sip.dar.configuration -Djavax.servlet.sip.ar.spi.SipApplicationRouterProvider
Oracle WebLogic Server SIP Container	-Dwlss.udp.listen.on.ephemeral -Dwlss.udp.lb.masquerade -Dweblogic.management.discover -Dweblogic.RootDirectory
Oracle WebLogic SIP Server	-Dwlss.udp.listen.on.ephemeral -Dwlss.udp.lb.masquerade -Dweblogic.management.discover -Dweblogic.RootDirectory -DWLSS.SNMPAgentPort
Installer	-Djava.io.tmpdir
WlssEchoServer	-Dwlss.ha.echoserver.port -Dwlss.ha.echoserver.logfile -Dreplica.host.monitor.enabled -Dwlss.ha.heartbeat.interval -Dwlss.ha.heartbeat.count -Dwlss.ha.heartbeat.SoTimeout

For more information about these, and other options, see *Oracle Fusion Middleware Command Reference for Oracle WebLogic Server*.

Supported Platforms, Protocols, RFCs and Standards

The following sections describe how Oracle WebLogic Server SIP Container complies with various specifications and RFCs:

- Section E.1, "Supported Configurations"
- Section E.2, "Supported SIP Clients"
- Section E.3, "Supported Load Balancer"
- Section E.4, "Supported Databases"
- Section E.5, "Overview of Oracle WebLogic Server SIP Container Standards Alignment"
- Section E.6, "Java Sun Recommendation (JSR) Standards Compliance"
- Section E.7, "IETF RFC Compliance"
- Section E.8, "3GPP R6 Specification Conformance"

E.1 Supported Configurations

Oracle WebLogic Server SIP Container is supported for production use on the operating system, hardware, and JVM combinations described shown in http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html.

Note: Production deployment is supported only for Linux and UNIX platforms. Windows platforms are supported only for development purposes.

Oracle WebLogic Server SIP Container has similar requirements to Oracle WebLogic Server 10g Release 3. The following items are required in addition to the basic WebLogic Server requirements:

- Gigabit Ethernet connections are required between engine and SIP data tier servers for most production deployments.
- Dual NICs are required to provide failover capabilities in a production environment.
- Additional RAM is required to support the throughput requirements of most production installations.

- A compatible Load Balancer is required for production installations.

E.2 Supported SIP Clients

Oracle WebLogic Server SIP Container and the included sample applications support the following SIP softphone:

- Oracle Communicator (default client)

E.3 Supported Load Balancer

Oracle WebLogic Server SIP Container has been tested to run with F5 Networks, Inc. BIG-IP load balancer, versions 9.0.5 through 9.2.3 Build 34.3.

See the Oracle Technology Network (<http://www.oracle.com/global/index.html>) for information about support for other load balancer solutions.

E.4 Supported Databases

Oracle WebLogic Server SIP Container has been tested to run with the following databases for storing long-lived session data and for replicating call state information across regional sites (geographic persistence):

- Oracle Database 10g (10.2.0.3 or higher)
- Oracle Database 11g (11.1.0.6 or higher)

E.5 Overview of Oracle WebLogic Server SIP Container Standards Alignment

Oracle WebLogic Server SIP Container is developed with special attention to Internet Engineering Task Force and 3rd Generation Partnership Project specifications. Feature development is prioritized according to general market trends, both observed and predicted. In cases where certain specifications are obsolete or where Internet drafts are formalized as 'Request For Comments' standards, Oracle WebLogic Server SIP Container places priority on compliance with those specifications. In cases where specifications are part of a larger release plan, as with the 3GPP, Oracle prioritizes compliance with the latest ratified release (in this case, Release 6). This should not be presumed to mean that the product is not compliant with subsequent versions of component specifications, although this document does not summarize compliance with those specifications.

E.6 Java Sun Recommendation (JSR) Standards Compliance

Oracle WebLogic Server SIP Container is compliant with Java EE version 5.0 and the corresponding Java EE component specifications.

Oracle WebLogic Server SIP Container is further enhanced by the addition of a SIP Servlet container defined by JSR 289: "SIP Servlet API."

Oracle WebLogic Server SIP Container has executed all related Test Compatibility Kits (TCKs) and has met the formal requirements established by Sun Microsystems for formal public statements of compliance.

