

Oracle® Communications Messaging Server

System Administrator's Guide

Release 8.0.1

E63711-02

February 2016

Copyright © 2015, 2016, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Preface	xxix
Audience	xxix
Related Documents	xxix
Document Revision History	xxix
Documentation Accessibility	xxx
Part I Monitoring and Managing Messaging Server	
1 Messaging Server System Administration Overview	
About Messaging Server	1-1
Messaging Server Configuration Methods	1-1
Overview of Messaging Server Administration Tasks	1-1
About Messaging Server Administration Tools	1-2
Directory Placeholders Used in This Guide	1-2
2 Overview of Messaging Server Unified Configuration	
What Is Messaging Server Unified Configuration?	2-1
Unified Configuration Files	2-2
Enabling Unified Configuration in Messaging Server	2-4
To Determine if Unified Configuration Is Deployed	2-4
Understanding Unified Configuration Limitations	2-5
Using the Repository of Previous Configurations	2-5
To List Configurations	2-5
To Compare Configurations	2-5
Using Legacy Configuration Tools with Unified Configuration	2-5
Separating Roles and Instances	2-6
More About Unified Configuration Options	2-6
Options That Have Passwords	2-7
Restricted Options	2-7
Obsolete Options	2-7
Option Relationships	2-8
Unified Configuration Option Names	2-8
Example of Legacy Configuration and Unified Configuration	2-9
Using Recipes	2-10
To Run a Recipe	2-10

Helpful Commands	2-10
To Show Settings.....	2-11
To Get Help.....	2-11

3 Configuring General Messaging Capabilities

Modifying Your Passwords	3-1
Managing Mail Users, Mailing Lists and Domains	3-2
Overview of Messaging Server and LDAP	3-2
To Remove a User from Messaging Server by Using Delegated Administrator.....	3-3
To Remove a Domain from Messaging Server using Delegated Administrator.....	3-3
Starting and Stopping Services	3-3
To Start and Stop Messaging Server Services in an HA Environment.....	3-3
To Start and Stop Messaging Server Services in a non-HA Environment.....	3-4
To Start Up, Shut Down, or View the Status of Messaging Services.....	3-5
Starting and Stopping a Messaging Server Running in MTA-only Mode.....	3-6
Automatic Restart of Failed or Unresponsive Services	3-7
Overview of Messaging Server Monitoring Processes	3-7
Automatic Restart in High Availability Deployments	3-8
Scheduling Automatic Tasks	3-9
Overview of Scheduling Automatic Tasks.....	3-9
Scheduler Examples.....	3-9
Pre-defined Automatic Tasks	3-10
Configuring a Greeting Message	3-10
To Create a New User Greeting	3-10
To Set a Per-Domain Greeting Message	3-10
Setting a User-Preferred Language	3-12
Overview of Setting a User-Preferred Language	3-12
To Set a Domain Preferred Language	3-13
To Specify a Site Language	3-13
Encryption Settings	3-13
Setting a Failover LDAP Server	3-13

4 Configuring and Administering Multiplexor Services

Multiplexor Services in Unified Configuration Overview	4-1
Multiplexor Services	4-2
Multiplexor Benefits	4-2
About Messaging Multiplexor	4-3
How the Messaging Multiplexor Works	4-4
Encryption (SSL) Option	4-5
Certificate-Based Client Authentication	4-5
To Enable Certificate-based Authentication for Your IMAP or POP Service	4-5
User Pre-Authentication.....	4-6
MMP Virtual Domains	4-6
About SMTP Proxy	4-7
Setting Up the Messaging Multiplexor	4-7
Before You Configure MMP	4-8
Multiplexor Configuration.....	4-8

To Configure the MMP.....	4-8
Multiplexor Configuration Options	4-9
Starting the Multiplexor	4-9
Modifying an Existing MMP	4-9
Configuring MMP with SSL or Client Certificate-Based Login.....	4-9
To Configure MMP with SSL	4-9
To Configure MMP with Client Certificate-based Login	4-10
A Sample Topology	4-10
MMP Tasks.....	4-11
To Configure Mail Access with MMP	4-11
To Set a Failover MMP LDAP Server.....	4-11

5 MTA Concepts

The MTA Functionality.....	5-1
MTA Architecture and Message Flow Overview.....	5-3
Dispatcher and SMTP Server (Slave Program).....	5-3
The Dispatcher.....	5-5
Creation and Expiration of Server Processes	5-5
To Start and Stop the Dispatcher	5-6
Rewrite Rules.....	5-6
Channels.....	5-7
Master and Slave Programs.....	5-7
Channel Message Queues	5-8
Channel Definitions	5-9
The MTA Directory Information	5-10
The Job Controller.....	5-10
To Start and Stop the Job Controller	5-12
On Demand Mail Relay	5-12
Priority Message Handling.....	5-13

6 MTA Address Translation and Routing

The Direct LDAP Algorithm and Implementation.....	6-1
Domain Locality Determination	6-1
Rewrite Rule Machinery	6-1
Domain Map Determination of Domain Locality	6-2
Caching of Domain Locality Information	6-3
Error Handling	6-3
Pattern for Domain Check Rewrite Rule	6-4
Putting It All Together	6-4
Alias Expansion of Local Addresses	6-4
Alias Checking with LDAP URLs	6-5
The \$V Metacharacter.....	6-5
Calling a Mapping From a URL.....	6-6
The \$R Metacharacter.....	6-7
Determining the Attributes to Fetch	6-7
Handling LDAP Errors	6-8

Sanity Checks on the LDAP Result	6-8
Support for Vanity Domains	6-8
Support for Catchall Addresses	6-8
Processing the LDAP Result	6-9
Object Class Check.....	6-9
Entry Status Checks	6-10
UID Check.....	6-11
Message Capture.....	6-11
Seeding the Reversal Cache.....	6-11
Mail Host and Routing Address.....	6-12
Miscellaneous Attribute Support.....	6-12
Delivery Options Processing.....	6-13
Additional Metacharacters for Use in Delivery Options	6-15
Delivery Option Defaults.....	6-16
Start and End Date Checks	6-16
Optin and Presence Attributes.....	6-16
Sieve Filter Handling.....	6-17
Deferred Processing Control	6-17
Group Expansion Attributes	6-17
To Modify Group Membership Attribute Syntax	6-20
Address Reversal	6-21
Asynchronous LDAP Operations.....	6-22
Settings Summary	6-23
Processing Multiple Different LDAP Attributes with the Same Semantics	6-23

7 About MTA Services

MTA Configuration Overview	7-1
Compiling the MTA Configuration.....	7-2
Changes for Compiled MTA Configurations.....	7-2
Recompiling the MTA Configuration	7-2

8 Message Transfer Agent Command-line Utilities

MTA Commands	8-1
Command Descriptions	8-2
imsimta cache.....	8-2
imsimta chbuild.....	8-3
imsimta cnbuild.....	8-5
imsimta counters	8-7
imsimta crdb	8-8
imsimta find	8-11
imsimta kill.....	8-12
imsimta process	8-12
imsimta program.....	8-12
imsimta purge.....	8-14
imsimta qclean.....	8-14
imsimta qm.....	8-15
imsimta qm Options	8-26

imsimta qtop	8-28
imsimta refresh	8-29
imsimta reload	8-30
imsimta renamedb	8-31
imsimta restart	8-31
imsimta return	8-32
imsimta run	8-32
imsimta shutdown	8-33
imsimta start	8-33
imsimta stop	8-34
imsimta submit	8-34
imsimta test	8-35
imsimta test -domain	8-42
imsimta version	8-45
imsimta view	8-45

9 Configuring Rewrite Rules

Before You Begin	9-1
Editing Rewrite Rules in Unified Configuration	9-1
Rewrite Rule Structure	9-3
Rewrite Rule Patterns and Tags	9-4
A Rule to Match Percent Hacks	9-6
A Rule to Match Bang-Style (UUCP) Addresses	9-6
A Rule to Match Any Address	9-6
Tagged Rewrite Rule Sets	9-6
Rewrite Rule Templates	9-7
Ordinary Rewriting Templates, A%B@C or A@B	9-7
Repeated Rewrites Template, A%B	9-7
Specified Route Rewriting Templates, A (B)C@D or A (B)C	9-8
Case Sensitivity in Rewrite Rule Templates	9-8
How the MTA Applies Rewrite Rules to an Address	9-8
Step 1. Extract the First Host or Domain Specification	9-9
Step 2. Scan the Rewrite Rules	9-10
Step 3. Rewrite Address According to Template	9-11
Step 4. Finish the Rewrite Process	9-12
Rewrite Rule Failure	9-12
Syntax Checks After Rewrite	9-12
Handling Domain Literals	9-12
Template Substitutions and Rewrite Rule Control Sequences	9-13
Username and Subaddress Substitution, \$U, \$0U, \$1U	9-15
Host/Domain and IP Literal Substitutions, \$D, \$H, \$nD, \$nH, \$L	9-16
Literal Character Substitutions, \$\$, \$%, \$@	9-16
LDAP Query URL Substitutions, \$]..[.....	9-16
General Database Substitutions, \$(...)	9-17
Apply Specified Mapping, \${...}	9-18
Customer-supplied Routine Substitutions, \$[...]	9-18
Single Field Substitutions, \$&, \$!, \$*, \$#	9-19

Unique String Substitutions.....	9-20
Source-Channel-Specific Rewrite Rules (\$M, \$N).....	9-20
Destination-Channel-Specific Rewrite Rules (\$C, \$Q).....	9-20
Direction-and-Location-Specific Rewrite Rules (\$B, \$E, \$F, \$R).....	9-21
Host-Location-Specific Rewrites (\$A, \$P, \$S, \$X).....	9-21
Changing the Current Tag Value, \$T.....	9-22
Controlling Error Messages Associated with Rewriting (\$?).....	9-22
Handling Large Numbers of Rewrite Rules	9-23
Testing Rewrite Rules	9-23
Rewrite Rules Example	9-23

10 Using Predefined Channels

Predefined Channels	10-1
To Deliver Messages to Programs Using the Pipe Channel	10-2
To Configure the Native (/var/mail) Channel	10-3
To Temporarily Hold Messages Using the Hold Channel	10-4
The Conversion Channel	10-4
MIME Overview.....	10-5
Message Construction.....	10-5
MIME Headers.....	10-5
Selecting Traffic for Conversion Processing.....	10-6
To Control Conversion Processing.....	10-7
Conversion Channel Information Flow.....	10-9
To Use Conversion Channel Environmental Variables.....	10-9
To Use Conversion Channel Output Options.....	10-12
Headers in an Enclosing MESSAGE/RFC822 Part.....	10-14
To Call Out to a Mapping Table from a Conversion Entry.....	10-14
To Bounce, Delete, Hold, Retry Messages Using the Conversion Channel Output.....	10-15
To Bounce Messages.....	10-16
To Conditionally Delete a Message or Its Parts.....	10-16
To Hold a Message.....	10-16
To Cause Messages to Be Reprocessed.....	10-16
Conversion Channel Example.....	10-17
Automatic Arabic Character Set Detection.....	10-20
To Automatically Detect Arabic Character Sets.....	10-20
Character Set Conversion and Message Reformatting	10-21
Character Set Conversion.....	10-22
Message Reformatting.....	10-24
Service Conversions.....	10-28

11 LMTP Delivery

Overview of LMTP	11-1
LMTP Delivery Features	11-1
Support for LMTP Client and Server to Detect and Respond to Certain Conditions.....	11-2
Messaging Processing in a Two-Tiered Deployment Without LMTP	11-2
Messaging Processing in a Two-Tiered Deployment With LMTP	11-3
LMTP Architecture	11-4

Configuring LMTP	11-4
Before You Begin.....	11-5
To Configure the Front-end MTA Relay with LMTP.....	11-5
To Configure Back-End Stores with LMTP and a Minimal MTA.....	11-8
LMTP Protocol as Implemented	11-11

12 Vacation Automatic Message Reply

Vacation Autoreply Overview	12-1
Configuring Autoreply	12-2
To Configure Autoreply on the Back-end Store System.....	12-2
To Configure Autoreply on a Relay.....	12-2
To Share Autoreply Information Between Relays.....	12-4
Vacation Autoreply Theory of Operation	12-4
Vacation Autoreply Attributes	12-5
Other Auto Reply Tasks and Issues	12-6
To Send Autoreply Messages for Email That Have Been Automatically Forwarded from Another Mail Server.....	12-6

13 Mail Filtering and Access Control

PART 1. MAPPING TABLES	13-1
Controlling Access with Mapping Tables	13-2
Access Control Mapping Tables - Operation.....	13-2
Access Control Mapping Table Flags	13-3
Input vs Output Flags.....	13-3
Output Flag Argument Order.....	13-4
Send Access and Mail Access Mapping Tables	13-5
SEND_ACCESS and ORIG_SEND_ACCESS Mapping Tables.....	13-8
MAIL_ACCESS and ORIG_MAIL_ACCESS Mapping Tables.....	13-9
FROM_ACCESS Mapping Table	13-11
PORT_ACCESS Mapping Table	13-16
IP_ACCESS Mapping Table	13-20
When Access Controls Are Applied	13-20
To Test Access Control Mappings	13-21
To Limit Specified IP Address Connections to the MTA	13-22
To Add SMTP Relaying	13-22
Allowing SMTP Relaying for External Sites.....	13-24
Configuring SMTP Relay Blocking	13-24
Differentiate Between Internal and External Mail.....	13-25
Differentiate Authenticated Users' Mail.....	13-26
To Add Distinguishing Authenticated Submissions.....	13-26
Prevent Mail Relay.....	13-27
To Use DNS Lookups Including RBL Checking for SMTP Relay Blocking.....	13-27
Handling Large Numbers of Access Entries	13-29
Controlling the Envelope From: Address in Mappings Strings and Mailing List Named Options and LDAP Attributes	13-31
PART 2. MAILBOX FILTERS	13-31

Sieve Filter Support	13-31
Sieve Filtering Overview	13-33
To Create User-level Filters	13-33
To Create Channel-level Filters	13-33
To Create a Channel-level Filter.....	13-34
Managing Sieve Scripts	13-35
Location of Managed Sieves	13-35
Script Name Storage	13-36
Stored Script Semantics.....	13-36
ManageSieve Service Configuration	13-37
Configuring the ManageSieve Service.....	13-37
Supported Channel Options.....	13-38
Supported ManageSieve-Channel-Specific Options.....	13-39
ManageSieve Server Use of PORT_ACCESS Mapping.....	13-40
Basic Command/Response Test Script.....	13-40
To Create MTA-Wide Filters	13-42
Routing Discarded Messages Out the FILTER_DISCARD Channel	13-42
To Debug User-level Filters	13-43
<i>imsimta test -exp</i> Output.....	13-44
<i>imsimta test -exp</i> Syntax.....	13-45

14 Using and Configuring MeterMaid for Access Control

Technical Overview	14-1
Theory of Operations	14-1
Options for MeterMaid	14-2
Limit Excessive IP Address Connections Using Metermaid – Example	14-7
Configuring check_metermaid.so Clients to Access Multiple MeterMaid Servers	14-9
Considerations for Distributing Load Across Multiple MeterMaid Servers.....	14-9
Configuring check_metermaid.so to Access Multiple MeterMaid Servers.....	14-9

15 Implementing Greylisting by Using MeterMaid

What is Greylisting and How Does It Work?	15-1
Basic Greylisting Implementation	15-2
Using Greylisting with Messaging Server.....	15-2
Enhancing Greylisting Functionality	15-3
Preloading the Greylisting Table with Outbound Transactions	15-4
Matching a Range of IP Addresses.....	15-4
Simplifying the Sender Address	15-5
Providing an Opt-In Mechanism	15-5
Whitelisting Based on User's Addressbook	15-6
Combining Functionality: A Complex Example	15-7
Mapping Table Notes	15-10

16 MeterMaid Reference

<i>configutil</i> Options	16-1
Table Types	16-1

greylisting Tables	16-1
simple Tables	16-1
throttle Tables	16-1
check_metermaid.so Reference	16-2
adjust Routine	16-3
adjust_and_test Routine	16-3
fetch Routine	16-4
greylisting Routine.....	16-5
remove Routine	16-6
store Routine	16-6
test Routine.....	16-7
throttle Routine.....	16-8
17 Message Tracking and Recall	
18 Administering Event Notification Service	
ENS Publisher in Messaging Server.....	18-1
Configuring the ENS Publisher in Unified Configuration	18-1
Administering Event Notification Service	18-2
Starting and Stopping ENS.....	18-2
Event Notification Service Configuration Options	18-2
ENS SSL Support.....	18-2
ENS Support for Password Based Authentication.....	18-3
19 Messaging Server Specific Event Notification Service Information	
Event Notification Types and Options	19-1
Event Types.....	19-1
Options	19-2
Mandatory Event Reference Options	19-3
Optional Event Reference Options	19-3
Available Options for Each Event Type.....	19-4
Payload	19-5
Payload Configuration Options	19-5
Examples	19-5
Implementation Notes.....	19-6
20 Event Notification Service C API Reference	
API Overview.....	20-1
API Basic Usage	20-2
New Client API ens_sopen.....	20-3
API Usage Notes	20-3
21 Configuring IMAP IDLE	
Benefits of Using IMAP IDLE	21-1
Configuring IMAP IDLE with ENS in Unified Configuration.....	21-1

Prerequisites for Configuring IMAP IDLE with ENS	21-2
To Configure IMAP IDLE with ENS	21-2
To Disable IMAP IDLE	21-2

22 Lemonade Profile 1 Support

Introduction to Lemonade	22-1
Lemonade Features	22-2
Support for BURL	22-2
IMAP URLAUTH Support	22-3
IMAP CATENATE Support	22-3
IMAP Conditional Store Operation Support	22-3
IMAP ANNOTATE Support	22-4
Controlling IMAP CAPABILITIES Vector	22-4
Support for SMTP Submission Service Extension for Future Message Release	22-4

23 Managing Logging

Overview of Logging	23-1
What Is Logging and How Do You Use it?	23-1
Types of Logging Data	23-2
Types of Messaging Server Log Files	23-2
Tools for Managing Logging	23-4
Tracking a Message Across the Various Log Files	23-4
Managing MTA Message and Connection Logs	23-5
Understanding the MTA Log Entry Format	23-6
Enabling MTA Logging.....	23-9
Specifying Additional MTA Logging Options	23-10
MTA Message Logging Examples	23-12
Enabling Dispatcher Debugging	23-23
Managing Message Store, Admin, and Default Service Logs	23-24
msconfig Logging Options	23-25
Understanding Service Log Characteristics	23-26
Understanding Service Log File Format	23-28
Defining and Setting Service Logging Options	23-29
Searching and Viewing Service Logs	23-31
Working With Service Logs	23-32
Implementing and Configuring Message Store Transaction Logging	23-33
Overview of Message Store Transaction Logging	23-33
Message Store Transaction Logging Log Entries	23-34
Configuring Message Store Transaction Logging	23-35
Message Store Transaction Log Examples	23-36
Other Message Store Logging Features	23-39
Message Store Logging Examples	23-39
Using Message Store Log Messages	23-40
MMP Logging	23-40