E.7 IETF RFC Compliance

The following table lists the Oracle WebLogic Server SIP Container level of compliance to common Internet Engineering Task Force (IETF) Requests for Comment (RFCs) and Internet drafts. The level of compliance is defined as follows:

- **Yes**—Indicates that Oracle WebLogic Server SIP Container directly supports the feature or specification.
- **Yes (Platform)**—Indicates Oracle WebLogic Server SIP Container can host applications or components that implement the RFC. However, the RFC or feature has no impact on the transaction layer of the protocol or on the behavior of the SIP Servlet container.

Table E-1 Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
761	DoD Standard Transmission Control Protocol	Yes	See http://www.ietf.org/rfc/rfc761.txt
768	User Datagram Protocol	Yes	See http://www.ietf.org/rfc/rfc768.txt
1157	A Simple Network Management Protocol (SNMP)	Yes	Oracle WebLogic Server SIP Container supports SNMP V2c traps. See http://www.ietf.org/rfc/rfc1157.txt
1847	Security Multiparts for MIME: Multipart/Signed and Multipart/Encrypted	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that consume or generate signed or encrypted multipart MIME objects. See http://www.ietf.org/rfc/rfc1847.txt
1901	Introduction to Community-based SNMPv2	Yes	Oracle WebLogic Server SIP Container supports SNMP V2c traps. See http://www.ietf.org/rfc/rfc1901.txt
1905	Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)	Yes	Oracle WebLogic Server SIP Container supports SNMP V2c traps. See http://www.ietf.org/rfc/rfc1905.txt
1906	Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)	Yes	Oracle WebLogic Server SIP Container supports SNMP over both TCP and UDP. See http://www.ietf.org/rfc/rfc1906.txt
1907	Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)	Yes (Platform)	See http://www.ietf.org/rfc/rfc1907.txt
2183	Communicating Presentation Information in Internet Messages: The Content-Disposition Header Field	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc2183.txt

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
2246	The TLS Protocol Version 1.0	Yes	Oracle WebLogic Server SIP Container supports TLS. See http://www.ietf.org/rfc/rfc2246.txt
2327	SDP: Session Description Protocol	Yes	Oracle WebLogic Server SIP Container supports applications that consume or generate SDP. See http://www.ietf.org/rfc/rfc2327.txt
2460	Internet Protocol, Version 6 (IPv6) Specification	Yes	See http://www.ietf.org/rfc/rfc2460.txt
2543	SIP: Session Initiation Protocol (v1)	Yes	Oracle WebLogic Server SIP Container supports backward compatibility as described in this specification. See http://www.ietf.org/rfc/rfc2543.txt
2616	Hypertext Transfer Protocol -- HTTP 1.1	Yes	See http://www.ietf.org/rfc/rfc2616.txt
2617	HTTP Authentication: Basic and Digest Access Authentication	Yes	See http://www.ietf.org/rfc/rfc2617.txt
2782	A DNS RR for specifying the location of services (DNS SRV)	Yes	See http://www.ietf.org/rfc/rfc2782.txt
2806	URLs for Telephone Calls	Yes	See http://www.ietf.org/rfc/rfc2806.txt
2848	The PINT Service Protocol: Extensions to SIP and SDP for IP Access to Telephone Call Services	Yes (Platform)	Note that implementing PINT services implies a pre-IMS architecture. Although Oracle favors the 3GPP/TISPAN architecture and approach to class 4/5 Service Emulation and does not advocate PINT, it is possible to implement PINT service elements using Oracle WebLogic Server SIP Container. See http://www.ietf.org/rfc/rfc2848.txt
2916	E.164 number and DNS	Yes	See http://www.ietf.org/rfc/rfc2916.txt
2960	Stream Control Transmission Protocol	Yes	SCTP supported only for Diameter traffic. See http://www.ietf.org/rfc/rfc2960.txt
2976	The SIP INFO Method	Yes	See http://www.ietf.org/rfc/rfc2976.txt
3204	MIME media types for ISUP and QSIG Objects	Yes (Platform)	Oracle WebLogic Server SIP Container does not directly consume or generate ISUP and QSIG objects, but it supports applications that consume or generate these objects. See http://www.ietf.org/rfc/rfc3204.txt
3261	SIP: Session Initiation Protocol	Yes	See http://www.ietf.org/rfc/rfc3261.txt

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
3262	Reliability of Provisional Responses in the Session Initiation Protocol (SIP)	Yes	See http://www.ietf.org/rfc/rfc3262.txt
3263	Session Initiation Protocol (SIP): Locating SIP Servers	Yes	See http://www.ietf.org/rfc/rfc3263.txt
3264	An Offer/Answer Model with Session Description Protocol (SDP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3264.txt
3265	Session Initiation Protocol (SIP)-Specific Event Notification	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3265.txt
3266	Support for IPv6 in Session Description Protocol (SDP)	Yes	See http://www.ietf.org/rfc/rfc3266.txt
3268	Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS)	Yes (Platform)	Oracle WebLogic Server SIP Container supports cryptographic services, but specific algorithms that are used are subject to local availability and export control. See http://www.ietf.org/rfc/rfc3268.txt
3311	The Session Initiation Protocol (SIP) UPDATE Method	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3311.txt
3312	Integration of Resource Management and Session Initiation Protocol (SIP).	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3312.txt
3313	Private Session Initiation Protocol (SIP) Extensions for Media Authorization	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3313.txt
3323	A Privacy Mechanism for the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3323.txt
3325	Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks	Yes	See http://www.ietf.org/rfc/rfc3325.txt
3326	The Reason Header Field for the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3326.txt

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
3327	Session Initiation Protocol (SIP) Extension Header Field for Registering Non-Adjacent Contacts.	Yes (Platform)	See http://www.ietf.org/rfc/rfc3327.txt
3351	User Requirements for the Session Initiation Protocol (SIP) in Support of Deaf, Hard of Hearing and Speech-impaired Individuals	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3351.txt
3372	Session Initiation Protocol for Telephones (SIP-T): Context and Architectures	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3372.txt
3388	Grouping of Media Lines in the Session Description Protocol (SDP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3388.txt
3420	Internet Media Type message/sipfrag		See http://www.ietf.org/rfc/rfc3420.txt
3428	Session Initiation Protocol (SIP) Extension for Instant Messaging	Yes	See http://www.ietf.org/rfc/rfc3428.txt
3455	Private Header (P-Header) Extensions to the Session Initiation Protocol (SIP) for the 3rd-Generation Partnership Project (3GPP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3455.txt
3489	STUN-Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)	Yes	See http://www.ietf.org/rfc/rfc3489.txt
3515	The Session Initiation Protocol (SIP) Refer Method.	Yes	See http://www.ietf.org/rfc/rfc3515.txt
3524	Mapping of Media Streams to Resource Reservation Flows	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3524.txt

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
3556	Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3556.txt
3578	Mapping of Integrated Services Digital Network (ISDN) User Part (ISUP) Overlap Signalling to the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification, but it does not provide an ISUP interface. See http://www.ietf.org/rfc/rfc3578.txt
3581	An Extension to the Session Initiation Protocol (SIP) for Symmetric Response Routing	Yes	See http://www.ietf.org/rfc/rfc3581.txt
3589	Diameter Command Codes for Third Generation Partnership Project (3GPP) Release 5	Yes	See http://www.ietf.org/rfc/rfc3589.txt
3588	Diameter Base Protocol	Yes	See http://www.ietf.org/rfc/rfc3588.txt
3605	Real Time Control Protocol (RTCP) attribute in Session Description Protocol ((SDP)		See http://www.ietf.org/rfc/rfc3605.txt
3608	Session Initiation Protocol (SIP) Extension Header Field for Service Route Discovery During Registration.	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification, but it does not provide a means of storing the ServiceRoute established during registration. This functionality can be implemented as part of the application. See http://www.ietf.org/rfc/rfc3608.txt
3665	Session Initiation Protocol (SIP) Basic Call Flow Examples.	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3665.txt
3666	Session Initiation Protocol (SIP) Public Switched Telephone Network (PSTN) Call Flows	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3666.txt
3680	A Session Initiation Protocol (SIP) Event Package for Registrations	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3680.txt

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
3689	General Requirements for Emergency Telecommunication Service (ETS)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3689.txt
3690	IP Telephony Requirements for Emergency Telecommunication Service (ETS)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3690.txt
3702	Authentication, Authorization, and Accounting Requirements for the Session Initiation Protocol (SIP)	Yes	Oracle WebLogic Server SIP Container version supports JDBC and LDAP. See http://www.ietf.org/rfc/rfc3702.txt
3725	Best Current Practices for Third Party Call Control (3pcc) in the Session Initiation Protocol (SIP)	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3725.txt
3761	The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)	Yes	See http://www.ietf.org/rfc/rfc3761.txt
3764	Enumservice Registration for Session Initiation Protocol (SIP) Addresses-of-Record	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3764.txt
3824	Using E.164 numbers with the Session Initiation Protocol (SIP)	Yes	See http://www.ietf.org/rfc/rfc3824.txt
3840	Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3840.txt
3841	Caller Preferences for the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3841.txt
3853	S/MIME Advanced Encryption Standard (AES) Requirement for the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3853.txt
3856	Presence Event Package for the Session Initiation Protocol (SIP)	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3856.txt