24	Monitoring Messaging Server	
	Automatic Monitoring and Restart	24-1
	Daily Monitoring Tasks	24-1
	Checking Postmaster Mail	24-2
	Monitoring and Maintaining the Log Files	24-2
	Setting Up the msprobe Utility	24-2
	Utilities and Tools for Monitoring	24-2
25	Monitoring the MTA	
	Monitoring the Size of the Message Queues	25-1
	Symptoms of Message Queue Problems	25-1
	To Monitor the Size of the Message Queues	25-1
	Checking for Held messages	25-2
	Monitoring Rate of Delivery Failure	25-2
	Symptoms of Rate of Delivery	25-2
	To Monitor the Rate of Delivery Failure	25-2
	Monitoring Inbound SMTP Connections	25-2
	Symptoms of Unauthorized SMTP Connections	25-2
	To Monitor Inbound SMTP Connections	25-2
	Monitoring the Dispatcher and Job Controller Processes	25-3
	Symptoms of Dispatcher and Job Controller Processes Down	25-3
	To Monitor Dispatcher and Job Controller Processes	25-3
26	SNMP Support	
	SNMP Implementation	26-1
	SNMP Operation in the Messaging Server	26-2
	Configuring SNMP Support for Oracle Solaris 10	26-3
	Net-SNMP Configuration	26-3
	Messaging Server Subagent Configuration	26-4
	Running as a Standalone SNMP Agent	26-5
	Monitoring Multiple Instances of Messaging Server	26-6
	Using Standalone Agents for High-availability Failover	26-6
	Distinguishing Multiple Instances Through SNMP v3 Context Names	26-6
	Messaging Server's Net-SNMP-based SNMP Subagent Options	26-7
	Monitoring from an SNMP Client	26-8
	SNMP Information from the Messaging Server	26-8
	applTable	26-9
	assocTable	26-10
	mtaTable	26-11
	mtaGroupTable	26-11
	mtaGroupAssociationTable	26-13
	mtaGroupErrorTable	26-13
27	Short Message Service (SMS)	
	Introduction	27-1

One-Way SMS.....	27-2
Two-Way SMS.....	27-2
Requirements.....	27-3
SMS Channel Theory of Operation.....	27-3
Directing Email to the Channel.....	27-3
The Email to SMS Conversion Process.....	27-4
Sample Email Message Processing.....	27-7
The SMS Message Submission Process.....	27-7
Site-defined Address Validity Checks and Translations.....	27-10
Site-defined Text Conversions.....	27-11
Message Header Entries.....	27-11
Message Body Entries.....	27-12
Example SMS Mapping Table.....	27-12
SMS Channel Configuration.....	27-14
Adding an SMS Channel.....	27-14
Adding the Channel Definition and Rewrite Rules.....	27-15
To Add Channel Definition and Rewrite Rules.....	27-15
Controlling the Number of Simultaneous Connections.....	27-16
Setting SMS Channel Options.....	27-16
Available Options.....	27-17
Email to SMS Conversion Options.....	27-19
SMS Gateway Server Option.....	27-22
SMS Options.....	27-22
SMPP Options.....	27-27
Localization Options.....	27-29
Formatting Templates.....	27-32
Adding Additional SMS Channels.....	27-33
Adjusting the Frequency of Delivery Retries.....	27-34
Sample One-Way Configuration (MobileWay).....	27-34
Debugging.....	27-35
Configuring the SMS Channel for Two-Way SMS.....	27-35
SMS Gateway Server Theory of Operation.....	27-36
Function of the SMS Gateway Server.....	27-36
Behavior of the SMPP Relay and Server.....	27-37
Remote SMPP to Gateway SMPP Communication.....	27-37
SMS Reply and Notification Handling.....	27-38
Routing Process for SMS Replies.....	27-39
SMS Gateway Server Configuration.....	27-39
Setting Up Bidirectional SMS Routing.....	27-40
Set the SMS Address Prefix.....	27-40
Set the Gateway Profile.....	27-40
Configure the SMSC.....	27-41
Enabling and Disabling the SMS Gateway Server.....	27-41
Starting and Stopping the SMS Gateway Server.....	27-41
SMS Gateway Server Configuration File.....	27-41
Configuring Email-To-Mobile on the Gateway Server.....	27-42
A Gateway Profile.....	27-42

An SMPP Relay	27-42
An SMPP Server	27-43
Configuring Mobile-to-Email Operation	27-43
Configure a Mobile-to-Email Gateway Profile	27-43
Configure a Mobile-Email SMPP Server	27-44
Configuration Options	27-44
Global Options.....	27-45
Thread Tuning Options.....	27-45
Historical Data Tuning.....	27-46
Miscellaneous	27-47
SMPP Relay Options.....	27-47
SMPP Server Options	27-49
Gateway Profile Options.....	27-50
Configuration Example for Two-Way SMS.....	27-54
SMS Gateway Server Storage Requirements	27-56
SMS Configuration Examples	27-58
28 Configuring Messaging Server for One-Way SMS	
29 Configuring Messaging Server for Two-Way SMS	
30 Using the iSchedule Channel to Handle iMIP Messages	
Inviting Users on Internal and External Calendar Systems Background	30-1
Manually Accepting External Invitations.....	30-1
Automatically Accepting External Invitations	30-2
Message Server iMIP Configuration Overview	30-2
Configuring the iSchedule Channel for iMIP Messages in Unified Configuration	30-3
Using the iSchedule Recipe to Automate Configuring the iSchedule Channel in Unified Configuration 30-3	
Manually Configuring the iSchedule Channel in Unified Configuration	30-4
Verifying the Calendar Server Configuration.....	30-5
Modifying iSchedule Channel Options.....	30-6
To Enable or Disable iMIP Message Processing.....	30-6
To Modify the iSchedule Service URL	30-6
Configuring the iSchedule Channel in Legacy Configuration	30-6
Troubleshooting the iSchedule Configuration	30-7
31 Handling sendmail Clients	
To Create the sendmail Configuration File on Oracle Solaris 8 Platforms.....	31-1
To Create the sendmail Configuration File on Oracle Solaris 9 Platforms.....	31-2
32 Handling Forged Email by Using the Sender Policy Framework	
Theory of Operations.....	32-1
SPF Limitations.....	32-3
SPF Pre-Deployment Considerations	32-3
Setting up the Technology.....	32-3

Reference Information	32-3
Testing SPF by Using spfquery	32-5
Syntax.....	32-5
Example with Debugging Enabled.....	32-6
Handling Forwarded Mail in SPF by Using the Sender Rewriting Scheme (SRS).....	32-7
33 Message Store Directory Layout	
Message Store Directory Layout.....	33-1
34 Shared Folders Overview	
35 Monitoring LDAP Directory Server	
Symptoms of slapd Problems	35-1
To Monitor slapd	35-1
36 Monitoring System Performance	
Monitoring End-to-end Message Delivery Times	36-1
Monitoring CPU Usage	36-1
37 Monitoring the Message Store	
General Message Store Monitoring Procedures.....	37-1
Checking Hardware Space.....	37-2
Checking Log Files	37-2
Checking User IMAP/POP/Webmail Session by Using Telemetry	37-2
Checking stored Processes	37-3
Checking Database Log Files.....	37-4
Checking User Folders	37-4
Checking for Core Files	37-4
Monitoring <i>imapd</i> , <i>popd</i> and <i>httpd</i>	37-4
Symptoms of <i>imapd</i> , <i>popd</i> and <i>httpd</i> Problems	37-4
To Monitor <i>imapd</i> , <i>popd</i> and <i>httpd</i>	37-4
Monitoring <i>stored</i>	37-5
Symptoms of <i>stored</i> Problems	37-5
To Monitor <i>stored</i>	37-5
Monitoring the State of Message Store Database Locks	37-6
Symptoms of Message Store Database Lock Problems	37-6
To Monitor Message Store Database Locks.....	37-6
To Monitor Mailbox Quotas and Usage	37-6
To Monitor Message Store Database Statistics with <i>imcheck</i>	37-7
Gathering Message Store Counter Statistics by Using <i>counterutil</i>	37-7
To Get a Current List of Available Counter Objects	37-7
<i>counterutil</i> Output	37-8
Gathering Alarm Statistics by Using <i>counterutil</i>	37-8
IMAP, POP, HTTP, and MMP Connection Statistics by Using <i>counterutil</i>	37-9
Disk Usage Statistics by Using <i>counterutil</i>	37-10
Server Response Statistics	37-10

38 Monitoring User Access to the Message Store

39 Message Archiving

Microsoft Exchange Envelope Journaling	39-1
Archiving Overview	39-2
Message Archiving Systems: Compliance and Operational.....	39-2

40 Unified Messaging

Using Messaging Server to Manage Unified Messaging	40-1
What is the Challenge?	40-1
The Oracle Solution	40-1
Open Standards and Regulatory Requirements.....	40-2
Architectural Overview of a Unified Messaging Application.....	40-2
Message Deposit.....	40-2
Message Retrieval via Telephone User Interface	40-4
Message Retrieval via PC.....	40-6
Message Retrieval Through an IMAP Client	40-6
Message Retrieval Through Convergence.....	40-7
Designing and Coding Your Unified Messaging Application.....	40-9
Planning the Message-Type Configuration.....	40-9
Coding and Configuring Your UM System	40-9
Mailbox Administration and Operations	40-11
Sample IMAP Sessions Using Message-Type Flags.....	40-12
Administering Quotas for Message Types.....	40-13
Expiring Messages by Message Type.....	40-14
Delivering Notifications for Message Types	40-16
Notifications for Particular Message States.....	40-17
How Do You Implement Notifications for Message Types?	40-17
Notification Properties for Message Types	40-19
Additional Unified Messaging Support Features	40-21
Set IMAP Flag Based on Header Value at Delivery.....	40-21
Modifications to IMAP Commands to Provide Message Counts.....	40-21
IMAP Unauthenticate.....	40-21
Modify IMAP APPEND to bypass quotas	40-21
SMTP Future Release.....	40-22

41 Messaging Server Command-Line Reference

configure Command	41-1
Syntax.....	41-1
Options	41-1
Examples	41-2
configtoxml Command.....	41-2
Syntax.....	41-2
Options	41-2
Example	41-3
Notes on the configtoxml Command.....	41-4

msconfig Command	41-4
Syntax.....	41-4
Options	41-5
Notes on the msconfig Command	41-7
Option Name Changes in Unified Configuration	41-7
Using the msconfig Command in Edit Mode	41-8

Part II Improving Performance

42 Messaging Server Tuning and Best Practices

43 Tuning the mboxlist Database Cache in Unified Configuration

Setting the Mailbox Database Cache Size	43-1
To Adjust the Mailbox Database Cache Size.....	43-3
To Monitor the Mailbox Database Cache Size	43-3

44 Best Practices for Messaging Server and ZFS

Before You Begin	44-1
Configuration Recommendations for ZFS and Messaging Server	44-1
mboxlist Database, Message File and Index Cache Files Overview	44-1
Index Cache Record File System	44-2
Access Time Record	44-2
ZFS Pool Space Utilization.....	44-2
To Configure ZFS and Messaging Server	44-2
ZFS Administration Recommendations	44-3
ZFS Reference	44-3

Part III Troubleshooting

45 Troubleshooting the MTA

Troubleshooting Overview	45-1
Standard MTA Troubleshooting Procedures	45-1
Check the MTA Configuration.....	45-2
Check the Message Queue Directories.....	45-2
Check the Ownership of Critical Files.....	45-2
Check that the Job Controller and Dispatcher Are Running	45-3
Check the Log Files	45-4
Running a Channel Program Manually	45-5
Starting and Stopping Individual Channels	45-5
To Stop Outbound Processing (dequeueing) for a Specific Channel	45-5
To Stop Inbound Processing from a Specific Domain or IP Address (Enqueueing to a Channel)	45-6
An MTA Troubleshooting Example	45-6
Identify the Channels in the Message Path	45-7
Manually Start and Stop Channels to Gather Data.....	45-7
Common MTA Problems and Solutions	45-10

TLS Problems	45-10
Changes to Configuration Files or MTA Databases Do Not Take Effect.....	45-11
The MTA Sends Outgoing Mail but Does Not Receive Incoming Mail.....	45-11
Dispatcher (SMTP Server) Won't Start Up	45-11
Timeouts on Incoming SMTP Connections.....	45-11
To Identify the Causes of Timeouts on Incoming SMTP Connections	45-11
Messages Are Not Dequeued	45-12
Creating a New Channel.....	45-14
MTA Messages Are Not Delivered.....	45-15
Messages are Looping	45-16
Diagnosing and Cleaning up .HELD Messages	45-17
Received Message is Encoded	45-19
Server-Side Rules (SSR) Are Not Working.....	45-19
Testing Your SSR Rules.....	45-20
Common Syntax Problems	45-20
Slow Response After Users Press Send Email Button	45-20
Asterisks in the Local Parts of Addresses or Received Fields	45-20
Abnormal Job Controller Terminations Seen in <code>job_controller</code> Logs.....	45-20
General Error Messages	45-21
Errors in <code>mm_init</code>	45-21
Compiled Configuration Version Mismatch.....	45-23
Swap Space Errors.....	45-23
File Open or Create Errors	45-24
Illegal Host/Domain Errors	45-24
Errors in SMTP channels, <code>os_smtp_*</code> errors.....	45-25

46 Troubleshooting the Message Store

Repairing Mailboxes and the Mailboxes Database (<code>reconstruct</code> Command)	46-1
Reduced Message Store Performance	46-1
Convergence Not Loading Mail Page.....	46-2
Command Using Wildcard Pattern Does Not Work	46-2
Unknown/invalid Partition	46-3
User Mailbox Directory Problems	46-3
Store Daemon Not Starting	46-4
User Mail Not Delivered Due to Mailbox Overflow	46-4
IMAP Events Become Slow	46-4

Part IV Java Message Queue Notifications

47 Java Message Queue Notification

Two Notification Messaging Services.....	47-1
Notification Plug-ins	47-2
Benefits of Using JMQ Notification	47-2
Publishing to a Topic or a Queue.....	47-3
Using Multiple JMQ Notification Plug-ins.....	47-3
Configuring Options for a Notification Plug-in	47-4

Configuring Conditional Notifications for Specified Users.....	47-4
Sending Conditional Notifications to Distributed Message Queue Destinations	47-4
JMQ Notification	47-5
Setting JMQ Notifications in Unified Configuration	47-5
Where to Go for More Information	47-5

48 Java Message Queue Notification Messages and Properties

Notification Messages	48-1
Rules and Guidelines for Notification Messages	48-2
Notifications for Particular Message Types	48-3
Setting Different Options on a per-messagetype Basis.....	48-4
Default Values of the <i>notifytarget</i> Options	48-4
Notification Message Properties	48-5
Standard Notification Message Properties	48-5
Properties Specific to Particular Notification Messages	48-6
Syntax for newflags and oldflags Properties	48-8
Properties Carried with Each Notification Message	48-8

49 Configuring a JMQ Notification Service (Tasks and Examples)

Planning for Your JMQ Notification Service.....	49-1
To Configure an Instance of the JMQ Notification Plug-in	49-2
Before You Begin	49-2
Specifying Notification Messages that Use More Than One <i>msconfig</i> Option	49-5
To Configure New-Message and Updated-Message Notifications with Message Headers and Message Bodies.....	49-5
To Configure Deleted-Message Notifications with Message Headers	49-5
Configuring Notifications for Changes in Message Status	49-6
To Enable Notifications When Message-Flags Have Changed	49-6
To Enable Conditional Notifications for Specified Users.....	49-7
To Configure Conditional Notifications to Be Sent to Different Message Queue Destinations.....	49-7

50 Enabling JMQ Event Notification (Example)

Event Notifications in Messaging Server Overview	50-1
Enable Java Message Queue (JMQ).....	50-1
Configure and Enable the Messaging Server <i>jmnotify</i> Plugin.....	50-2
Verify the JMQ Broker	50-2
Related Information.....	50-4

Part V Managing the Message Store and Mailboxes

51 Managing Mailboxes

To Manage Mailboxes with <i>mboxutil</i>	51-1
Examples	51-2
To Move Mailboxes to a Different Disk Partition.....	51-2
To Remove Orphan Accounts	51-3

To Find a Mailbox's Directory Using hashdir	51-3
To Find Out How Many Users Have Read Messages in a Shared Folder	51-3

52 Backing Up and Restoring the Message Store

Mailbox Backup and Restore Overview	52-1
To Create a Mailbox Backup Policy	52-2
Peak Business Loads	52-2
Full and Incremental Backups	52-2
Parallel or Serial Backups	52-2
To Create Backup Groups	52-3
Pre-defined Backup Group	52-4
To Run the imbackup Utility	52-4
Running the imbackup Utility	52-4
Incremental Backup	52-4
Excluding Bulk Mail When You Perform Backups	52-5
To Restore Mailboxes and Messages	52-5
Considerations for Partial Restore	52-5
To Restore Messages from a Mailbox that Has Been Incrementally Backed-up	52-7
To Use StorageTek Enterprise Backup Software	52-8
To Back Up Data By Using StorageTek Enterprise Backup Software	52-8
Restoring Data Using StorageTek Enterprise Backup Software	52-9
To Use a Third Party Backup Software (Besides StorageTek Enterprise Backup Software)	52-10
Troubleshooting Backup and Restore Problems	52-10
Message Store Disaster Backup and Recovery	52-11

53 Administering Very Large Mailboxes

Very Large Mailboxes Overview	53-1
The Structure of a Mailbox	53-1
Mailbox Size Limit	53-2
Mailbox Migration	53-2
Pre-Deployment Preparations	53-2
Checking Mailbox Data	53-3

54 Message Store Message Expiration

imexpire Theory of Operation	54-1
To Deploy the Message Expiration Feature	54-2
To Define Message Expiration Policy	54-2
Examples of Message Expiration Policy	54-2
To Set Rules Implementing Message Expiration Policy	54-3
Expiration Rules Guidelines	54-3
Localized Mailbox Names in imexpire	54-5
Setting imexpire Rules Textually	54-6
Example imexpire Rules	54-7
Setting imexpire Folder Patterns	54-7