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
3857	Watcher Information Event Template-Package for the Session Initiation Protocol (SIP)	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3857.txt
3858	Extensible Markup Language (XML) Based Format for Watcher Information	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3858.txt
3863	Presence Information Data Format (PIDF)	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3863.txt
3891	The Session Initiation Protocol (SIP) 'Replaces' Header	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3891.txt
3892	The Session Initiation Protocol (SIP) Referred-By Mechanism	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3892.txt
3893	Session Initiation Protocol (SIP) Authenticated Identity Body (AIB) Format	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3893.txt
3903	Session Initiation Protocol (SIP) Extension for Event State Publication	Yes	See http://www.ietf.org/rfc/rfc3903.txt
3911	The Session Initiation Protocol (SIP) "Join" Header	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3911.txt
3959	The Early Session Disposition Type for the Session Initiation Protocol (SIP)		See http://www.ietf.org/rfc/rfc3959.txt
3960	Early Media and Ringing Tone Generation in the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc3960.txt
3966	The tel URI for Telephone Numbers	Yes	See http://www.ietf.org/rfc/rfc3966.txt
4028	Session Timers in the Session Initiation Protocol (SIP)	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc4028.txt
4032	Update to the Session Initiation Protocol (SIP) Preconditions Framework	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc4032.txt

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
4244	An Extension to the Session Initiation Protocol (SIP) for Request History Information	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc4244.txt
4320	Actions Addressing Identified Issues with the Session Initiation Protocol's (SIP) Non-INVITE Transaction		See http://www.ietf.org/rfc/rfc4320.txt
4321	Problems Identified Associated with the Session Initiation Protocol's (SIP) Non-INVITE Transaction		See http://www.ietf.org/rfc/rfc4321.txt
4474	Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)	Yes	See http://www.ietf.org/rfc/rfc4474.txt .
4479	Data Model for Presence	Yes	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc4479.txt .
4480	RPID: Rich Presence Extensions to the Presence Information Data Format	Implicitly Supported	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc4480.txt .
4481	Timed Presence Extensions to the Presence Information Data Format (PIDF) to Indicate Status Information for Past and Future Time Intervals	Implicitly Supported	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc4481.txt .
4483	A Mechanism for Content Indirection in Session Initiation Protocol (SIP) Messages	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See http://www.ietf.org/rfc/rfc4483.txt .
4825	The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)	Partially Supported	No support for partial document manipulation (XPath)

Table E-1 (Cont.) Oracle WebLogic Server SIP Container IETF Compliance

RFC or Specification Number	Title	Compliant?	Additional Information
4827	Extensible Markup Language (XML) Configuration Access Protocol (XCAP) Usage for Manipulating Presence Document Contents	Yes	Replaces draft-ietf-simple-xcap-pidf-manipulation-usage-02
draft-levy-sip-diversion-08	Diversion Indication in SIP	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See https://datatracker.ietf.org/public/dindex.cgi?command=id_detail&id=6002
draft-donovan-mmusic-183-00	SIP 183 Session Progress Message Draft	Yes (Platform)	Oracle WebLogic Server SIP Container supports applications that conform to this specification. See https://datatracker.ietf.org/public/dindex.cgi?command=id_detail&id=4308
draft-ietf-sip-gruu-15	Obtaining and Using Globally Routable User Agent (UA) URIs (GRUU) in the Session Initiation Protocol (SIP)	Yes	See http://www.ietf.org/internet-drafts/draft-ietf-sip-gruu-15.txt
Presence Authorization Rules	draft-ietf-simple-presence-rules-04	Yes	http://tools.ietf.org/html/draft-ietf-simple-presence-rules-10
OMA extensions to the presence data model	OMA-TS-Presence_SIMPLE-V1_0-20051122-C	Yes	
OMA extensions to geopriv common policy	OMA-TS-XDM_Core-V1_0-20051122-C	Yes	
OMA extensions to presence rules	OMA-TS-Presence_SIMPLE_XDM-V1_0-20051122-C	Yes	

E.8 3GPP R6 Specification Conformance

Table E-2 summarizes the ability of Oracle WebLogic Server SIP Container to support implementation of the enablers or application functions identified by each applicable 3GPP Release 6 specification.

Other than the exceptions noted, Oracle WebLogic Server SIP Container does not impose any restrictions on implementing applications or functions that are compliant with those associated with the Application Server entity described in the specification. In some cases, applications must implement support for SIP methods or headers. The default behavior of the Oracle WebLogic Server SIP Container Sip Servlet Container is

to pass unrecognized headers, request methods and payloads to SIP Servlets using normal SIP Servlet API procedures.

Table E-2 3GPP R6 Specification Conformance

Specification	Comments
3GPP TS 23.228: "IP Multimedia Subsystem (IMS); Stage 2 (Release 6)"	<ul style="list-style-type: none"> ■ No comments.
3GPP TS 24.229: "IP Multimedia Call Control Protocol based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3 (Release 6)"	<ul style="list-style-type: none"> ■ Oracle WebLogic Server SIP Container does not enforce the requirement that only one p-charging-function-address header per SIP request as described in sub-section 5.7.1.2. Oracle WebLogic Server SIP Container does enforce uniqueness. ■ Oracle WebLogic Server SIP Container does not provide privacy support as described in sub-section 5.7.3.
3GPP TS 23.141: "Presence Service; Architecture and Functional description (Release 6)"	<ul style="list-style-type: none"> ■ Oracle WebLogic Server SIP Container does not natively support the Ph (MAP), Pl (LIF-MLP), Px (DIAMETER Cx/Dx), Pg (phase 4, 3GPP Release 5), Pc (CAMEL phase 4, 3GPP Release 5) or Pr, Pk and Pp (RADIUS) reference points. ■ Oracle WebLogic Server SIP Container does not support IPv6 as required for the Presence User Agent (Peu) reference point as required in sub-section 4.3.1.
3GPP TS 23.218: "IP Multimedia (IM) session handling; IM call model; Stage 2 (Release 6)"	<ul style="list-style-type: none"> ■ No comments.
3GPP TS 24.247 "Messaging using the IP Multimedia (IM) Core Network (CN) subsystem; Stage 3 (Release 6)"	<ul style="list-style-type: none"> ■ Oracle WebLogic Server SIP Container does not provide support for the Message Session Relay Protocol (MSRP), although it is presumed that an MSRP relay will typically be implemented as a Media Resource Function in the IMS architecture.
3GPP TS 24.841: "Presence service based on Session Initiation Protocol (SIP); Functional models, information flows and protocol details (Release 6)"	<ul style="list-style-type: none"> ■ Oracle WebLogic Server SIP Container does not provide support for IETF RFC 3310: "Hypertext Transfer Protocol (HTTP) Digest Authentication Using Authentication and Key Agreement (AKA)".
3GPP TS 24.109: "Bootstrapping interface (Ub) and Network application function interface (Ua); Protocol details (Release 6)"	<ul style="list-style-type: none"> ■ Oracle WebLogic Server SIP Container does not provide support for IETF RFC 3310: "Hypertext Transfer Protocol (HTTP) Digest Authentication Using Authentication and Key Agreement (AKA)". ■ Oracle WebLogic Server SIP Container supports the 'X-3GPP-Asserted-Identity extension-header' for use in applying access control and authorization constraints within the integrated security framework.
3GPP TS 29.328: "IP Multimedia Subsystem (IMS) Sh interface; Signalling flows and message contents"	<ul style="list-style-type: none"> ■ No comments.
3GPP TS 29.329: "Sh interface based on the Diameter protocol; Protocol details"	<ul style="list-style-type: none"> ■ No comments.

Table E-2 (Cont.) 3GPP R6 Specification Conformance

Specification	Comments
3GPP TS 32.299: "Telecommunication management; Charging management; Diameter charging applications"	<ul style="list-style-type: none">■ No comments.
3GPP TS 33.222: "Generic Authentication Architecture (GAA); Access to network application functions using Hypertext Transfer Protocol over Transport Layer Security (HTTPS) (Release 6)"	<ul style="list-style-type: none">■ Oracle WebLogic Server SIP Container supports the Application Server role in the GAA.

Using Oracle WebLogic SIP Container Export/Import

The following sections describe how to use Oracle WebLogic Server SIP Container Export and Import functions:

- [Section F.1, "Export"](#)
- [Section F.2, "Import"](#)

Note: All examples in this document use Linux commands; adapt commands as appropriate for other operating systems.