55	Configuring Message Expiration (Tasks)	
	To Set imexpire Rules Textually	55-1
	To Set Expiration Rules by Using the msconfig Command.....	55-2
	To Set imexpire Folder Patterns	55-3
	To Schedule Message Expiration and Logging Level	55-3
	Expire and Purge Log and Scheduling Options	55-3
	To Set imexpire Logging Levels.....	55-4
	To Exclude Specified Users from Message Expiration	55-5
56	Configuring POP, IMAP, and HTTP Services	
	General Configuration	56-1
	Enabling and Disabling Services.....	56-2
	Specifying Port Numbers	56-2
	Ports for Encrypted Communications.....	56-2
	IMAP Over SSL	56-2
	POP Over SSL	56-3
	HTTP Over SSL	56-3
	Service Banner	56-3
	Login Requirements	56-3
	To Set the Separator for POP Clients.....	56-3
	To Allow Log In without Using the Domain Name	56-4
	Password-Based Login	56-4
	Certificate-Based Login	56-4
	Performance Options	56-5
	Number of Processes	56-5
	Number of Connections per Process.....	56-6
	Number of Threads per Process.....	56-6
	Dropping Idle Connections	56-7
	Logging Out HTTP Clients.....	56-7
	Client Access Controls	56-7
	To Configure POP Services	56-8
	To Configure IMAP Services	56-9
	Configuring IMAP IDLE	56-10
	To Configure the mshttpd Process for Use by Convergence	56-10
	Configuring Your HTTP Service.....	56-11
57	Handling Message Store Overload	
	Overview of Managing Message Store Load	57-1
	Message Store Load Throttling.....	57-1
	Job Controller Stress Handling	57-1
	Default Job Controller Configuration.....	57-2
58	Managing Message Store Partitions and Adding Storage	
	Message Store Partition Overview.....	58-1
	To Add a Message Store Partition.....	58-2
	To Change the Default Message Store Partition.....	58-2

Adding More Physical Disks to the Message Store	58-3
59 Managing Message Store Quotas	
Message Store Quota Overview	59-1
Quota Overview	59-1
Quota Theory of Operations.....	59-2
Message Store Quota Attributes and Options	59-3
To Specify a Default User Quota	59-5
To Specify Individual User Quotas	59-6
To Specify Domain Quotas	59-6
To Set Up Quota Notification	59-6
To Disable Quota Notification.....	59-7
To Enable or Disable Quota Enforcement	59-8
To Enable Quota Enforcement at the User level.....	59-8
To Perform Quota Enforcement at the Domain Level.....	59-8
Disabling Quota Enforcement	59-8
To Set a Grace Period	59-9
Netscape Messaging Server Quota Compatibility Mode	59-9
60 Managing Message Types in the Message Store	
To Configure Message Types	60-1
Sending Notification Messages for Message Types	60-2
Administering Quotas by Message Type	60-3
Before You Set Message-Type Quotas	60-3
Methods of Setting Message-Type Quotas	60-3
Example of a Message-Type Quota Root.....	60-3
Expiring Messages by Message Type	60-4
Example: Sample Rules for Expiring Different Message Types.....	60-5
61 Managing Shared Folders	
Specifying Sharing Attributes for Private Shared Folders	61-1
To Create a Public Shared Folder	61-2
To Grant Folder Access Rights Based on Group Membership	61-3
To Set or Change a Shared Folder's Access Control Rights	61-3
Shared Folder Examples.....	61-4
Enabling or Disabling Listing of Shared Folders	61-5
Setting Up Distributed Shared Folders	61-5
Setting Up Distributed Shared Folders-Example	61-5
Monitoring and Maintaining Shared Folder Data	61-6
To Monitor Shared Folder Usage.....	61-6
To List Users and Their Shared Folders.....	61-7
To Remove Inactive Users	61-7
To Set Access Rights	61-7

62 Upgrading the Message Store

Architecture and Components	62-1
Message Store Component Version Compatibilities	62-3
Upgrading the Mailboxes	62-4
Upgrading and Downgrading the Berkeley Database (BDB).....	62-4
Database BTREE File	62-5
Database Log Files	62-5
IMAPD, MSHTTPD and Convergence	62-5
Upgrading From Messaging Server 32-bit to 64-bit.....	62-6
Migrating from x86 to SPARC	62-6
stored -r.....	62-6
ims_db_upgrade	62-6
Downgrading.....	62-6
Significant Changes in the Message Store Between Versions	62-6
Changes from Messaging Server 6.3 to Messaging Server 7.0	62-7
Changes to <i>store.idx</i>	62-7
Message Store Maintenance Queue and <i>impurge</i>	62-7
Mailbox Self-Healing (Auto-Repair)	62-8
Changes from Messaging Server 7 to Messaging Server 7 Update 1	62-8
Berkeley Database Upgrade	62-8
Changes from Messaging Server 7 Update 1 to Messaging Server 7 Update 5	62-8
Changes to the Owner's Seen and Deleted Flags	62-8
Immediate flag update and state sharing.....	62-9
Change to the <i>service.imap.capability.condstore</i> option	62-9
Changes to the Berkeley Database.....	62-9
Changes to <i>mboxlist</i> and <i>lockdir</i> BDB environments.....	62-9

63 Message Store Automatic Recovery On Startup

Overview of Automatic Recovery on Startup	63-1
Automatic Startup and Recovery Theory of Operations	63-1
Error Messages Signifying that reconstruct is Needed.....	63-2
Message Store Database Snapshot Theory of Operations	63-2
Message Store Database Snapshot Interval and Location.....	63-3
Message Store Database Snapshot Options	63-3

64 Message Store Maintenance Queue

Message Store Maintenance Queue Overview.....	64-1
Displaying the Maintenance Queue.....	64-2
Deleting, Expunging, Purging, and Cleaning Up Messages.....	64-2
Mailbox Self Healing (Auto Repair).....	64-3
Maintenance Queue Configuration Options	64-3
The impurge Command	64-3

65 Message Store Message Type Overview

About Message Type	65-1
Planning the Message-Type Configuration	65-1

Defining and Using Message Types.....	65-2
Message Types in IMAP Commands	65-2
66 Migrating Mailboxes to a New System	
Tools Summary for Relocating Messaging Server Users to a New Mailhost.....	66-1
Migrating Mailboxes from an x86 Host to a SPARC Host	66-2
Moving Mailboxes to Another Messaging Server While Online	66-2
Advantages	66-2
Disadvantages.....	66-2
Incremental Mailbox Migration While Online.....	66-2
Online Migration Overview	66-3
To Migrate User Mailboxes from One Messaging Server to Another While Online	66-4
To Move Mailboxes Using an IMAP Client.....	66-7
To Move Mailboxes by Using the <i>imsimport</i> Command	66-8
67 Monitoring Disk Space	
Overview	67-1
Symptoms of Insufficient Disk Space	67-1
Monitoring Disk Space	67-2
Monitoring the Message Store.....	67-2
Monitoring Message Store Partitions	67-2
68 Protecting Mailboxes from Deletion or Renaming	
69 Reducing Message Store Size Due to Duplicate Storage	
Relinker Theory of Operations.....	69-2
Using relinker in the Command Line Mode	69-2
Using Relinker in the Realtime Mode	69-3
Configuring Relinker	69-4
70 Specifying Administrator Access to the Message Store	
Overview.....	70-1
Adding an Administrator Entry	70-1
Modifying or Deleting an Administrator Entry	70-2
71 Valid Message Store UIDs and Folder Names	
Message Store User ID	71-1
Message Store Mailbox Name for Commands	71-1
Valid UIDs.....	71-1
72 Message Store Automatic Failover with Database Replication	
Overview of Message Store Database Replication.....	72-1
Configuration Options	72-3
Configuration Options	72-3

Command-line Utilities	72-4
Configuring Message Store Database Replication	72-5
To Configure a Three Node Cluster for HA.....	72-5
To Change the DB Replication Local Instance Port	72-5
Message Store Automatic Failover.....	72-6
Basic Requirements	72-6
Overview of Message Store Automatic Failover	72-6
Configuring Message Store Automatic Failover	72-7
To Configure the LMTP Server	72-7
To Configure the Client.....	72-8
73 Administering Message Store Database Snapshots (Backups)	
To Specify Message Store Database Snapshot Interval and Location.....	73-1
Message Store Database Snapshot Recovery and Verification	73-2
Message Store Database Snapshot Rolling Backup	73-2
Message Store Database Recovery	73-3
74 Messaging Server and Tiered Storage Overview	
Overview of Messaging Server Storage.....	74-1
Message Store and ZFS	74-2
How the Message Store Works	74-2
Messaging Server Disk Throughput.....	74-3
Messaging Server Disk Capacity	74-4
Disk Sizing for MTA Message Queues	74-4
MTA Message Queue Performance.....	74-4
MTA Message Queue Availability	74-5
MTA Message Queue Available Disk Sizing	74-5
Performance Considerations for a Message Store Architecture.....	74-5
Messaging Server Directories (General Recommendations for Storage).....	74-5
MTA Queue Directory.....	74-6
Messaging Server Log Directory	74-6
Mailbox Database Files.....	74-6
Message Store Index Files	74-6
Message Files	74-6
Mailbox List Database Temporary Directory	74-7
Multiple Store Partitions	74-8
Setting Disk Stripe Width	74-9
MTA Performance Considerations.....	74-9
MTA and RAID Trade-offs.....	74-10
Background: Communication Services Logical Architectures Overview	74-10
Two-tiered Logical Architecture	74-10
Benefits of a Two-tiered Architecture	74-11
Horizontal Scalability Strategy	74-12
Scaling Front-end and Back-end Services	74-12
Implementing Local Message Transfer Protocol (LMTP) for Messaging Server.....	74-13
Background: "How Email Works" Introduction to Messaging Server.....	74-13
What Does Messaging Server Enable Users to Do?	74-13

A User Decides to Send an Email	74-14
User Receives an Email	74-17
User Access Mailbox.....	74-18

75 Message Store Command Reference

configutil.....	75-1
counterutil.....	75-5
deliver	75-6
hashdir	75-8
imarchive.....	75-9
imcheck.....	75-10
imdbverify	75-13
imexpire.....	75-14
iminitquota	75-17
immonitor-access	75-17
impurge	75-22
imquotacheck	75-23
imsasm	75-29
imsbackup.....	75-32
imsconnutil	75-33
imscripter	75-35
imsexport.....	75-39
imsimport.....	75-40
imsrestore.....	75-42
mboxutil	75-44
mkbackupdir	75-50
msprobe.....	75-52
msuserpurge	75-53
readership	75-54
reconstruct.....	75-56
refresh	75-60
rehostuser.....	75-61
relinker	75-63
start-msg	75-65
stop-msg	75-66
stored.....	75-67

Preface

This guide explains how to administer Oracle Communications Messaging Server and its accompanying software components.

Audience

This document is intended for system administrators whose responsibility includes Messaging Server. This guide assumes you are familiar with the following topics:

- Messaging protocols
- Oracle Directory Server Enterprise Edition and LDAP
- System administration and networking
- General deployment architectures

Related Documents

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

- *Messaging Server Installation and Configuration Guide*: Provides instructions for installing and configuring Messaging Server.
- *Messaging Server Reference*: Provides additional information for using and configuring Messaging Server.
- *Messaging Server Release Notes*: Describes the new features, fixes, known issues, troubleshooting tips, and required third-party products and licensing.
- *Messaging Server Security Guide*: Provides guidelines and recommendations for setting up Messaging Server in a secure configuration.

Document Revision History

The following table lists the revision history for this guide:

Version	Date	Description
E63711-02	February 2016	Minor formatting and text changes.
E63711-01	September 2015	Initial release.

Documentation Accessibility

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

Access to Oracle Support

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

Part I

Monitoring and Managing Messaging Server

Part I of the *Messaging Server System Administrator's Guide* describes how to monitor and manage Oracle Communications Messaging Server.

Part I contains the following chapters:

- Messaging Server System Administration Overview
- Overview of Messaging Server Unified Configuration
- Configuring General Messaging Capabilities
- Configuring and Administering Multiplexor Services
- MTA Concepts
- MTA Address Translation and Routing
- About MTA Services
- Message Transfer Agent Command-line Utilities
- Configuring Rewrite Rules
- Using Predefined Channels
- LMTP Delivery
- Vacation Automatic Message Reply
- Mail Filtering and Access Control
- Using and Configuring MeterMaid for Access Control
- Implementing Greylisting by Using MeterMaid
- MeterMaid Reference
- Message Tracking and Recall
- Administering Event Notification Service
- Messaging Server Specific Event Notification Service Information
- Event Notification Service C API Reference
- Configuring IMAP IDLE
- Lemonade Profile 1 Support
- Managing Logging
- Monitoring Messaging Server
- Monitoring the MTA

- SNMP Support
- Short Message Service (SMS)
- Configuring Messaging Server for One-Way SMS
- Configuring Messaging Server for Two-Way SMS
- Using the iSchedule Channel to Handle iMIP Messages
- Handling sendmail Clients
- Handling Forged Email by Using the Sender Policy Framework
- Message Store Directory Layout
- Shared Folders Overview
- Managing Mailboxes
- Monitoring LDAP Directory Server
- Monitoring System Performance
- Monitoring the Message Store
- Monitoring User Access to the Message Store
- Message Archiving
- Unified Messaging
- Messaging Server Command-Line Reference

Messaging Server System Administration Overview

This chapter provides an overview of Oracle Communications Messaging Server, and describes the basic administration tasks and tools used to perform those tasks.

About Messaging Server

Oracle Communications Messaging Server is an extensible framework of cooperative modules that creates an enterprise-wide, open standards-based, scalable electronic message handling system. This system is the combination of message user and transfer agents, message stores, and access units that together provide electronic messaging.

Messaging Server provides a number of messaging capabilities, including:

- **Message Store.** Provides a modular and scalable repository of user messages.
- **Message Transfer Agent.** Responsible for routing, transfer, and delivery of internet mail messages. Oracle Communications Messaging Server includes a fast, scalable, and flexible MTA that replaces the Sendmail utility bundled with most UNIX systems.
- **Message Access.** Provides client access to messages over standard protocols IMAP and POP3. HTTP-based access is offered for web clients and Lemonade Profile 1 is supported for mobile devices.

Messaging Server Configuration Methods

When you install Messaging Server, you must decide if you want to use Unified Configuration or legacy configuration. Unified Configuration is an improved, streamlined process to configure and administer Messaging Server. Unlike in legacy configurations, Unified Configuration uses validation to verify configuration accuracy, and employs a single tool to configure the entire Messaging Server configuration (with a few exceptions). For more information, see "[Overview of Messaging Server Unified Configuration](#)."

Note: This guide describes how to administer Messaging Server by using Unified Configuration.

Overview of Messaging Server Administration Tasks

A Messaging Server administrator is responsible for the day-to-day tasks of maintaining and managing Messaging Server and its users. The tasks also include

managing Messaging Server components and potentially other Unified Communications Suite components.

You perform the following tasks as a Messaging Server administrator:

- Stopping and starting Messaging Server
- Managing user accounts
- Monitoring Messaging Server
- Tuning Messaging Server performance
- Migrating data to Messaging Server
- Backing up and restoring files
- Troubleshooting Messaging Server

About Messaging Server Administration Tools

Messaging Server provides the **msconfig** command-line utility for administering the server. For more information, see "[Messaging Server Command-Line Reference](#)." In addition, Messaging Server provides command-line utilities to manage the message store and the mail transfer agent (MTA). See "[Message Transfer Agent Command-line Utilities](#)" and "[Message Store Command Reference](#)" for more information.

Directory Placeholders Used in This Guide

Table 1–1 lists the placeholders that are used in this guide:

Table 1–1 Messaging Server Directory Placeholders

Placeholder	Directory
<i>MessagingServer_home</i>	Specifies the installation location for the Messaging Server software. The default is /opt/sun/comms/messaging64 .
<i>ConfigRoot</i>	Specifies the location of the configuration files. The default is /var/opt/sun/comms/messaging64/config .
<i>DelegatedAdmin_home</i>	Specifies the installation location for the Delegated Administrator software. The default is /opt/sun/comms/da .

Overview of Messaging Server Unified Configuration

This chapter introduces Oracle Communications Messaging Server Unified Configuration, describes its capabilities, and provides initial guidelines on how to transition from a legacy configuration. Unified Configuration provides the ability to configure Messaging Server in a way that is much less error-prone and much easier to script and reuse across multiple hosts.

What Is Messaging Server Unified Configuration?

Unified Configuration is an improved process to configure and administer Messaging Server. Unlike in legacy configurations, Unified Configuration uses validation to verify configuration accuracy, and employs a single tool to configure the entire Messaging Server configuration (with a few exceptions).

Table 2-1 describes how Unified Configuration improves upon issues with legacy configuration.

Table 2-1 Legacy Versus Unified Configuration

Legacy Configuration	Unified Configuration Improvement
Dealing with many configuration files (with inconsistent formats) and hand-editing them can lead to errors and invalid configurations.	Unified Configuration “unifies” configuration management. One tool, “ msconfig Command ,” administers Messaging Server configuration. Validation checking prevents introducing some configuration errors.
Configuration settings themselves are often complicated and not straight-forward.	Unified Configuration reduces redundancy and host specific-configurations, so that, for example, you can use the same settings for many options among the MMP, MTA, and message store configurations.
When problems arise, there are support challenges due to the many Messaging Server configuration files. Additionally, because passwords are contained in the configuration files, it makes it difficult for customers to just send these files to Oracle Support without first removing the passwords.	Unified Configuration uses only three text files to store configuration data, with most data stored in the config.xml file. Passwords are stored in a separate file, removing the need for customers to edit configuration files before sending to Oracle Support. In addition, Unified Configuration provides an audit trail of configuration changes. The changes (currently the last 100 changes) are actually stored in a repository, referred to as a <i>graveyard</i> . Storing of changes further enables you to restore an entire configuration, and to roll-back and roll-forward between configurations. See “ Unified Configuration Files ” for more information.

Table 2–1 (Cont.) Legacy Versus Unified Configuration

Legacy Configuration	Unified Configuration Improvement
<p>It is difficult to separate instance-specific details from details shared by functionally similar machines.</p>	<p>Unified Configuration separates tasks for single instances (referred to as <i>instance</i>) of Messaging Server from global tasks for a group (referred to as <i>role</i>) of Messaging Server machines. The intention is that the role contain configuration information suitable for sharing with other hosts that have the same function in the deployment. At present, there is no mechanism to automatically share role configuration.</p> <p>See "Separating Roles and Instances" for more information.</p>
<p>Customers must create their own procedures and scripts with their own tools to manage a deployment.</p>	<p>New customers can write automation scripts by using the Unified Configuration <i>recipe language</i>. The recipe language introduces the ability to robustly change a configuration in a reproducible fashion in a way that can be sensitive to what was previously in the configuration. When you use recipes, you are able to make use of the Unified Configuration history and administrative undo features. In addition, Oracle can use the recipe language to automate configuration changes that are otherwise complex, interconnected, and require lots of documentation. The SpamAssassin.rcp, HAConfig.rcp, and LMTPSingleSystem.rcp recipes, available in the <i>MessagingServer_home/lib/recipes</i> directory, are good examples.</p> <p>See "Using Recipes" for more information.</p>

Unified Configuration Files

Table 2–2 describes the Unified Configuration file names, file management tool, file format, character set, XML schema, ownership, and file permissions. The Unified Configuration files are located in the *ConfigRoot* directory, by default, */var/opt/sun/comms/messaging64/config*.