F.1 Export

This section details the steps you must take in order to export your OWLSC data.

F.1.1 Export the Database Data from the Current Environment

Note: The first four steps in this procedure are performed on your RDBMS machine.

Export data by using one of the Oracle database export utilities. Before executing the export commands, on the existing database machine, create a directory on the file system where the exported data dump files and log files will be stored.

```
$ mkdir ~/owlcs_data_pump_dir
```

1. Change directory to the \$ORACLE_HOME/bin directory of your database installation and execute the following commands. These commands can be executed as the *sys* user or the *system* user.
2. Create DATA_PUMP_DIR in your database using the following commands:

- \$./sqlplus sys@<db_service> as sysdba
- Enter password for sys user at the prompt.
- At the SQL prompt, execute the following commands:

```
SQL> create or replace directory DATA_PUMP_DIR as '<full_path_to_the_data_pump_dir_created_above>';
SQL> commit;
SQL> quit
```

3. Export data in the subscriber data services schema.

- `$./expdp "'sys@<db_service> as sysdba'" SCHEMAS=<current_prefix>_orasdpds DIRECTORY=DATA_PUMP_DIR DUMPFILE=<current_prefix>_orasdpds.dmp LOGFILE=<current_prefix>_orasdpds.log`
- Enter password for sys user at the prompt.

Note: The value of the DIRECTORY parameter in the command line above should be specified as DATA_PUMP_DIR and not as the actual directory on the file system. When the command is done, verify that two files, one corresponding to DUMPFILE and the other corresponding to LOGFILE are created in your data pump directory.

4. Export data in the XDMS schema.

- `$./expdp "'sys@<db_service> as sysdba'" SCHEMAS=<current_prefix>_orasdpds DIRECTORY=DATA_PUMP_DIR DUMPFILE=<current_prefix>_orasdpds.dmp LOGFILE=<current_prefix>_orasdpds.log`
- Enter password for sys user at the prompt.

Note: The value of the DIRECTORY parameter in the command line above should be specified as DATA_PUMP_DIR and not as the actual directory on the file system. When the command is done, verify that two files, one corresponding to DUMPFILE and the other corresponding to LOGFILE are created in your data pump directory.

Exporting of the Location Service schema is not required, since location data is recreated when clients sign in to the server.

5. Complete OWSM Policy Migration. Start WLST by running the following command:

```
$ORACLE_HOME/common/bin/wlst.sh.
```

6. Connect to the local WLS instance by running the following command:

```
wls:/offline> connect('weblogic','weblogic','127.0.0.1:7001')
```

In this example, 'weblogic'/'weblogic' are sample WLS admin username/password. Replace them with the real values in your environment. The port may change if you have another instance of WLS running (this is the WLS AdminServer port).

7. Run the following WLST commands to export the policies and assertion templates. Replace "wlcs_server1" with your OWLSC instance name.

```
wls:/base_domain/serverConfig>
exportMetadata(application='wsm-pm',server='wlcs_
server1',docs='/assertiontemplates/**',toLocation='/tmp/owsmexport/')
wls:/base_domain/serverConfig>
exportMetadata(application='wsm-pm',server='wlcs_
server1',docs='/policies/**',toLocation='/tmp/owsmexport/')
```

8. Exit the WLST command line tool by running the following command:

```
wls:/base_domain/serverConfig> exit()
```

9. (Performed on the machine in which the OWLSC instance is installed [your middleware machine]) Next, export the *Credential Store*. OWSM stores client policy username and password credentials and keystore passwords in the credential store. Copy `<domain>/config/fmwconfig/cwallet.sso` from current machine to the new machine. If this is already performed as part of OPSS migration, then you can omit this step.
10. Export UMS-related details to the new environment (if UMS is installed in your environment). Export UMS-related details to your new environment (if UMS is installed in your environment). If the User Messaging Service is used in the current and new environments, perform the following step:
 - Start WLST by running the following command:

```
$ORACLE_HOME/common/bin/wlst.sh
```

Note: This is the ORACLE_HOME on the middleware instance. By default that is `<middleware_home>/as11gr1wlcs1` directory, where `<middleware_home>` is the directory where OWLSC is installed.

11. Run the following WLST commands to download the user messaging preferences from the backend database to the specified xml file:

```
wls:/offline> manageUserMessagingPrefs(operation='download',
filename='/tmp/userprefs-dump.xml', url='t3://localhost:8001',
username='weblogic', password='weblogic')
```

Note: In the above sample 'weblogic'/'weblogic' are sample WLS admin username/password combinations. Replace them with the real values in your environment. 8001 is the Managed Server Port where UMS is running. Replace it accordingly with the appropriate value.

12. Exit the WLST command line tool by running the following command:

```
wls:/offline> exit()
```

F.2 Import

This section details the steps you must take in order to import your OWLSC data.

1. Import Database data into the Production Environment. This is achieved by using one of the Oracle database import utilities. Before executing the import commands, on the new database machine, create a directory on the file system where the data dump files to be imported and log files will be stored:

```
$ mkdir ~/owlcs_data_pump_dir
```

2. Copy (ftp or remote copy) all dmp files from the data pump directory on your old database machine to the above directory.
3. Change directory to the `$ORACLE_HOME/bin` of your production database installation and execute the following commands. These commands can be executed as the `sys` user or the `system` user.

- a. Create DATA_PUMP_DIR in your production database with the following commands:

```
$ ./sqlplus sys@<db_service> as sysdba
```

Enter password for sys user at the prompt

At the SQL prompt, execute the following command:

```
SQL> create or replace directory DATA_PUMP_DIR as '<full_path_to_the_data_pump_dir_created_above>'
SQL> commit;
SQL> quit
```

- b. Drop all sequences in the <production_prefix>_orasdpds schema with the following commands:

```
# $ ./sqlplus sys@<db_service> as sysdba
```

Enter password for sys user at the prompt

At the SQL prompt, execute the following commands:

```
SQL> alter session set current_schema=<production_prefix>_orasdpds;
SQL> drop sequence ACCOUNT_SEQ;
SQL> drop sequence CREDENTIALS_SEQ;
SQL> drop sequence PRIVATE_IDENTITY_SEQ;
SQL> drop sequence PUBLIC_IDENTITY_SEQ;
SQL> drop sequence REALM_SEQ;
SQL> drop sequence ROLE_SEQ;
SQL> commit;
SQL> quit
```

- c. Import data into the subscriber data services schema

```
$ ./impdp "sys@<db_service> as sysdba" REMAP_SCHEMA=<test_prefix>_orasdpds:<production_prefix>_orasdpds TABLE_EXISTS_ACTION=REPLACE DIRECTORY=DATA_PUMP_DIR DUMPFILE=<test_prefix>_orasdpds.dmp LOGFILE=<production_prefix>_orasdpds.log
```

Enter password for sys user at the prompt.

Note that the value of the DIRECTORY parameter in the command line above should be specified as DATA_PUMP_DIR and not as the actual directory on the file system. Ignore error that user already exists.

- d. Export data into the XDMS schema

```
$ ./expdp "sys@<db_service> as sysdba" REMAP_SCHEMA=<test_prefix>_orasdpdxms:<production_prefix>_orasdpdxms TABLE_EXISTS_ACTION=REPLACE DIRECTORY=DATA_PUMP_DIR DUMPFILE=<test_prefix>_orasdpdxms.dmp LOGFILE=<production_prefix>_orasdpdxms.log
```

Enter password for sys user at the prompt.

Note that the value of the DIRECTORY parameter in the command line above should be specified as DATA_PUMP_DIR and not as the actual directory on the file system. Ignore error that user already exists.

Importing of the Location Service schema is not required, since location data is recreated when clients sign in to the server.

4. Copy the /tmp/owsmexport directory to the new machine.
5. Start WLST by running the following command:

```
$ORACLE_HOME/common/bin/wlst.sh
```

6. Connect to the local WLS instance by running the following command:

```
wls:/offline> connect('weblogic','weblogic','127.0.0.1:7001')
```

Note: In the above sample 'weblogic'/'weblogic' are sample WLS admin username/password. Replace them with the real values in your environment. The port may change if you have another instance of WLS running (this is the WLS AdminServer port).