Table 2–2 Unified Configuration File Properties

Configuration File	Description	File Management Tool	File Format	Char Set	File Ownership	Recommended File Permissions
restricted.cnf	Contains protected Messaging Server UID and GID information.	Text editor	option=value	ASCII	root	0644

Table 2–2 (Cont.) Unified Configuration File Properties

Configuration File	Description	File Management Tool	File Format	Char Set	File Ownership	Recommended File Permissions
xpass.xml	Contains obfuscated passwords (<i>BASE64</i> encoded). This is the only file within Unified Configuration where password information is stored.	msconfig utility	XML 1.0	UTF-8	mailsrv	0600
config.xml	Contains most of the non-password configuration information. In addition, when necessary, you could send this entire file to Oracle Support to help with resolving problems.	msconfig utility	XML 1.0	UTF-8	mailsrv	0640
configlib.xml	Contains static, default mapping tables that are mostly concerned with character sets and language issues, including: <ul style="list-style-type: none"> ▪ DISPOSITION_LANGUAGE ▪ DOMAIN_DC ▪ LANGUAGE_LOCALES ▪ LDAP_USERS_LANGUAGE ▪ LDAP_USERS2_LANGUAGE ▪ NOTIFICATION_LANGUAGE 	Managed by Oracle (that is, the file is not to be edited)	XML 1.0	UTF-8	mailsrv	0644

Notes:

- When you perform the initial Messaging Server configuration, the **restricted.cnf** file sets the UID under which to run Messaging Server. After initial configuration, there should rarely be a need to edit this file. A legacy configuration can also use the **restricted.cnf** file for enhanced security.
- Never edit the **configlib.xml** file. Doing so causes an unsupported configuration.

Note: About MTA Tailor Options

In legacy configuration, you use the MTA tailor file of option settings (**imta_tailor**) to set various MTA installation and operational parameters. In Unified Configuration, the MTA tailor file is obsolete and no longer used. Unified Configuration replaces the MTA tailor options that specified locations of MTA directories or files with rationalized, consistent locations, which are based off the installation main location and located by using the **SERVERROOT** environment variable. Legacy configuration MTA tailor options that set other sorts of MTA operational parameters have typically been replaced with Unified Configuration options of the form **mta.option-name**.

Enabling Unified Configuration in Messaging Server

There are two ways to enable Unified Configuration:

1. **When migrating to Unified Configuration:** Use the *MessagingServer_home/bin/configtoxml* program to migrate a legacy configuration to Unified Configuration. When you run **configtoxml**, your old configuration is converted to Unified Configuration.
 - The legacy configuration is saved in the *ConfigRoot/legacy-config* directory.
 - If necessary, you can use the **configtoxml -undo** command to restore a saved legacy configuration.
2. **For a new Messaging Server instance:** Run the **configure** command to enable Unified Configuration by default.
 - The presence of a **config.xml** file in the **config** directory indicates that Unified Configuration is enabled.
 - To generate a legacy configuration instead of a Unified Configuration, you can use the **configure -noxml** command.
 - When you perform a fresh installation of Messaging Server and choose to configure a Unified Configuration, you cannot revert that Unified Configuration to a legacy configuration. If, however, you upgrade Messaging Server and convert to a Unified Configuration (by running the **configtoxml** command), you can revert back to the legacy configuration.
 - A Unified Configuration is more simple than a legacy configuration. In addition, where appropriate, modern default values are established and seldom used features are removed (for example, the **tcp_tas** channel is not present in Unified Configuration).
 - Legacy configuration files such as **dispatcher.cnf**, **option.dat**, and so on, are ignored when Unified Configuration is enabled.

To Determine if Unified Configuration Is Deployed

The following example shows how to determine if Unified Configuration is deployed on your system:

```
cd /opt/sun/comms/messaging64/bin
imsimta version
...
Using /opt/sun/comms/messaging64/config/config.xml
SunOS host2.example.com 5.10 Generic_142901-03 i86pc i386 i86pc
```

In this example, the presence of the **config.xml** file indicates that Unified Configuration has been enabled on this host.

If you are using a compiled configuration and see in that the status is not compiled, you should recompile the configuration. For example:

```
/opt/sun/comms/messaging64/bin/imsimta version
...
Using /opt/sun/comms/messaging64/config/config.xml (not compiled)
SunOS host1.example.com 5.10 Generic_147441-09 i86pc i386 i86pc

/opt/sun/comms/messaging64/bin/imsimta cbuild
```

For more information, see "Compiling the MTA Configuration."

Understanding Unified Configuration Limitations

In general, Unified Configuration has consolidated all the various Messaging Server files. Nevertheless, the current Unified Configuration implementation has a few limitations:

- The channel sieves and channel header trimming option files have not yet been converted to XML.
- Some files, such as localized templates for DSNs and NDNs, might remain in their current format and not be converted to XML.
- The content of the **conversions** file is a mono-block in XML.

Using the Repository of Previous Configurations

The repository of previous configurations, known as the *graveyard*, is stored in the *ConfigRoot/old-configs/* directory. The move from current configuration to the graveyard is performed when a new configuration is written to disk. The graveyard maintains the most recent 100 configurations. With the graveyard, you can restore an old configuration by reverting to a previous configuration. Furthermore, you can compare differences between any two configurations, for example, between the active configuration and a previous configuration, or two old configurations.

To List Configurations

- Use the **msconfig history** command to show a list of configurations currently in the graveyard.

To Compare Configurations

- To compare configurations, use the **msconfig differences *m* *n*** command, where *m* and *n* are the numbers of the previous configurations from the **history** command that you want to compare.

Using Legacy Configuration Tools with Unified Configuration

Once Unified Configuration is enabled, legacy configuration tools might work differently than in previous releases. Specifically:

- Scripts that directly alter legacy configuration files do not work in a Unified Configuration.

- The **configutil** command can still set, get, and delete a legacy configuration. Additionally, in Unified Configuration, **configutil** options can also automatically perform:
 - Option name translations
 - Option value translations
 - Option value validations

Note: Use of the **configutil** command is deprecated in Unified Configuration mode. Some configuration changes that were previously possible with the **configutil** command are only possible by using the **msconfig** command, for example, some changes to notification configuration.

- The "mkbackupdir" command, which creates and synchronizes the backup directory with the information in the message store, works with Unified Configuration.
- The "imsimta program" command, which manipulates the program delivery options, does not work with Unified Configuration. Instead, you must use the "msconfig Command."
 - This command issues an error and exits with 1 when Unified Configuration is being used.
- All other utilities only consume options and continue to work with both Unified Configuration and legacy configurations.

Separating Roles and Instances

Unified Configuration separates tasks for single instances (referred to as *instance*) of Messaging Server from global tasks for a group (referred to as *role*) of Messaging Server machines. The intention is that the role contain configuration information suitable for sharing with other hosts that have the same role in the deployment.

Note: At present, there is no mechanism to automatically share role configuration.

Any configuration option can be an instance setting, a role setting, or both. When the same option is in both the role and instance, the instance value takes precedence. Both the **configure** and **msconfig** commands put a given setting in either the instance or the role based on the likely scope for the option. Normally, you use the default location (instance or role) determined by the **msconfig** command and not explicitly specify one or the other.

Initial configuration generates an instance and role, as does migrating from a legacy configuration to Unified Configuration.

More About Unified Configuration Options

This section provides information about password, restricted, and obsolete options.

- [Options That Have Passwords](#)
- [Restricted Options](#)

- [Obsolete Options](#)
- [Option Relationships](#)
- [Unified Configuration Option Names](#)

Options That Have Passwords

The password options have the following characteristics:

- Options can be marked as being a password.
- By default, password values are not displayed.
- The passwords are stored in obfuscated form in the `xpass.xml` file. Because Messaging Server stores passwords in a separate file, you do not need to edit configuration files to remove them before sharing with Oracle Support.

Restricted Options

Some configuration options are marked “restricted” by Oracle. Additionally, the XML schema exposes the entire configuration, so there are no longer any hidden configuration options.

Options may be restricted for several reasons, including:

- The option has complex and subtle consequences and would cause harm in all but a very few rare circumstances.
- The option might be a legacy option that should not be used in new systems.
- The option might be a placeholder for a feature that has not yet been implemented.

Restricted options have the following characteristics:

- The `msconfig` command displays a warning when you attempt to set a restricted option. In addition, the `msconfig` command requires an extra step to actually set the restricted option.
- The restriction is noted within the configuration file itself, which helps you to be aware of any special circumstances. For example:

```
<delimiter_char v="127" xannotation="RESTRICTED USAGE OPTION: user remark"
xauthor="dcn@example.com" xmtime="2010-05-12T17:42:19-08:00"/>
```

The **user remark** text is any optional remark added by the administrator when the configuration was updated. The “**RESTRICTED USAGE OPTION**” is text inserted by Oracle into the remark field when the option is restricted.

Caution: If you set a restricted option without being advised to do so by Oracle Support, your configuration is considered unsupported by Oracle.

Obsolete Options

When an option has been marked by Oracle as being obsolete, the configuration no longer uses it. However, you cannot remove it from the XML Schema as that would make existing configurations invalid.

When marked as obsolete, the option:

- Remains in the XML Schema
- Can no longer be set or changed
- Can only be deleted from a configuration

Option Relationships

Unified Configuration enables some relationships between options to be expressed so that when a particular option is set, other unnecessary options can be automatically removed. For example, if you set the **mx** option on a channel (for MX mail forwarding records), any of the **nomx**, **randommx**, and other related “mx” options are removed. In addition, Unified Configuration uses the concept of default relationships to help with configuration. For example, option X and option Y might have a default relationship such that when X is not set, the value is taken from Y; or when Y is not set, then Y’s default value is *value*. Furthermore, Unified Configuration has the capability to know in which release an option became available and warn when a certain configuration is not release-suitable. In general, option relationships help to reduce configuration mistakes.

Unified Configuration Option Names

Unified Configuration uses a “unified” option naming convention that is reminiscent of legacy **configutil** option names.

In general, this option naming convention uses the following structure:

[role.]group[.sub-group | .sub-group].option

[instance.]group[.sub-group | .sub-group].option

The following example shows a *group.sub-group.option* convention:

```
imap.logfile.flushinterval
```

In this example, **imap** is the group, **logfile** is the sub-group, and **flushinterval** is the option.

This example shows a *group.option* convention:

```
mta.mm_debug
```

In this example, **mta** is the group and **mm_debug** is the option.

Characteristics about option names to keep in mind:

- Many groups only appear once (for example, **imap** and **pop**).
- Some groups may appear many times. For example:

```
channel
mapping
sectoken
alias
task
```
- The group or sub-group can include a *:name* portion used for “named” groups. For example:

```
channel:tcp_local.slave_debug
partition:primary.path
```

Characteristics about instances and roles to keep in mind:

- An option in the “instance” overrides the same option in the “role.” For example, IMAP is effectively disabled by this configuration:

```
instance.imap.enable = 0
role.imap.enable = 1
```

- There actually is no option called **imap.enable**. It is either **role.imap.enable** or **instance.imap.enable**.
- When setting options, you typically do not specify either “role” or “instance.” The **msconfig** command applies heuristics to determine whether “role” or “instance” applies. Here is a sample, basic **config.xml** file that shows how the configuration uses instance and role:

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<xconfig...>
<role name="store">
<base>
<defaultdomain v="example.com"/> role.base.defaultdomain
</base>
<imap>
<enable v="1"/> role.imap.enable
<numprocesses v="2"/> role.imap.numprocesses
</imap>
<mapping name="ABC">
<rule pattern="x*y" template="$N"/> role.mapping:ABC.rule
<rule pattern="*" template="$Y"/> role.mapping:ABC.rule
</mapping>
</role>
<instance name="ims" roleref="store">
<base>
<hostname v="wassonite.example.com"/> instance.base.hostname
</base>
<mta>
<mm_debug v="5"/> instance.mta.mm_debug
</mta>
</instance>
</xconfig>
```

In the preceding example, some option names that you would see upon listing them with the **msconfig show** command are displayed in bold. Also, you can see that the default domain (**defaultdomain**), number of IMAP processes (**numprocesses**), and mappings (**mapping name**) have been defined for the store role; and that the host name (**hostname**) and MTA logging debug level (**mm_debug**) have been set for the store instance.

Tip: Use the **configutil -H** command to translate the legacy **configutil** option names to Unified Configuration names. For example:

```
configutil -H -o logfile.imap.expirytime
Configuration option: logfile.imap.expirytime
Unified Config Name: imap.logfile.expirytime
```

Example of Legacy Configuration and Unified Configuration

Unified Configuration greatly simplifies the configuration process, as shown in this example of configuring the SMTP server for debugging.

In a legacy configuration, you need to perform the following steps:

1. Edit the **imta.cnf** file and modify the **tcp_local** channel entry:

```
tcp_local identnonnumeric inner loopcheck maysaslserver maytlsserver mx \
pool SMTP_POOL remotehost sasls witchchannel tcp_auth smtp sourcespamfilter1 \
switchchannel master_debug
tcp_local-daemon
```

2. Edit the **option.dat** file and add the following option:

```
mm_debug=3
```

3. Edit (or create if it does not exist) the **tcp_local_option** file and add the following option:

```
trace_level=2
```

4. Check permissions on all files, especially the **tcp_local_option** file.

In Unified Configuration, the equivalent steps to the preceding task are the following:

```
% msconfig
msconfig> set channel:tcp_local.slave_debug
msconfig# set mm_debug 3
msconfig# set channel:tcp_local.options.trace_level 2
msconfig# write
```

Using Recipes

You use recipe files, which are expressed by using a programming language, to automate configuration tasks, typically by scripting them. Recipes are located in the *MessagingServer_home/lib/recipes* directory. The primary inspiration for the recipe language is the Icon programming language designed by Ralph Griswald. As such, it supports C-like expressions, operators, and assignments, Sieve-like conditionals, and loops. The available data types are integers, strings, and lists.

Recipes typically operate in three phases. First, a number of checks are done to make sure the right conditions exist for the recipe to be effective. Next the recipe asks a number of questions to determine exactly what changes should be made. Finally, the recipe implements the requested changes. Note that while this is the typical ordering, recipes are not constrained to use it and may use other approaches if appropriate.

By using the recipe language, you can more easily script complex configuration changes. For more information on writing recipes, see the help text for the recipe language by typing **msconfig -help** and choosing help for the **Recipe_language** topic, or refer to the discussion on recipe language in the *Messaging Server Reference*.

To Run a Recipe

To run a recipe, type the following command:

```
msconfig run <recipe_name>
```

Helpful Commands

This section provides some helpful commands to get started with Unified Configuration.

- [To Show Settings](#)
- [To Get Help](#)

To Show Settings

- Use the **msconfig show** command to display current settings. For example, to show all currently enabled options:

```
/opt/sun/comms/messaging64/bin/msconfig show *enable
role.watcher.enable = 1
role.schedule.enable = 1
role.store.enable = 1
role.store.purge.enable = 1
role.imap.enable = 1
role.pop.enable = 1
role.mta.enable = 1
role.dispatcher.service:SMTP.enable = 1
role.dispatcher.service:SMTP_SUBMIT.enable = 1
role.mmp.enable = 0
role.ens.enable = 1
role.http.enable = 1
```

To Get Help

- Use the **msconfig help** command.

Configuring General Messaging Capabilities

This chapter describes the general Oracle Communications Messaging Server tasks, such as starting and stopping services and configuring directory access by using command-line utilities. Tasks specific to administering individual Messaging Server services, such as POP, IMAP, HTTP, and SMTP, are described later in this guide.

Modifying Your Passwords

If you set up a number of administrators with the same password during the initial Messaging Server configuration, you might want to change the passwords of those administrators.

Table 3–1 shows the password options that are set up during initial runtime configuration. Use the **msconfig** command to make changes to the Messaging Server configuration, or **ldapmodify** to update information stored in Directory Server.

Table 3–1 Passwords Set in Messaging Server Initial Runtime Configuration

Option	Description
base.ugldapbindcred	Password for the Messaging Server LDAP user/group access account (base.ugldapbinddn). Use msconfig to change.
base.proxyadminpass	Password for the Proxy Administrator account (base.proxyadmin), which is used to provide proxy authentication access to end-user mailboxes. Use msconfig to change.
http.smtpauthpassword	Password used when mshttpd submits mail to the MTA. Set by initial configuration to the same password as base.proxyadminpass . Use msconfig to change.
SSL passwords for key files	Passwords that are stored in the xpass.xml file. Use the msconfig set -prompt "sectoken:Internal (Software) Token" command to change. This command causes msconfig to prompt for the password without an echo.
Admin Account credentials	The "admin" account is both in the service administrator group by default and is a store admin by default. You are prompted for this password during initial configuration. By default, the "admin" account is used for proxy and SMTP authentication, so this password needs to match the settings for base.proxyadminpass and http.smtpauthpassword .

Table 3-1 (Cont.) Passwords Set in Messaging Server Initial Runtime Configuration

Option	Description
Messaging End User Administrator	This is the LDAP user for this specific host. The base.ugldapbindcred entry and the "Messaging End User Administrator" actually refer to the same password, which is set both in the option and in the userPassword attribute for that user in the LDAP directory. The password is generated randomly by initial configuration and is only used by one single Messaging Server host to bind to the LDAP directory server to perform searches.

The following example uses the **proxyadminpass** option to change the password of the Proxy Administrator account. You should not set passwords from the command line, so this example shows using **msconfig** in interactive mode.

```
msconfig
msconfig> set -prompt proxyadminpass
Password:
Verify:
msconfig# write
msconfig> exit
```

Managing Mail Users, Mailing Lists and Domains

User, mailing list, and domain information is stored as entries in an LDAP directory. An LDAP directory can contain a wide range of information about an organization's employees, members, clients, or other types of individuals that in one way or another "belong" to the organization. These individuals constitute the *users* of the organization.

Topics in this section:

- [Overview of Messaging Server and LDAP](#)
- [To Remove a User from Messaging Server by Using Delegated Administrator](#)
- [To Remove a Domain from Messaging Server using Delegated Administrator](#)

Overview of Messaging Server and LDAP

In the LDAP directory, the information about users is structured for efficient searching, with each user entry identified by a set of attributes. Directory attributes associated with a user can include the user's name and other identification, division membership, job classification, physical location, name of manager, names of direct reports, access permission to various parts of the organization, and preferences of various kinds.

In an organization with electronic messaging services, many if not all users hold mail accounts. Messaging Server stores copies of some account information (uid and quota in particular) on local servers. In general, the LDAP directory is considered authoritative for account information by Messaging Server. Once account information for a mail user is present in the LDAP directory, then the mail server named in the **mailHost** attribute automatically creates that user without any additional mail server specific configuration.

Creating and managing mail users and mailing lists consists of creating and modifying user and mailing list entries in the LDAP directory. This is done by using the Delegated Administrator GUI or command-line utilities, or by directly modifying the LDAP directory information.

Note: In general, the Messaging Server documentation does not describe how to directly modify the LDAP directory. Consult the Directory Server documentation for more information.

To Remove a User from Messaging Server by Using Delegated Administrator

1. Mark the user as deleted by running the **commadmin user delete** command. (See the discussion on removing users, groups, and services from a domain in the *Messaging Server Delegated Administrator's Guide* for more information.)
2. Remove services from the user. A service can be a mailbox or a calendar. For the current version of Messaging Server, the program is called "**msuserpurge**."
3. Permanently remove the user, by invoking the **commadmin domain purge** command.

To Remove a Domain from Messaging Server using Delegated Administrator

1. Mark the domain as deleted by running the **commadmin domain delete** command. (See the discussion on removing users, groups, and services from a domain in the *Messaging Server Delegated Administrator's Guide* for more information.)
2. Remove services from the users of that domain. A service can be a mailbox or a calendar. For Messaging Server, the program is called "**msuserpurge**."
3. Permanently remove the domain, by invoking the **commadmin domain purge** command.

Starting and Stopping Services

How you stop and start Messaging Server services depends on if they are installed in an HA environment.

Topics in this section:

- [To Start and Stop Messaging Server Services in an HA Environment](#)
- [To Start and Stop Messaging Server Services in a non-HA Environment](#)
- [To Start Up, Shut Down, or View the Status of Messaging Services](#)
- [Starting and Stopping a Messaging Server Running in MTA-only Mode](#)

To Start and Stop Messaging Server Services in an HA Environment

While Messaging Server is running under HA control, you cannot use the normal Messaging Server start, restart, and stop commands to control individual Messaging Server services. For example, if you attempt a **stop-msg** in an HA deployment, the system warns that it has detected an HA setup and informs you how to properly stop the system.

The appropriate HA start, stop, and restart commands are shown in [Table 3–2](#). There are no specific HA commands to individually start, restart, or stop other Messaging Server services (for example, SMTP). However, you can run a **stop-msgservice** command to stop/restart individual servers such as **imap**, **pop** or **sched**.

The finest granularity in Oracle Solaris Cluster (formerly known as Sun Cluster) is that of an individual resource. Because Messaging Server is known to Oracle Solaris

Cluster as a resource, the Oracle Solaris Cluster `scswitch` commands affect all Messaging Server services as a whole.

Table 3–2 Start, Stop, Restart in an Oracle Solaris Cluster 3.0/3.1 Environment

Action	Individual Resource	Entire Resource Group
Start	<code>scswitch -e -jresource</code>	<code>scswitch -Z -gresource_group</code>
Restart	<code>scswitch -n -jresourcescswitch -e -jresource</code>	<code>scswitch -R -gresource_group</code>
Stop	<code>scswitch -n -jresource</code>	<code>scswitch -F -gresource_group</code>

Table 3–3 Start, Stop, Restart in an Oracle Solaris Cluster 3.2 Environment

Action	Individual Resource	Entire Resource Group
Start	<code>clrs onlineresource</code>	<code>clrg onlineresource_group</code>
Restart	<code>clrs disableresourceclrs enableresource</code>	<code>clrg restartresource_group</code>
Stop	<code>clrs offlineresource</code>	<code>clrg disableresource_group</code>

Table 3–4 Start, Stop, Restart in Veritas 3.5, 4.0, 4.1 and 5.0 Environments

Action	Individual Resource	Entire Resource Group
Start	<code>hares -onlineresource_name -sys system_name</code>	<code>hagrp -onlinegroup_name -sys system_name</code>
Restart	<code>hares -offline resource_name -sys system_name</code> <code>hares -online resource_name -sys system_name</code>	<code>hagrp -offline service_group -sys system_name</code> <code>hagrp -online service_group -sys system_name</code>
Stop	<code>hares -offline resource_name -sys system_name</code>	<code>hagrp -offline service_group -sys system_name</code>

Table 3–5 Start, Stop, Restart in an Oracle Clusterware 12.1 Environment

Action	Individual Resource
Start	<code>crsctl start resource MS resource -n system</code>
Restart	<code>crsctl stop resource MS resource -n system</code> <code>crsctl start resource MS resource -n system</code>
Stop	<code>crsctl stop resource MS resource -n system</code>

To Start and Stop Messaging Server Services in a non-HA Environment

- Start and stop services from the command line by using the following commands:

```
MessagingServer_home/bin/start-msg
MessagingServer_home/bin/stop-msg
```

Though you can use the command template to start and stop services individually (`MessagingServer_home/bin/stop-msgservice` (where `service` can be `mta`, `imap`, `pop`, `store`, `http`, `ens`, `sched`, `purge`, `mfagent`, `snmp`, `mmp`, `sms`, `metermaid`, `cert`, `dispatcher`, `job_controller` or `watcher`)), do not do so except in specific tasks as described. Certain services have dependencies on other services and must be started in a prescribed order. Complications can arise when trying to start services on their own.

For this reason, you should start and stop all the services together by using the **start-msg** and **stop-msg** commands.

Note: You must first enable services before starting or stopping them. For more information, see ["Enabling and Disabling Services."](#)

Important: If a server process crashes, other processes might hang as they wait for locks held by the server process that crashed. If you are not using automatic restart (see ["Automatic Restart of Failed or Unresponsive Services"](#)), and if any server process crashes, it is generally safer to stop **all** processes, then restart **all** processes. This includes the POP, IMAP, and MTA processes, as well as the **stored** (message store) process, and any utilities that modify the message store, such as **mboxutil**, **deliver**, **reconstruct**, **readership**, or **upgrade**.

To Start Up, Shut Down, or View the Status of Messaging Services

Do not shut down individual services except in the specific tasks as described. Certain services have dependencies on other services and must be started in a prescribed order. Complications can arise when trying to start services on their own. For this reason, you should start and stop all the services together by using the **start-msg** and **stop-msg** commands.

However, when you make a configuration change that requires restart of a service, and you want to minimize service disruption, then use **stop-msgservice** followed by **start-msgservice**. For example, when changing an MTA option that requires a restart of the dispatcher but does not require a restart of the job_controller, it is better to run **stop-msg dispatcher; start-msg dispatcher** to avoid unnecessarily flushing the job_controller's cache. For more information, see ["start-msg"](#) and ["stop-msg."](#)

The services must be enabled to stop or start them. See ["To Specify What Services Can Be Started."](#)

To Specify What Services Can Be Started

By default the following services are started with **start-msg**:

```
start-msg
Connecting to watcher ...
Launching watcher ... 9347
Starting store server .... 9356
Checking store server status ..... ready
Starting purge server .... 9413
Starting imap server .... 9420
Starting pop server .... 9425
Starting http server .... 9437
Starting sched server ... 9451
Starting dispatcher server .... 9461
Starting job_controller server .... 9466
```

The complete list of scopes with an "enable" option is as follows. Some of these, such as **notifytarget**, **service**, and **task** are named scopes. For more information, see the discussion on scope syntax in ["msconfig Command."](#)

```
autorestart
notifytarget
ens
```

```
folderquota
imap
indexer
messagetype
msghash
dbreplicate
mta
metermaid
mmp
pab
pop
purge
relinker
schedule
service
smime
sms_gateway
snmp
store
task
typequota
watcher
http
```

Set both **imap.enable** and **imap.enablesslport** to *0* to disable IMAP. The same goes for POP and HTTP. See *Messaging Server Reference* for details.

Starting and Stopping a Messaging Server Running in MTA-only Mode

To start an MTA-only system, you should also start **imsched**. Before you do this, remove any scheduled jobs that are not appropriate to your installations.

imsched is an individual component of Messaging Server that must be started separately if you are not starting all of Messaging Server. If you start your MTA-only system by using **start-msg imta** or **start-msg mta**, then you do not run the **imsched** process.

To run messaging server in MTA mode only (no store, imap, or pop processes), you can either select the MTA to be only installed and configured during the Messaging Server configuration after initial install (*MessagingServer_home/bin/configure*), or manually disable the message store and **mshttp** process by using the following commands:

```
msconfig set store.enable 0
msconfig set http.enable 0
```

Once you have disabled HTTP and other store processes, you can then start Messaging Server by running the following command:

```
start-msg
Connecting to watcher ...
Launching watcher ... 4034
Starting ens server ... 4035
Starting sched server ... 4036
Starting dispatcher server .... 4038
Starting job_controller server .... 4042
```

All the appropriate processes are started, including **imsched** and **imta**. This way you do not have to remember to start the **sched** process.

Automatic Restart of Failed or Unresponsive Services

Topics in this section:

- [Overview of Messaging Server Monitoring Processes](#)
- [Automatic Restart in High Availability Deployments](#)

Overview of Messaging Server Monitoring Processes

Messaging Server provides two processes called **watcher** and **msprobe** that transparently monitor services and automatically restart them if they crash or become unresponsive (the services hangs). **watcher** monitors server crashes. **msprobe** monitors non-responsive server processes by checking their response times. When a server fails or stops responding to requests, it is automatically restarted. [Table 3-6](#) shows the services monitored by each utility.

Services Monitored by watcher and msprobe

Table 3-6 *Services Monitored by watcher and msprobe*

watcher (crash)	msprobe (unresponsive hang)
IMAP, POP, HTTP, job controller, dispatcher, message store (stored), imsched , MMP. (LMTP/SMTP servers are monitored by the dispatcher and LMTP/SMTP clients are monitored by the job_controller.) The watcher also monitors all processes that access the message store in such a way that they could hold outstanding message store locks when they crash. This includes ims_master , lmtplib_server , and store utilities.	IMAP, POP, HTTP, cert, job controller, message store (stored), imsched , ENS, LMTP, SMTP

Enabling the watcher (**watcher.enable 1**, default) monitors process failures and unresponsive services and logs error messages to the **default** log file indicating specific failures. To enable automatic server restart, use **msconfig** to set **base.autorestart.enable 1**. By default, this option is set to no (0).

If any of the message store services fail or freeze, all message store services that were enabled at start-up are restarted. For example, if **imapd** fails, at the least, **stored** and **imapd** are restarted. If other message store services were running, such as the POP or HTTP servers, then those are restarted as well, whether or not they failed.

Automatic restart also works if a message store utility fails or freezes. For example, if **mboxutil** fails or freezes, the system automatically restarts all the message store servers. However, it does not restart the utility. **msprobe** runs every 10 minutes. Service and process restarts are performed once within a 10-minute period (and are configurable by using **base.autorestart**). If a server fails more than once during this designated period of time, then the system stops trying to restart this server. If this happens in an HA system, Messaging Server is shut down and a failover to the other system occurs.

Whether or not **base.autorestart.enable** is enabled, the system still monitors the services and sends failure or non-response error messages to the console and *MessagingServer_home/data/log/watcher* listens to port 49994 by default, but this is configurable with the **watcher.port** option.

A watcher log file is generated in *MessagingServer_home/data/log/watcher*. This log file is not managed by the logging system (no rollover or purging) and records all server starts and stops. The following is an example log:

```
watcher process 13425 started at Mon June 4 11:23:54 2012

Watched 'imapd' process 13428 exited abnormally
Received request to restart: store imap pop http
Connecting to watcher ...
Stopping http server 13440 .... done
Stopping pop server 13431 ... done
Stopping pop server 13434 ... done
Stopping pop server 13435 ... done
Stopping pop server 13433 ... done
imap server is not running
Stopping store server 13426 .... done
Starting store server .... 13457
checking store server status ..... ready
Starting imap server ..... 13459
Starting pop server ..... 13462
Starting http server ..... 13471
```

See "[Monitoring Using msprobe and watcher Functions](#)" for more details on how to configure this feature.

msprobe is controlled by **imsched**. If **imsched** crashes, this event is detected by **watcher** and triggers a restart (if **autorestart** is enabled). However, in the rare occurrence of **imsched** hanging, you need to kill **imsched** with a **killimsched_pid** command, which causes the watcher to restart it.

Automatic Restart in High Availability Deployments

Table 3–7 shows the configuration options to be set for automatic restart in high availability deployments:

Table 3–7 HA Automatic Restart Options

Option	Description/ HA Value
watcher.enable	Enable watcher on start-msg startup. Default is enabled (1).
base.autorestart.enable	Enable automatic restart of failed or frozen (unresponsive) servers including IMAP, POP, HTTP, job controller, dispatcher, and MMP servers. Default is enabled).
base.autorestart.timeout	Failure retry time-out. If a server fails more than once during this designated period of time, then the system stops trying to restart this server. If this happens in an HA system, Messaging Server is shutdown and a failover to the other system occurs. The value (set in seconds) should be set to a period value longer than the msprobe interval. (See schedule.task : in the following section). Default is 600.
schedule.task:msprobe.crontab	msprobe runs schedule. A crontab style schedule string. Default is 5,15,25,35,45,55 * * * * lib/msprobe . To disable, run the following command: msconfig set schedule.task:msprobe.enable 0

Scheduling Automatic Tasks

Topics in this section include:

- [Overview of Scheduling Automatic Tasks](#)
- [Scheduler Examples](#)
- [Pre-defined Automatic Tasks](#)

Overview of Scheduling Automatic Tasks

Messaging Server provides a general task scheduling mechanism by using a process called **imsched**. It is intended for scheduling Messaging Server processes. It is enabled by setting the **schedule.task** option. If you modify the schedule, either restart the scheduler with the command **stop-msg sched** and **start-msg sched**, or refresh the scheduler process (**refresh sched**).

This option requires a command and a schedule on which to execute the command. The format is as follows:

```
schedule.task:<taskname>.crontab = <schedule>
```

where:

- *taskname* is the name of the command to run, for example, **expire**, **msprobe**, and so on.
- *schedule* is a non-empty string with the following format:

```
<minute> <hour> <day-of-month> <month-of-year> <day-of-week> <command args>
```
- *command args* can be any Messaging Server command and its arguments. Paths can be relative to *MessagingServer_home* or absolute paths. See "[Pre-defined Automatic Tasks](#)" for relative path examples.

minute hour day-of-month month-of-year day-of-week is the schedule for running the command. It follows the UNIX **crontab** time format.

The values are separated by a space or tab and can be 0-59, 0-23, 1-31, 1-12 or 0-6 (with 0=Sunday) respectively. Each time field can be either an asterisk (meaning all legal values), a list of comma-separated values, or a range of two values separated by a hyphen. Days can be specified by both day of the month and day of the week and both are required if specified. For example, setting the 17th day of the month and Tuesday only runs the command on the 17th day of a month when it is Tuesday.

If you modify scheduler, either restart the scheduler with the command **stop-msg sched** and **start-msg sched**, or refresh the scheduler by running **refresh sched**.

- To disable a scheduled task:

```
msconfig set schedule.task:<taskname>.enable = 0
refresh sched
```

Scheduler Examples

Run **imexpire** at 12:30am, 8:30am, and 4:30pm:

```
msconfig set schedule.task:expire.crontab "30 0,8,16 * * * bin/imexpire"
```

Run **imsbackup** Monday through Friday at midnight (12AM):

```
msconfig set schedule.task:msbackup.crontab "0 0 * * 1-5 bin/imsbackup -f"
```

```
backupfile /primary"
```

Pre-defined Automatic Tasks

At installation, Messaging Server creates, schedules and enables the following set of pre-defined automatic tasks:

The following automatic tasks are set and enabled for the message store:

```
schedule.task:expire.crontab = 0 23 * * * bin/imexpire
schedule.task:snapshot.crontab = 0 2 * * * bin/imdbverify -s -m
schedule.task:snapshotverify.crontab = 5,15,25,35,45,55 * * * * bin/imdbverify
```

The following automatic tasks are set and enabled for the MTA:

```
schedule.task:purge.crontab = 0 0,4,8,12,16,20 * * * bin/imsimta purge
schedule.task:return_job.crontab = 30 0 * * * lib/return_job
```

The following automatic task is set and enabled for the message store:

```
schedule.task:msprobe.crontab = 5,15,25,35,45,55 * * * * lib/msprobe
```

Configuring a Greeting Message

Messaging Server enables you to create an email greeting message to be sent to each new user.

Topics in this section:

- [To Create a New User Greeting](#)
- [To Set a Per-Domain Greeting Message](#)

To Create a New User Greeting

- To create a new-user greeting:

```
msconfig set base.welcomemsg <Message>
```

Where *Message* must contain a header (with at least a subject line), followed by \$\$, then the message body. The \$ represents a new line.

For example, to enable this option, you can set the following configuration variables:

```
msconfig set base.welcomemsg 'Subject: Welcome!! $$ example.com welcomes you to
the premier Internet experience in Dafandzadgad!'
```

Depending on the shell that you are using, it might be necessary to append a special character before \$ to escape the special meaning of \$. (\$ is often the escape character for the shell.) Alternatively, you can do this within the **msconfig** prompt so that you do not need to the \$. Simply run **msconfig**, then issue the **setoptionvalue** command.

To Set a Per-Domain Greeting Message

Whenever you create a new hosted domain, create per-domain greeting messages for your supported languages. If this is not done, the generic greeting message set by **base.welcomemsg** is sent.

You can set a greeting message for new users in each domain. The message can vary depending on the user's, the domain's, or the site's preferred language. This is done by

setting the **mailDomainWelcomeMessage** attribute in the desired LDAP domain entry. The attribute syntax is as follows:

```
mailDomainWelcomeMessage;lang-userprefLang
```

```
mailDomainWelcomeMessage;lang-domain_prefLang
```

```
mailDomainWelcomeMessage;lang-gen.sitelanguage
```

The following example sets the domain welcome message for English:

```
mailDomainWelcomeMessage;lang-en: Subject: Welcome!! $$Welcome to the mail system.
```

The following example sets the domain welcome message for French:

```
mailDomainWelcomeMessage;lang-fr: Subject: Bienvenue!! $$Bienvenue a example.org!
```

Using these examples, assume the following:

- The domain is *example.org*.
- A new user belongs to this domain.
- The user's preferred language is French as specified by the LDAP attribute **preferredlanguage**.
- The **example.org** domain has the above English and French welcome messages available.
- The site language is **en** as specified by **gen.sitelanguage**.

For a list of supported locales and their language value tag, see the Directory Server Reference.

When users log in for the first time, they receive the French greeting. If the French welcome message isn't available, they get the English greeting.

Greeting Message Theory of Operations

Greeting messages can be set by both the LDAP attribute **mailDomainWelcomeMessage**, the **base.welcomemsg** option, and the **message_language:langcode.welcomemsg** option. The **base.welcomemsg** option is the default, the **message_language:langcode.welcomemsg** option is language-code specific. The order in which a message is chosen, with the top one having the highest preference, is shown below:

```
mailDomainWelcomeMessage;lang-user_prefLang
```

```
mailDomainWelcomeMessage;lang-domain_prefLang
```

```
mailDomainWelcomeMessage;lang-gen.sitelanguage
```

```
mailDomainWelcomeMessage
```

```
base.welcomemsg;lang-"$user-prefLang"
```

```
base.welcomemsg;lang-"$domain-prefLang"
```

```
base.welcomemsg;lang-"$gen.sitelanguage"
```

```
base.welcomemsg
```

The algorithm works as follows: if there are no domains (or there are, but there is no per domain welcome message provisioned for them), a welcome message is configured with the **base.welcomemsg** option, if specified. If a user has a preferred language (set with the **preferredlanguage** LDAP attribute) and

base.welcomemsg;lang-user_prefLang is set, the user will receive that welcome message at the time of their first log in to the server. If **base.welcomemsg;lang-gen.sitelanguage** is set, and **preferredlanguage** is not set, but the site language is set (using **base.sitelanguage** option), user will receive that message. If no language tag option is set and a untagged **base.welcomemsg** is set, then that message will be sent to the user. If none of the values are set, user will not receive any welcome message.

If the user is in a domain, then similar to the discussion above, the user might receive one of **mailDomainWelcomeMessage;lang-xx**, depending on which one is available in the list and in the order given.

Example: Domain is **example.org**. The domain preferred language is German (de). But the new user in this domain has preferred language of Turkish (tr). Site language is English. The following values are available (**mailDomainWelcomeMessage** are attributes of the domain example.org):

```
mailDomainWelcomeMessage;lang-fr
mailDomainWelcomeMessage;lang-ja
base.welcomemsg;lang-de
base.welcomemsg;lang-en
base.welcomemsg
```

According to the algorithm, the message sent to the user is **base.welcomemsg;lang-de**.

Setting a User-Preferred Language

Topics in this section:

- [Overview of Setting a User-Preferred Language](#)
- [To Set a Domain Preferred Language](#)
- [To Specify a Site Language](#)

Overview of Setting a User-Preferred Language

Administrators can set a preferred language for the GUI and server-generated messages by setting the attribute **preferredLanguage** in the user's LDAP entry.

When the server sends messages to users outside of the server's administrative domain it does not know what their preferred language is unless it is responding to an incoming message with a preferred language specified in the incoming message's header. The header fields (**Accept-Language**, **Preferred-Language** or **X-Accept-Language**) are set according to attributes specified in the user's mail client.

If there are multiple settings for the preferred language, the server chooses the preferred language. For example, if a user has a preferred language attribute stored in the Directory Server and also has a preferred language specified in their mail client, the server chooses the preferred language in the following order:

1. The **Accept-Language** header field of the original message.
2. The **Preferred-Language** header field of the original message.
3. The **X-Accept-Language** header field of the original message.
4. The preferred language attribute of the sender (if found in the LDAP directory).

To Set a Domain Preferred Language

A domain preferred language is a default language specified for a particular domain. For example, you can specify Spanish for a domain called **mexico.example.org**. Administrators can set a domain preferred language by setting the attribute **preferredLanguage** in the domain's LDAP entry.

To Specify a Site Language

You can specify a default site language for your server as follows. The site language is used to send language-specific versions of messages if no user preferred language is set.

- *Command Line:* Specify a site language as follows:

```
msconfig set base.sitelanguage <value>
```

where *value* is one of the local supported languages. See the Directory Server documentation for a list of supported locales and the language value tag.

Encryption Settings

This is described in enabling SSL and selecting ciphers in the *Messaging Server Security Guide*, which also contains background information on all security and access-control topics for Messaging Server.

Setting a Failover LDAP Server

It is possible to specify more than one LDAP server for the user/group directory so that if one fails another takes over.

To set a failover LDAP server:

1. Set **base.ugldaphost** to the multiple replicated LDAP servers. For example:

```
msconfig set base.ugldaphost "ldap1.example.com ldap2.example.com:389"
```

2. If you are using a compiled MTA configuration then recompile the MTA configuration file.

```
imsimta cnbuild
```

3. Restart Messaging Server.

```
stop-msg
start-msg
```

Configuring and Administering Multiplexor Services

This chapter describes the Messaging Multiplexor (MMP) for standard mail protocols (POP, IMAP, and SMTP).

Multiplexor Services in Unified Configuration Overview

The MMP configuration is stored in the Unified Configuration. The following MMP configuration files are no longer used in Unified Configuration:

Table 4–1 Legacy MMP Configuration Files

File Type	Legacy File Names
POP SSL MMP Encryption File	PopProxyAService.cfg
POP Services Configuration Template	PopProxyAService-def.cfg
IMAP SSL MMP Encryption File	ImapProxyAService.cfg
IMAP Services Configuration Template	ImapProxyAService-def.cfg
Service Starting Configuration File	AService.cfg
Service Starting Configuration Template	AService-def.cfg
SMTP SSL MMP Encryption File	SmtproxyAService.cfg
SMTP Services MMP Configuration Template	SmtproxyAService-def.cfg

In Unified Configuration, you enable and modify the MMP configuration by running the **msconfig** command to set the appropriate MMP options. The **ServiceList** and **SSLports** options are gone in Unified Configuration. You now use the **imapproxy**, **popproxy**, and **smtproxy** configuration groups, and the **tcp_listen** option. Use the following commands to view the initial MMP configuration settings.

```
msconfig
msconfig> show mmp*
role.mmp.enable = 0
msconfig> show imapproxy*
role.imapproxy.connlimits = :20
```

```
role.imaproxy.tcp_listen:imaproxy1.tcp_ports = 143
msconfig> show popproxy*
role.popproxy.connlimits = :20
role.popproxy.tcp_listen:popproxy1.tcp_ports = 110
msconfig> show submitproxy*
role.submitproxy.connlimits = :20
```

In addition, the **ssl_ports** option works the like **tcp_ports** option but enables SSL services (thus fixing the problem in legacy configuration, where an SSL proxy service had to be listed in both **ServiceList** and **SSLPorts** options).

The following examples commands show how to update the MMP configuration:

- To enable MMP: **msconfig set mmp.enable 1**
- To change an IMAP proxy option: **msconfig set imaproxy.optionvalue**
- To change a POP proxy option: **msconfig set popproxy.optionvalue**
- To change an SMTP proxy option: **msconfig set smtpproxy.optionvalue**
- To set a certmap default option: **msconfig set base.certmap:default.optionvalue**

See the *Messaging Server Reference*, or option descriptions in the **msconfig** online help, for more information.

Multiplexor Services

A multiplexor is necessary to achieve horizontal scalability (the ability to support more users by adding more machines), because it provides a single domain name that can be used to connect indirectly to multiple mail stores. A multiplexor can also provide security benefits.

In Unified Configuration, MMP is no longer managed separately from Oracle Communications Messaging Server.

Multiplexor Benefits

Message stores on heavily used messaging servers can grow quite large. Spreading user mailboxes and user connections across multiple servers can therefore improve capacity and performance. In addition, it can be more cost-effective to use several small server machines than one large, high-capacity, multiprocessor machine.

If the size of your mail-server installation requires the use of multiple message stores, your organization can benefit in several ways from using the multiplexor. The indirect connection between users and their message stores, coupled with the ease of reconfiguration of user accounts among messaging servers allows for the following benefits:

- **Simplified User Management** Because all users connect to one server (or more, if you have separate multiplexor machines for POP, IMAP, SMTP, or web access), you can preconfigure email clients and distribute uniform login information to all users. This simplifies your administrative tasks and reduces the possibility of distributing erroneous login information.

For especially high-load situations, you can run multiple multiplexor servers with identical configurations and manage connections to them by DNS round robin or by using a load-balancing system. Because the multiplexors use information stored in the LDAP directory to locate each user's Messaging Server, moving a user to a new server is simple for the system administrator and transparent to the user. The administrator can move a user's mailbox from one Messaging Server

host to another, and then update the user's entry in the LDAP directory. The user's mail address, mailbox access, and other client preferences need not change.

- **Improved Performance** If a message store grows prohibitively large for a single machine, you can balance the load by moving some of the message store to another machine.

You can assign different classes of users to different machines. For example, you can choose to locate premium users on a larger and more powerful machine.

The multiplexors perform some buffering so that slow client connections (through a modem, for example) do not slow down the Messaging Server.

- **Decreased Cost** Because you can efficiently manage multiple Messaging Server hosts with a multiplexor, you might be able to decrease overall costs by purchasing several small server machines that together cost less than one very large machine.
- **Better Scalability** With the multiplexors, your configuration can expand easily. You can incrementally add machines as your performance or storage-capacity needs grow, without replacing your existing investment.
- **Minimum User Downtime.** Using the multiplexors to spread a large user base over many small store machines isolates user downtime. When an individual server fails, only its users are affected.
- **Increased Security** You can use the server machine on which the multiplexor is installed as a firewall machine. By routing all client connections through this machine, you can restrict access to the internal message store machines by outside computers. The multiplexors support both unencrypted and encrypted communications with clients.

About Messaging Multiplexor

The Messaging Multiplexor (MMP) is a specialized messaging server that acts as a single point of connection to multiple back-end messaging servers. With Messaging Multiplexor, large-scale messaging service providers can distribute POP and IMAP user mailboxes across many machines to increase message store capacity. All users connect to the single multiplexor server, which redirects each connection to the appropriate messaging server.

If you provide electronic mail service to many users, you can install and configure the Messaging Multiplexor so that an entire array of Messaging Server hosts appear to your mail users to be a single host.

The Messaging Multiplexor is provided as part of Messaging Server. You can install the MMP at the same time you install Messaging Server or other Communications Suite servers, or you can install the MMP separately at a later time. The MMP supports the following items:

- Both unencrypted and encrypted (SSL) communications with mail clients.
- Client certificate-based authentication, described in "[Certificate-Based Client Authentication](#)."
- User pre-authentication, described in "[User Pre-Authentication](#)."
- Virtual domains that listen on different IP addresses and automatically append domain names to user IDs, described in "[MMP Virtual Domains](#)."
- Multiple installations of the MMP on different servers.

- Enhanced LDAP searching.
- POP before SMTP service for legacy POP clients.


This section consists of the following subsections:

- [How the Messaging Multiplexor Works](#)
- [Encryption \(SSL\) Option](#)
- [Certificate-Based Client Authentication](#)
- [To Enable Certificate-based Authentication for Your IMAP or POP Service](#)
- [User Pre-Authentication](#)
- [MMP Virtual Domains](#)
- [About SMTP Proxy](#)

How the Messaging Multiplexor Works

The MMP is a multithreaded server that facilitates distributing mail users across multiple server machines. The MMP handles incoming client connections destined for other server machines (the machines on which user mailboxes reside). Clients connect to the MMP itself, which determines the correct server for the users, connects to that server, and then passes data between the client and server. This capability allows Internet service providers and other large installations to spread message stores across multiple machines (to increase capacity) while providing the appearance of a single mail host for users (to increase efficiency) and for external clients (to increase security). The following figure shows clients and servers in an MMP installation.

Figure 4–1 Clients and Servers in an MMP Installation



All POP, IMAP, and SMTP clients work with the Messaging Multiplexor. The MMP accepts connections, performs LDAP directory lookups, and routes the connections appropriately. As is typical with other mail server installations, each user is assigned a specific address and mailbox on a specific Messaging Server. However, all connections are routed through the MMP.

In more detail, these are the steps involved in establishing a user connection:

1. A user's client connects to the MMP, which accepts preliminary authentication information (user name).
2. The MMP queries the Directory Server to determine which Messaging Server contains that user's mailbox.

3. The MMP connects to the proper Messaging Server, replays authentication, then acts as a pass-through pipe for the duration of the connection.

Encryption (SSL) Option

Messaging Multiplexor supports both unencrypted and encrypted (SSL) communications between the Messaging Server(s) and their mail clients. The current version of Messaging Server supports the new certificate database format (**cert8.db**).

When SSL is enabled, the MMP IMAP supports both STARTTLS on the standard IMAP port and IMAP+SSL on port 993. The MMP can also be configured to listen on port 995 for POP+SSL.

In legacy configuration, to enable SSL encryption for your IMAP, POP, and SMTP services, you would uncomment the appropriate SSL settings from the **.cfg** files. In Unified Configuration, you use the **msconfig** command to set the appropriate options. You must also set the list of all IMAP, POP, and SMTP server ports regardless of whether or not they are secure. See "[Configuring MMP with SSL or Client Certificate-Based Login](#)" for details.

By default, SSL is not enabled. To enable SSL, you must install an SSL server certificate. Then, you should use the **msconfig** command to set the SSL options. For a list of the SSL options, see the **ssl*** options in the **msconfig** online help.

Certificate-Based Client Authentication

In Unified Configuration, the certificate mapping file (**certmap.conf**), which matches a client's certificate to the correct user in the Users/Groups Directory Server, is no longer used. Instead, you use the following **msconfig** command to set the appropriate options:

```
msconfig set base.certmap:default.<option> <value>
```

The **default** can be replaced with the certificate **issuerDN** to have configuration specific to that certificate. That replaces the other groups in the **certmap.conf** file.

To use certificate-based client authentication, you must also enable SSL encryption as described in "[Encryption \(SSL\) Option](#)."

You also have to configure a store administrator. You can use the mail administrator, but it is recommended that you create a unique user ID, such as **mmpstore** for this purpose so that you can set permissions as needed.

In Unified Configuration, the MMP supports the **dncomps** and **filtercomps** options. The values of these two options has the form *fromattr=toattr*. A *fromattr* value in a certificate's subjectDN can be used to form an LDAP query with the *toattr=value* element. For example, a certificate with a subjectDN of "**cn=Pilar Lorca, ou=pilar, o=example.com**" could be mapped to an LDAP query of "**(uid=pilar)**" with the line:

```
msconfig> set base.certmap:default.filtercomps ou=uid
```

To Enable Certificate-based Authentication for Your IMAP or POP Service

1. Decide on the user ID you intend to use as store administrator. While you can use the mail administrator for this purpose, it is recommended that you create a unique user ID for store administrator (for example, **mmpstore**).
2. Make sure that SSL encryption is (or will be) enabled as described in "[Encryption \(SSL\) Option](#)."

3. Configure the MMP to use certificate-based client authentication by specifying default **certmap** option in your configuration. For example:

```
msconfig set base.certmap:default.dncomps ""
```

4. Install at least one trusted CA certificate, as described in the discussion on installing certificates of trusted CAs in the *Messaging Server Security Guide*.

User Pre-Authentication

The MMP provides you with the option of pre-authenticating users by binding to the directory as the incoming user and logging the result.

Note: Enabling user pre-authentication reduces server performance.

The log entries are in the format: *datetime*(**sid 0xhex**) *user name***pre-authenticated - clientIPaddress, server IPaddress**

Where *date* is in the format **yyyymmdd**, *time* is in the time configured on the server in the format **hhmmss**, *hex* is the session identifier (**sid**) represented as a hexadecimal number, the **username** includes the virtual domain (if any), and the IP address is in dot-quad format.

MMP Virtual Domains

An MMP *virtual domain* is a set of configuration settings associated with one or more server IP addresses. The primary use of this feature is to provide different default domains for each server IP address. The **hosteddomains** option defaults to 1 (enabled).

A user can authenticate to the MMP with either a short-form userID or a fully qualified userID in the form **user@domain**. When a short-form userID is supplied, the MMP will append the **defaultdomain** setting, if specified. Consequently, a site which supports multiple hosted domains can permit the use of short-form user IDs simply by associating a server IP address and MMP virtual domain with each hosted domain.

To configure a virtual domain option, use the following command:

```
msconfig set vdomain:<IP-address>.<option> <value>
```

For example, to set the default domain, use the following command:

```
msconfig set vdomain:192.0.2.0.defaultdomain example.com
```

When set, virtual domain configuration option values override global configuration option values.

You can specify the following configuration options for a virtual domain:

```
authcachettl  
authenticationldapattributes  
authservice  
authservicettl  
binddn  
bindpass  
clientlookup  
crams  
debugkeys  
defaultdomain  
domainsearchformat  
ehlokeywords
```

```
failovertimeout
hosteddomains
ldapcachesize
ldapcachettl
mailhostattrs
popbeforesmtpkludgechannel
preauth
replayformat
restrictplainpasswords
searchformat
smtpproxypassword
smtprelays
ssladjustciphersuites
sslnicknames
storeadmin
storeadminpass
tcpaccess
tcpaccessattr
virtualdomaindelim
```

For more information on these configuration options, see the **msconfig** online help or the *Messaging Server Reference*.

Tip: To view the reference information for these configuration options, use the following URL:

```
http://msg.wikidoc.info/index.php/MMP_Reference#<optionname>
```

For example, to see the description for the **tcpaccess** option, use the following URL:

```
http://msg.wikidoc.info/index.php/MMP_Reference#tcpaccess
```

About SMTP Proxy

The MMP includes an SMTP submission proxy, which is disabled by default. Most sites do not need the SMTP proxy because Internet Mail standards already provide an adequate mechanism for horizontal scalability of SMTP (DNS MX records).

The SMTP proxy is useful for the security features it provides. First, the SMTP proxy is integrated with the POP proxy to implement the POP before SMTP authorization facility required by some legacy POP clients. For more information, see the discussion on using the MMP SMTP proxy in the *Messaging Server Installation and Configuration Guide*. In addition, an investment in SSL acceleration hardware can be maximized by using the SMTP proxy.

Setting Up the Messaging Multiplexor

During the initial runtime configuration of Messaging Server, you determined if you wanted to configure the MMP on a machine. You could either set it up on the same machine as your Messaging Server or set it up on a separate machine.

Note: MMP does not cache DNS results. A high quality caching DNS server on the local network is a requirement for a production deployment of Messaging Server.

The following sections describe how to set up the MMP:

- [Before You Configure MMP](#)
- [Multiplexor Configuration](#)
- [To Configure the MMP](#)
- [Multiplexor Configuration Options](#)
- [Starting the Multiplexor](#)
- [Modifying an Existing MMP](#)

Before You Configure MMP

Before configuring the MMP:

1. Choose the machine on which you will configure the MMP. It is best to use a dedicated machine for the MMP.

Note: It is recommended that the MMP not be enabled on a machine that is also running either the POP or IMAP servers. If you install MMP on the same machine as Messaging Server, you must make sure that the POP and IMAP servers are set to non-standard ports. That way, the MMP and Messaging Server ports do not conflict with one another.

2. On the machine where the MMP is to be configured, create a UNIX system user to be used by the MMP. This new user must belong to a UNIX system group. See the discussion on creating UNIX system users and groups in the *Messaging Server Installation and Configuration Guide*.
3. Set up the Directory Server and its host machine for use with Messaging Server, if they are not already set up. See the discussion on preparing directory server for Messaging Server configuration in the *Messaging Server Installation and Configuration Guide*.
4. If the MMP is upgraded before the back-end servers, set the **capability** option to match the response to the **capability** command from the older back end. See the discussion on the capability option in the *Messaging Server Reference* for more information.

Multiplexor Configuration

To configure the MMP, you must use the Messaging Server **configure** program, which gives you the option of enabling the Messaging Multiplexor. For detailed information about the configure program, see the *Messaging Server Installation and Configuration Guide*.

To Configure the MMP

1. Install Messaging Server software on the machine where you are installing and configuring the MMP.
2. Configure the MMP by creating the Messaging Server Initial Runtime Configuration. See the discussion on creating the initial Messaging Server runtime configuration in the *Messaging Server Installation and Configuration Guide*.

Multiplexor Configuration Options

You control how the MMP operates by specifying various configuration options in the Unified Configuration. See the *Messaging Server Reference* for more information.

Starting the Multiplexor

To start, stop, or refresh an instance of the Messaging Multiplexor, use the one of the commands in [Table 4-2](#). These commands are located in the *MessagingServer_home/bin* directory.

Table 4-2 MMP Commands

Option	Description
<code>start-msg mmp</code>	Starts the MMP (only if the MMP is enabled and one is not already running).
<code>stop-msg mmp</code>	Stops the most recently started MMP.
<code>refresh mmp</code>	Causes an MMP that is already running to refresh its configuration without disrupting any active connections.

Modifying an Existing MMP

1. To modify an existing instance of the MMP, use the `msconfig` command to edit the configuration as necessary.
2. The run either `refresh mmp` or `stop-msg mmp; start-msg mmp`.

Use the former only if you changed "refreshable" options and the latter if you changed any "non-refreshable" options.

Configuring MMP with SSL or Client Certificate-Based Login

This section describes how to configure MMP with SSL or client certificate-based login.

Note: It is assumed that the MMP is installed on a machine that does not have a Message Store or MTA.

To Configure MMP with SSL

1. Generate and install the certificate by using the `certutil` command. See the discussion on certificate based authentication for Messaging Server in the *Messaging Server Security Guide* for details.

2. Set the password used for the certificate file. For example:

```
msconfig
msconfig> set "sectoken:Internal (Software) Token.tokenpass" newpassword
msconfig> write
```

The default setting for this password was provided during initial configuration, but it might be different. It must match the password that was used when the certificate db was created by running the `certutil -N` command.

3. Set either `sslenable` on the relevant proxy (for STARTTLS) and/or set the `ssl_ports` on a `tcp_listen` for the appropriate proxy. In general, the default settings cover the remainder of the configuration and you do not need to be changed.

4. Start the MMP:
`MessagingServer_home/bin/start-msg`
5. If you do not want to use SSL between the MMP and the back-end server, then set the `sslbacksideport` option to 0 for `imapproxy` and `popproxy` as appropriate.

To Configure MMP with Client Certificate-based Login

If you want client certificate based login, do the following:

1. Get a copy of a client certificate and the CA certificate which signed it.
2. Import the CA certificate as a Trusted Certificate Authority (see the discussion on obtaining and managing certificates in the *Messaging Server Security Guide*).
3. Use the Store Administrator you created during your Messaging Server installation. For more information, see "[Specifying Administrator Access to the Message Store](#)."

4. Create a `certmap.conf` file for the MMP. For example:

```
msconfig> set base.certmap:default.dncomps ""
msconfig# set base.certmap:default.filtercomps "e=mail"
```

This means to search for a match with the `e` field in the certificate DN by looking at the mail attribute in the LDAP server.

5. Use the `msconfig` command to update the configuration with the following options:
 - a. Set `storeadmin` and `storeadminpass` to values from Step 3.
 - b. Set `usergroupdn` to the root of your Users and Groups tree.
6. If you want client certificates with POP3, repeat Step 5 for the `popproxy` group.
7. If the MMP is not already running, start it with the following command in the `MessagingServer_home/bin` directory:
start-msg mmp
8. Import the client certificate into your client. In Netscape Communicator, click the padlock (Security) icon, then select Yours under Certificates, then select Import a Certificate and follow the instructions.


Note: All your users have to perform this step if you want to use client certificates everywhere.

A Sample Topology

The fictional Example Corporation has two Messaging Multiplexors on separate machines, each supporting several Messaging Servers. POP and IMAP user mailboxes are split across the Messaging Server machines, with each server dedicated exclusively to POP or exclusively to IMAP. (You can restrict client access to POP services alone by removing the `imapproxy` entry from the MMP configuration. Likewise, you can restrict client access to IMAP services alone by removing the `popproxy` entry from the MMP configuration). Each Messaging Multiplexor also supports only POP or only IMAP. The LDAP directory service is on a separate, dedicated machine.

The following figure illustrates this topology.

Figure 4–2 Multiple MMPs Supporting Multiple Messaging Servers



MMP Tasks

This section describes the following miscellaneous MMP configuration tasks:

- [To Configure Mail Access with MMP](#)
- [To Set a Failover MMP LDAP Server](#)

To Configure Mail Access with MMP

The MMP does not make use of the **PORT_ACCESS** mapping table. If you wish to reject SMTP connections from certain IP addresses and you are using the MMP, you must use the **tcpaccessattr** option.

To Set a Failover MMP LDAP Server

It is possible to specify more than one LDAP server for the MMP so that if one fails another takes over. Modify the **ugldaphost** option. For example:

```
msconfig set ugldaphost "ldap1.example.com ldap2.example.com"
```

Note: Make sure there is a space between the host names in the preceding configuration, and because of that space, to enclose the hosts in quotation marks.

MTA Concepts

This chapter provides a conceptual description of the Oracle Communications Messaging Server MTA.

Topics:

- The MTA Functionality
- MTA Architecture and Message Flow Overview
- The Dispatcher
- Rewrite Rules
- Channels
- The MTA Directory Information
- The Job Controller
- On Demand Mail Relay
- Priority Message Handling

The MTA Functionality

The Message Transfer Agent, or *MTA* is a component of the Messaging Server (see [Figure 5-1, "High Level Message Store and MTA Architecture"](#)). At its most basic level, the MTA is a message router. It accepts messages from other servers, reads the address, and routes it to the next server on way to its final destination, typically a user's mailbox.

Over the years, a lot of functionality has been added to the MTA, and with it, size, power, and complexity. These MTA functions overlap, but, in general, can be classified as follows:

- **Routing.** Accepts a message, expands or transforms it if necessary (for example if it is an alias), and routes it to the next server, channel, program, file, or whatever. The routing function has been expanded to allow administrator specification of the internal and external mechanics of how messages are routed. For example, it is possible to specify things such as SMTP authentication, use of various SMTP commands and protocol, TCP/IP or DNS lookup support, job submission, process control and message queueing and so on.
- **Address Rewriting.** Envelope addresses are often rewritten as part of the routing process, but envelope or header addresses can also be rewritten to a more desired or appropriate form.

- **Filtering.** The MTA can filter messages based on address, domain, possible virus or spam content, size, IP address, header content, and so on. Filtered messages can be discarded, rejected, modified, sent to a file, sent to a program, or be sent to the next server on its way to a user mailbox.
- **Content Modification.** Message headers or content can be modified. Example: making a message readable to a specific client or in a specific character set or checking for spam or viruses.
- **Auditing.** Tracking who submitted what, where and when.

A number of subcomponents and processes support these functions and are shown in Figure 5–2, "MTA Architecture". This information describes these subcomponents and processes. In addition, a number of tools allow system administrators to enable and configure these functions. These include Unified Configuration options, mapping tables, channel options, channels, and rewrite rules. These are described in the following MTA information:

- [About MTA Services](#)
- [Channel Configuration in the *Messaging Server Reference*](#)
- [Configuring Rewrite Rules](#)
- [Integrating Spam and Virus Filtering Programs in the *Messaging Server Reference*](#)
- [LMTP Delivery](#)
- [Mail Filtering and Access Control](#)
- [Managing Logging](#)
- [Monitoring Messaging Server](#)
- [Security and Access Control in the *Messaging Server Security Guide*](#)
- [Troubleshooting the MTA](#)
- [Using Predefined Channels](#)
- [Vacation Automatic Message Reply](#)

Figure 5–1 High Level Message Store and MTA Architecture





Figure 5–2 MTA Architecture



MTA Architecture and Message Flow Overview

This section provides a short overview of MTA architecture and message flow (see Figure 5–2). The MTA is a highly complex component and this figure is a *simplified* depiction of messages flowing through the system. In fact, this picture is not a perfectly accurate depiction of all messages flowing through the system. For purposes of conceptual discussion, however, it must suffice.

Dispatcher and SMTP Server (Slave Program)

Messages enter the MTA from the Internet or intranet via SMTP sessions. When the MTA receives a request for an SMTP connection, the MTA *dispatcher* (a multithreaded connection dispatching agent), executes a *slave* program (`tcp_smtp_server`) to handle the SMTP session. The dispatcher maintains pools of multithreaded processes for each service. As additional sessions are requested, the dispatcher activates an SMTP server program to handle each session. A process in the Dispatcher's process pool may concurrently handle many connections. Together the dispatcher and slave program perform a number of different functions on each incoming message. Three primary functions are:

- **Message blocking.** Messages from specified IP addresses, mail addresses, ports, channels, header strings and so on, may be blocked (see "Mail Filtering and Access Control").
- **Address changing.** Incoming **From:** or **To:** addresses may be rewritten to a different form.
- **Channel enqueueing.** Addresses are run through the rewrite rules to determine which channel the message should be sent.

For more information, see "[The Dispatcher](#)."

Routing and Address Rewriting

SMTP servers enqueue messages, but so can a number of other channels including, the conversion channel and reprocess channel. A number of tasks are achieved during this phase of delivery, but the primary tasks are:

- Alias expansion.
- Running the addresses through the rewrite rules which do two things:
 - Rewrite the domain part of addresses into a desired format.
 - Direct messages to the appropriate channel queue.

Channel

The channel is the fundamental MTA component used for message processing. A channel represents a message connection with another system (for example, another MTA, another channel, or the local message store). As mail comes in, different messages require different routing and processing depending on the message's source and destination. For example, mail to be delivered to a local message store is processed differently from mail to be delivered to the Internet, which is processed differently from mail to be sent to another MTA within the mail system. Channels provide the mechanism for customizing the processing and routing required for each connection. In a default installation, the majority of messages go to a channels handling Internet, intranet, and local messages.

Specialized channels for specific situations can also be created. For example, suppose that a certain Internet domain processes mail very slowly causing mail addressed to this domain to clog up the MTA. A special channel could be created to provide special handling for messages addressed to the slow domain, thus relieving the system of this domain bottleneck.

The domain part of the address determines to what channel the message is enqueued. The mechanism for reading the domain and determining the appropriate channel is called the rewrite rules (see "[Rewrite Rules](#)").

Channels typically consist of a channel queue and a channel processing program called a *master program*. After the slave program delivers the message to the appropriate channel queue, the master program performs the desired processing and routing. Here is an example of a channel entry:

```
tcp_intranet smtp mx single_sys subdirs 20 noreverse maxjobs 7 SMTP_POOL \  
maytlssserver allowswitchchannel sasls witchchannel tcp_auth \  
tcp_intranet-daemon
```

The first word, in this case **tcp_intranet** is the channel name. The last word is called the channel tag. The words in between are called channel options (formerly called channel keywords) and specify how messages are to be processed. Hundreds of different options enable messages to be processed in many ways. See the discussion on channel options in the *Messaging Server Reference* for a complete description of channel options.

Message Delivery

After the message is processed, the master program sends the message to the next stop along the message's delivery path. This may be the intended recipient's mailbox, another MTA, or even a different channel. Forwarding to another channel is not shown in the picture, but is a common occurrence.

The local parts of addresses and received fields are typically 7-bit characters. If the MTA reads 8-bit characters in the in these fields, it replaces each 8-bit character with asterisks.

The Dispatcher

The Dispatcher is a multithreaded dispatching agent that permits multiple multithreaded server processes to share responsibility for SMTP connection services. When using the Dispatcher, it is possible to have several multithreaded SMTP server processes running concurrently, all handling connections to the same port. In addition, each server may have one or more active connections.

The Dispatcher acts as a central receiver for the TCP ports listed in its configuration. For each defined service, the Dispatcher may create one or more SMTP server processes to handle the connections after they are established.

In general, when the Dispatcher receives a connection for a defined TCP port, it checks its pool of available worker processes for the service on that port and chooses the best candidate for the new connection. If no suitable candidate is available and the configuration permits it, the Dispatcher may create a new worker process to handle this and subsequent connections. The Dispatcher may also create a new worker process in expectation of future incoming connections. There are several configuration options which may be used to tune the Dispatcher's control of its various services, and in particular, to control the number of worker processes and the number of connections each worker process handles.

For more information, see the discussion on the dispatcher configuration file in the *Messaging Server Reference*.

Creation and Expiration of Server Processes

Automatic housekeeping facilities within the Dispatcher control the creation of new and expiration of old or idle server processes. The basic options that control the Dispatcher's behavior are **service:name.min_procs** and **service:name.max_procs**. The **service:name.min_procs** option provides a guaranteed level of service by having a number of server processes ready and waiting for incoming connections. The **service:name.max_procs** option, on the other hand, sets an upper limit on how many server processes may be concurrently active for the given service.

It is possible that a currently running server process might not be able to accept any connections because it is already handling the maximum number of connections of which it is capable, or because the process has been scheduled for termination. The Dispatcher may create additional processes to assist with future connections.

The **min_conns** and **max_conns** options provide a mechanism to help you distribute the connections among your server processes. The **min_conns** option specifies the number of connections that flags a server process as "busy enough," while the **max_conns** option specifies the "busiest" that a server process can be.

In general, the Dispatcher creates a new server process when the current number of server processes is less than the value of the **service:name.min_procs** option or when all existing server processes are "busy enough" (the number of currently active connections each has is at least **min_conns**).

If a server process is killed unexpectedly, for example, by the UNIX system **kill** command, the Dispatcher still creates new server processes as new connections come in.

For information about configuring the Dispatcher, see the discussion on the dispatcher configuration file in the *Messaging Server Reference*.

To Start and Stop the Dispatcher

To start the Dispatcher, run the following command:

```
start-msg dispatcher
```

This command subsumes and makes obsolete any other *start-msg* command that was used previously to start up a component of the MTA that the Dispatcher has been configured to manage. Specifically, you should no longer use **imsimta start smtp**. An attempt to execute any of the obsoleted commands causes the MTA to issue a warning.

To shut down the Dispatcher, run the following command:

```
stop-msg dispatcher
```

What happens with the server processes when the Dispatcher is shut down depends upon the underlying TCP/IP package. If you modify your MTA configuration or options that apply to the Dispatcher, you must restart the Dispatcher so that the new configuration or options take effect.

To restart the Dispatcher, run the following command:

```
imsimta restart dispatcher
```

Restarting the Dispatcher has the effect of shutting down the currently running Dispatcher, then immediately starting a new one.

Rewrite Rules

Rewrite rules determine the following:

- How to rewrite the domain part of an address into its proper or desired format.
- To which channel the message should be enqueued after the address is rewritten.

Each rewrite rule consists of a *pattern* and a *template*. The pattern is a string to match against the domain part of an address. The template specifies the actions to take if the domain part matches the pattern. It consists of two things:

1. A set of instructions (that is, a string of control characters) specifying how the address should be rewritten.
2. The name of the channel to which the message shall be sent. After the address is rewritten, the message is enqueued to the destination channel for delivery to the intended recipient.

Here is an example of a rewrite rule:

```
example.org $U$D@tcp_example-daemon
```

example.org is the domain pattern. Any message with the address containing **example.org** will be rewritten as per the template instructions (**\$U\$D**). **\$U** specifies that the rewritten address use the same user name. **%** specifies that the rewritten address use the same domain separator. **\$D** specifies that the rewritten address use the same domain name that was matched in the pattern. **@tcp_example-daemon** specifies that the message with its rewritten address be sent to the channel called **tcp_example-daemon**. See Configuring Rewrite Rules in Unified Configuration for more details.

For more information about configuring rewrite rules, see the MTA configuration file in the *Messaging Server Reference* and "[Configuring Rewrite Rules](#)."

Channels

The channel is the fundamental MTA component that processes a message. A channel represents a connection with another computer system or group of systems. The actual hardware connection or software transport or both may vary widely from one channel to the next.

Channels perform a variety of functions, including:

- Transmitting messages to remote systems, deleting them from their queue after they are sent
- Accepting messages from remote systems, placing them in the appropriate channel queues
- Delivering messages to the local message store
- Delivering messages to programs for special processing

Messages are enqueued by channels on the way into the MTA and dequeued on the way out. Typically, a message enters by one channel and leaves by another. A channel might dequeue a message, process the message, or enqueue the message to another MTA channel.

This section consists of the following subsections:

- [Master and Slave Programs](#)
- [Channel Message Queues](#)
- [Channel Definitions](#)

Master and Slave Programs

Generally (but not always), a channel is associated with two programs: master and slave. The slave program accepts messages from another system and adds them to a channel's message queue. The master program transfers messages from the channel to another system.

For example, an SMTP channel has a master program that transmits messages and a slave program that receives messages. These are, respectively, the SMTP client and server.

The master channel program is typically responsible for outgoing connections where the MTA has initiated the operation. The master channel program:

- Runs in response to a local request for processing.
- Dequeues the message from the channel message queue.
- If the destination format is not the same format as the queued message, performs conversion of addresses, headers, and content, as necessary.
- Initiates network transport of the message.

The slave channel program typically accepts incoming connections where the MTA is responding to an external request. The slave channel program:


- Runs in response to an external event or upon local demand.

- Enqueues a message to a channel. The target channel is determined by passing envelope addresses through a rewrite rule.

For example, [Figure 5-3](#) shows two channel programs, Channel 1 and Channel 2. The slave program in Channel 1 receives a message from a remote system. It looks at the address, applies rewrite rules as necessary, then based on the rewritten address enqueues the message to the appropriate channel message queue.


The master program dequeues the message from the queue and initiates network transport of the message. Note that the master program can only dequeue messages from its own channel queue.

Figure 5-3 Master and Slave Program Interaction



Although a typical channel has both a master and a slave program, it is possible for a channel to contain only a slave program *or* a master program. For example, the **ims-ms** channel supplied with Messaging Server contains only a master program because this channel is responsible only for dequeuing messages to the local message store, as shown in [Figure 5-4](#).

Figure 5-4 ims-ms Channel



Channel Message Queues

All channels have an associated message queue. When a message enters the messaging system, a slave program determines to which message queue the message is enqueued. The enqueued messages are stored in message files in the channel queue directories. By default, these directories are stored at the following location: `MessagingServer_home/data/queue/channell*`. For more information, see the discussion on message queue sizing in the *Messaging Server Installation and Configuration Guide*.

Caution: Do not add any files or directories in the MTA queue directory. When using a separate file system for the MTA queue directories, create a subdirectory under that mount point and specify that subdirectory in the **SERVERROOT** environment variable. Sites may change it by using either a symbolic link or using it as a file system mount point. The default value is: **IMTA_ROOT:data/queue/**.

Channel Definitions

In Unified Configuration, use the **msconfig edit channels** and **msconfig edit rewrites** commands to view and edit the channel block. (See "[Configuring Rewrite Rules](#)" for information about setting up channels.)

A channel definition contains the name of the channel followed by an optional list of keywords that define the configuration of the channel, and a unique channel tag, which is used in rewrite rules to route messages to the channel. Channel definitions are separated by single blank lines. Comments, but no blank lines, may appear inside a channel definition. The following represents the channel format.