7. Run the following WLST commands to replace the policies and assertion templates. Replace *wlcs_server1* with your OWLSC instance name:

```
wls:/base_domain/serverConfig>
deleteMetadata(application='wsm-pm',server='wlcs_
server1',docs='/assertiontemplates/**')
wls:/base_domain/serverConfig>
deleteMetadata(application='wsm-pm',server='wlcs_server1',docs='/policies/**')
wls:/base_domain/serverConfig>
importMetadata(application='wsm-pm',server='wlcs_
server1',docs='/assertiontemplates/**',fromLocation='/tmp/owsmexport/')
wls:/base_domain/serverConfig>
importMetadata(application='wsm-pm',server='wlcs_
server1',docs='/policies/**',fromLocation='/tmp/owsmexport/')
```

8. Exit the WLST command line tool by running the following command:

```
wls:/base_domain/serverConfig> exit()
```

Note: Regarding importing OWSM Keystore:

- Private keys will differ between current and new environments. So, they do not need to be migrated.
- Public keys, intermediate certs, and root certs may be migrated from current to new environments. Use `java keytool export` and `import` commands to move them. After doing so, review if these certs are applicable in the new environment based on the clients invoking the services in the new environment.
- Review if new public keys of client certs, intermediate CA certs or root CA certs must be added to the new keystore based on the clients invoking the services in the new environment.

For more details, see *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

9. Go to the new machine.

10. Start WLST by running the following command:

```
$ORACLE_HOME/common/bin/wlst.sh
```

11. Run the following WLST command to upload the User Messaging Preferences from file to the backend database:

```
wls:/offline> manageUserMessagingPrefs(operation='upload',  
filename='/tmp/userprefs-dump.xml', url='t3://localhost:8001',  
username='weblogic', password='weblogic')
```

Note: In the above sample 'weblogic'/'weblogic' are sample WLS admin username/password. Replace them with the real values in your environment. 8001 is the Managed Server Port where UMS is running. Replace it accordingly with the appropriate value.

12. Observe the message displayed for successful upload. Exit the WLST command line tool by running the following command:

```
wls:/offline> exit()
```

Note: For different options on performing download or upload, execute help ('manageUserMessagingPrefs') at the wls:/offline> prompt. User devices provisioned in the LDAP store are dynamic; the assumption is that both the current and new environments will point to the same LDAP store or will be re-configured to use the same set of information.

Index

A

Administration Console, 1-4
 configuring Container properties with, 2-1
Administration Server
 best practices, 1-6

C

channels
 SIP and SIPS, 3-4
 TCP and TLS, 3-7
configuration
 locking and persisting, 2-2
configuration locks, managing, 2-3
configurations
 single NIC, 3-9
custom channels
 properties, 3-5

D

datatier.xml, 4-2
Default Application Router (DAR), 1-3
deployed SIP Servlets
 upgrading, 8-1
Diameter
 configuration, 1-3
diameter
 configuration, C-1
 domain, 6-2
 nodes and relays, 6-1
diameter applications
 configuring, 6-9
Diameter Console, 6-4
diameter nodes
 configuring, 6-7
Diameter, resource, 1-2
DNS
 support, 3-3
DNS name, 3-1

E

engine tier, 7-3
Enterprise Deployment, 3-8
export and import

Oracle WebLogic Server SIP Container, F-1

G

Geo-Redundancy, 4-9

I

IP_ANY, 3-7
IPv4, 3-2
IPv6, 3-2

J

JMX, 2-3
JMX-compliant MBeans, 1-5
jrocket, 5-31

L

Listen Address, 3-1
Listen Port, 3-1
load balancer, 7-2
 SIP-aware, 1-2
load balancers
 configuring, 3-2, 3-13
 multi-homed, 3-2
log levels, setting, 1-5

M

MBean
 communication services, 2-4
MBeans
 creating and deleting, 2-5

N

NAT (Network Address Translation), 3-14
network resources
 managing, 3-1
NTP
 configuring, 2-6

O

Oracle WebLogic Server SIP Container

- architecture, 7-1
- common configuration tasks, 1-7
- configuration, 3-14
- custom resources, 1-2
- export and import, F-1
- introduction, 1-1
- SIP Data Tier partitions, 4-1

S

- SCTP, 6-6
- servers
 - stopping and starting, 1-5
- SIP Container
 - configuring, 2-1
- SIP Data Tier
 - configuration reference, B-1
- SIP data tier, 7-3
- SIP Servlet Container
 - configuration reference, A-1
- SIP Servlet container, 1-2
- SIP timers, 2-6
- startup
 - command options, D-1

T

- timer affinity, 2-5
- timers, 2-5
- troubleshooting, 5-1

W

- WebLogic Scripting Tool (WLST), 1-5
- WebLogic Server platform, 1-1
- WLST
 - configuring, 2-4
 - invoking, 2-4