```
[blank line]
! sample channel definition
<Channel_Name> <keyword1> <keyword2>
<Channel_Tag>
[blank line]
```

Collectively, the channel definitions are referred to as the channel host table. An individual channel definition is called a channel block. In the following example, the channel host table contains three channel definitions or blocks.

```
! test.cnf - An example configuration file.
!
! Rewrite Rules
.
.
.

! BEGIN CHANNEL DEFINITIONS
! FIRST CHANNEL BLOCK
l
local-host

! SECOND CHANNEL BLOCK
a_channel defragment charset7 usascii
a-daemon

! THIRD CHANNEL BLOCK
b_channel noreverse notices 1 2 3
b-daemon
```

A typical channel entry looks something like the following, shown in legacy format as would display using the **msconfig edit channels** command:

```
tcp_intranet smtp mx single_sys subdirs 20 noreverse maxjobs 7 SMTP_POOL \
maytlsserver allowswitchchannel sasls witchchannel tcp_auth
tcp_intranet-daemon
```

The first word, in this case **tcp_intranet**, is the channel name. The last word, in this case **tcp_intranet-daemon**, is called the *channel tag*. The channel tag is the name used by rewrite rules to direct messages. The words in between the channel name and

channel tag are called channel options (formerly channel keywords) and specify how the message is to be processed. Hundreds of different keywords allow messages to be processed in many ways. A complete listing of channel keywords is listed and described in the discussion on channel options in the *Messaging Server Reference*.

The channel host table defines the channels Messaging Server can use and the names of the systems associated with each channel.

On UNIX systems, the first channel block in the file always describes the local channel, **l**. (An exception is a **defaults** channel, which can appear before the local channel.) The local channel is used to make routing decisions and for sending mail sent by UNIX mail tools.

You can also set global options for channels or set options for a specific channel. For more information on the option files and the TCP/IP (SMTP) channel option files, see the *Messaging Server Reference*. For details on configuring channels see the *Messaging Server Reference*. For more information about creating MTA channels, see "[About MTA Services](#)."

- [Master and Slave Programs](#)
- [Channel Message Queues](#)
- [Channel Definitions](#)

The MTA Directory Information

For each message that it processes, the MTA needs to access directory information about the users, groups, and domains that it supports. This information is stored in an LDAP directory service. The MTA directly accesses the LDAP directory. This is fully described in "[MTA Address Translation and Routing](#)."

The Job Controller

Each time a message is enqueued to a channel, the Job Controller ensures that there is a job running to deliver the message. This might involve starting a new job process, adding a thread, or simply noting that a job is already running. If a job cannot be started because the job limit for the channel or pool has been reached, the Job Controller waits until another job has exited. When the job limit is no longer exceeded, the Job Controller starts another job.

Channel jobs run inside processing pools within the Job Controller. A pool can be thought of a "place" where the channel jobs are run. The pool provides a computing area where a set of jobs can operate without vying for resources with jobs outside of the pool. For more information on pools, see the discussion on processing pools for channel execution jobs in the *Messaging Server Reference*.

Job limits for the channel are determined by the **channel:name.maxjobs** option. Job limits for the pool are determined by the **job_controller.job_pool:name.job_limit** option for the pool.

Messaging Server normally attempts to deliver all messages immediately. If a message cannot be delivered on the first attempt, however, the message is delayed for a period of time determined by the appropriate **backoff** option. As soon as the time specified in the **backoff** option has elapsed, the delayed message is available for delivery, and if necessary, a channel job is started to process the message.

The Job Controller's in-memory data structure of messages currently being processed and awaiting processing typically reflects the full set of message files stored on disk in the MTA queue area. However, if a backlog of message files on disk builds up enough

to exceed the Job Controller's in-memory data structure size limit, then the Job Controller tracks in memory only a subset of the total number of messages files on disk. The Job Controller processes only those messages it is tracking in memory. After a sufficient number of messages have been delivered to free enough in-memory storage, the Job Controller automatically refreshes its in-memory store by scanning the MTA queue area to update its list of messages. The Job Controller then begins processing the additional message files it just retrieved from disk. The Job Controller performs these scans of the MTA queue area automatically.

In previous versions of Messaging Server, the Job Controller read all the files in the queue directory in the order in which they are found. It now reads several channel queue directories at once. This makes for much more reasonable behavior on startup, restart, and after **max_cache_messages** has been exceeded. The number of directories to be read at once is controlled by the Job Controller option **rebuild_parallel_channels**. This can take any value between 1 and 100. The default is 12.

If your site routinely experiences heavy message backlogs, you might want to tune the Job Controller by using the **max_cache_messages** option. By increasing the **max_cache_messages** option value to allow Job Controller to use more memory, you can reduce the number of occasions when message backlogs overflow the Job Controller's in-memory cache. This reduces the overhead involved when the Job Controller must scan the MTA queue directory. Keep in mind, however, that when the Job Controller does need to rebuild the in-memory cache, the process will take longer because the cache is larger. Note also that because the Job Controller must scan the MTA queue directory every time it is started or restarted, large message backlogs mean that starts or restarts of the Job Controller will incur more overhead than starts or restarts when no such backlog exists.

You do not want to overwhelm the job controller by keeping information about huge numbers of messages in memory. For this reason, there has to be a and upper and lower limit. The number specified by **max_cache_messages** is the number of messages that the job controller will hold in memory. It will get this high if there are new messages delivered, for instance ones received by **tcp_smtp_server**. Beyond this number, messages are queued (put on disk), but not put into the job controller memory structure. The job controller notices this condition and when the number of messages in memory drops below half this maximum, it starts scanning the disk queues for more messages. It always looks for untried messages "ZZ..." files first, then previously tried messages.

In addition, the job controller limits the number of messages reclaimed from disk. It only reads from disk up to three-quarters of the **max_cache_messages** to allow for headroom for new messages (if messages are being reclaimed from disk, they have been delayed, which is an undesirable state).

Furthermore, you want to avoid cluttering up the memory structure with delayed messages (those that cannot be processed yet). When a message is delayed because it cannot be delivered immediately (a delivery attempt has failed if the number of messages the job controller knows about is greater than 5/8 of **max_cache_messages** and the number of delayed messages is greater than 3/8 of **max_cache_messages**) the message is forgotten until the next sweep of the on disk structures, which will be when the number of messages drops below 1/2 **max_cache_messages**.

The only obvious problems with having **max_cache_messages** too small is that the scheduling of jobs will become suboptimal. The scanning of the disk queues is also a bit simplistic. If you have huge numbers of messages backlogged in both the **tcp_local** and **ims_ms** queues, then the rebuild thread finds all the messages for one channel first, then the ones for the next channel. This can result in alarmed administrators

reporting that they've fixed one issue, but are only seeing only one specific channel dequeuing.

This is not a problem. There is a memory cost of approximately 140 bytes for each message. Having a message limit of 100000, you are limiting the job controller data structures to about 20 Megabytes (there are other data structures representing jobs, channels, destination hosts and so on). This is insignificant on a big server.

All the named objects in the job controller are tracked in a hash table. This is sized at the next power of 2 bigger than **max_cache_messages**, and is never re-sized. Each entry in this hash table is a pointer, so we are looking at a memory usage of four times **max_cache_messages** rounded up to a power of two. Being a hash table, this tends all to be in memory as the hash function is supposed to be random. This is another 0.5 Megabytes in the default case.

For information about pools and configuring the Job Controller, see "[Recompiling the MTA Configuration](#)" and the discussion on configuring message processing and delivery in the *Messaging Server Reference*.

To Start and Stop the Job Controller

To start the Job Controller, run the following command:

```
start-msg job_controller
```

To shut down the Job Controller, run the following command:

```
stop-msg job_controller
```

To restart the Job Controller, run the following command:

```
imsimta restart job_controller
```

Restarting the Job Controller has the effect of shutting down the currently running Job Controller, then immediately starting a new one.

On Demand Mail Relay

Support for On Demand Mail Relay as specified in RFC 2645 has been completed. This support piggybacks off existing support for SMTP **TURN** and doesn't involve any new options. Specifically, the **turn_in** channel option now enables both **TURN** and **ATRN**.

Note that the optional parameter to the **ATRN** command is only allowed if the **ODMR** channel is marked **single_sys** or **single**. This is so messages sent to a particular email domain can be reliably retrieved without getting messages sent to other domains.

It is strongly recommended that each administrative domain that requires **ATRN** service be configured as a separate channel rather than relying on **ATRN**'s domain selection capabilities. Specifically, the recommended configuration process to provide relay on demand for a new administrative domain foo is as follows:

1. Create a new channel for the domain foo:

```
tcp_foo mustsaslserver single_sys slave smtp turn_in  
tcp_auth-daemon
```

Or in **msconfig**:

```
set channel:tcp_foo.official_host_name tcp_foo-daemon  
set channel:tcp_foo.mustsaslserver  
set channel:tcp_foo.single_sys
```

```
set channel:tcp_foo.slave
set channel:tcp_foo.smtp
set channel:tcp_foo.turn_in
```

Note the presence of the **slave** channel option. This prevents the channel from trying to perform regular SMTP deliveries. This can be removed if such delivery attempts are desired.

Also note the presence of **single_sys**. **multiple** can be used instead and will be more efficient if multiple email domains are involved and the **ATRN** command is always used without an argument.

2. Create rewrite rules for all email domains associated with the foo administrative domain, for example:

```
foo.example.com    $U%D@tcp_foo-daemon
foo.example.org    $U%D@tcp_foo-daemon
```

Or in **msconfig**:

```
set rewrite.rule foo.example.com $U%D@tcp_foo-daemon
set rewrite.rule foo.example.org $U%D@tcp_foo-daemon
```

3. Create a regular email account for foo with whatever name and credentials are desired. The account should include the LDAP attribute **mailSubmitChannel** with the value **tcp_foo**.
4. (optional) Create a dispatcher configuration identical to the regular service on port 25 except that it listens on port 366, the On Demand Mail Relay port.
5. (optional) Add a rule to the **FROM_ACCESS** mapping to block all mail coming from the ODMR port unconditionally.
6. (optional) Add a rule to the **FROM_ACCESS** mapping to block all mail from the **tcp_foo** channel unconditionally. This will prevent the administrative account from being used to send authenticated mail.

At this point **ATRN** should be usable by connecting, authenticating as the newly created user, and issuing an **ATRN**.

Priority Message Handling

This chapter describes Messaging Server's support of the **MT-PRIORITY** SMTP extension defined in RFC 6710. **MT-PRIORITY** values are always in the range -9 to 9. Priority 0 is the default.

The support for this extension consists of the following parts:

- The **mtprioritiesallowed** and **mtprioritiesrequired** source channel options. Both accept either one- or two-integer arguments. Two-integer arguments specify the range. You can specify the arguments in any order. [Table 5-1](#) shows the channel options and their descriptions.

Table 5–1 MT-PRIORITY SMTP Extension Support

Channel Option	Description
mtprioritiesallowed <i>int1</i> [<i>int1</i>]	Specifies the range of MT-PRIORITY values that will be accepted. MT-PRIORITY values outside this range will be adjusted up or down so they fall within the allowed range. If a single argument is given it specifies the highest priority value that will be accepted. The default if this option is not specified is for the MT-PRIORITY extension not to be offered and for MT-PRIORITY options not to be accepted.
mtprioritiesrequired <i>int1</i> [<i>int2</i>]	Specifies the range of MT-PRIORITY that will be accepted for enqueue. If a single argument is given it specifies the lowest priority value that will be accepted. The message will be rejected if the specified MT-PRIORITY value or the default value of 0 falls outside the required range.

- An **INCLUDE_MTPRIORITY** MTA option. This is a bit-encoded option. The default value is 0. [Table 5–2](#) shows the bits, corresponding values, and descriptions of the **INCLUDE_MTPRIORITY** MTA option.

Table 5–2 INCLUDE_MTPRIORITY MTA Option

Bit	Value	Description
0	1	Appends the MT-PRIORITY and expected message size as separate fields to the FROM_ACCESS mapping probe immediately after any INCLUDE_SPARES values.
1	2	Appends the MT-PRIORITY and expected message size as separate fields to the FORWARD mapping probe immediately after the conversion tag field.
2	4	Appends the MT-PRIORITY and expected message size as separate fields to the ORIG_SEND_ACCESS mapping probe immediately after any INCLUDE_SPARES values.
3	8	Appends the MT-PRIORITY and expected message size as separate fields to the SEND_ACCESS mapping probe immediately after any INCLUDE_SPARES values.
4	16	Appends the MT-PRIORITY and expected message size as separate fields to the ORIG_MAIL_ACCESS mapping probe immediately after any INCLUDE_SPARES values.
5	32	Appends the MT-PRIORITY and expected message size as separate fields to the MAIL_ACCESS mapping probe immediately after any INCLUDE_SPARES values.
6	64	Appends the MT-PRIORITY and actual message size values in the form: ;MT-PRIORITY=<value>;BLOCKS=<value> to the conversion mapping probe immediately after any tag= clause.
7	128	Append the MT-PRIORITY and expected message size as separate fields to any domain catchall mapping probe immediately after the conversion tag.

The expected message size is the size of the queued message entry for internal channels. It is the value given by the SMTP **SIZE** extension for incoming SMTP channels. The size is given in MTA blocks.

- A **LOG_MTPRIORITY** MTA option. This is a bit-encoded option. The default is 0. [Table 5–3](#) shows the bits, corresponding values, and descriptions of the **LOG_MTPRIORITY** MTA option. An **mp** element is used in the new XML log format to log the priority information. For more information about the XML log format, see ["Managing Logging."](#)

Table 5–3 LOG_MTPRIORITY MTA Option

Bit	Value	Description
0	1	Enables logging of the MT-PRIORITY associated with each transaction. The MT-PRIORITY appears immediately after the message sensitivity and before the header-based priority in each log entry.
1	2	Message transfer priority appears in the LOG_ACTION mapping table probe, immediately after the sensitivity field and before the header-based priority field.

- Rewrite rule metacharacters that test both the current **MT-PRIORITY** value and the expected message size. For each metacharacter, *n* is a signed integer value. Note that the sign is required even when *n* is positive. The expected message size is the size of the queued message entry for internal channels. It is the value given by the **SMTP SIZE** extension for incoming **SMTP** channels. The size is given in **MTA** blocks.

Table 5–4 shows the rewrite rule metacharacters and their descriptions.

Table 5–4 Rewrite Rule Metacharacters

Metacharacter	Description
\$n<P	Rule succeeds only if <i>n</i> is less than the current MT-PRIORITY .
\$n<=P	Rule succeeds only if <i>n</i> is less than or equal to the current MT-PRIORITY .
\$n>P	Rule succeeds only if <i>n</i> is greater than the current MT-PRIORITY .
\$n>=P	Rule succeeds only if <i>n</i> is greater than or equal to the current MT-PRIORITY .
\$n=P	Rule succeeds only if <i>n</i> is equal to the current MT-PRIORITY .
\$n<>P	Rule succeeds only if <i>n</i> is not equal to the current MT-PRIORITY .
\$n<B	Rule succeeds only if <i>n</i> is less than the expected message size.
\$n<=B	Rule succeeds only if <i>n</i> is less than or equal to the expected message size.
\$n>B	Rule succeeds only if <i>n</i> is greater than the expected message size.
\$n>=B	Rule succeeds only if <i>n</i> is greater than or equal to the expected message size.
\$n=B	Rule succeeds only if <i>n</i> is equal to the expected message size.
\$n<>B	Rule succeeds only if <i>n</i> is not equal to the expected message size.

- A **-mtpriority** switch added to **imsimta test -rewrite**, **imsimta calc**, and **imsimta test -expression** utilities. A single integer argument is required specifying the initial **MT-PRIORITY** value.
- A Sieve environment item, **vnd.oracle.mt-priority**. This item returns the current **MT-PRIORITY** value as a string.
- A nonstandard Sieve action, **setmtpriority**. This action accepts a single integer or string argument and sets the current **MT-PRIORITY** to the argument value. This action is only allowed in system-level Sieves and the argument must be in the -9 to 9 range of valid **MT-PRIORITY** values.
- A bit defined in the **MESSAGE_SAVE_COPY_FLAGS** MTA option. Bit 3 (value 8), if set, causes the **MT-PRIORITY** value for the current message to be included, delimited by vertical bars, immediately after the conversion tag.
- An **MTPRIORITY_POLICY** MTA option. This option is used to specify the priority handling policy the MTA has been configured to support. This name is

announced in the SMTP EHLO response on any channel where the **MT-PRIORITY** extension is enabled. The default is the empty string, meaning that no policy is announced.

MTA Address Translation and Routing

This chapter describes the flow of data through the MTA using direct LDAP data access.

The Direct LDAP Algorithm and Implementation

The following sections describe direct LDAP processing:

- Domain Locality Determination
- Alias Expansion of Local Addresses
- Processing the LDAP Result
- To Modify Group Membership Attribute Syntax

Domain Locality Determination

Starting with an address of the form *user@domain*, the address translation and routing process first checks to see if *domain* is local.

Topics in this section:

- Rewrite Rule Machinery
- Domain Map Determination of Domain Locality
- Caching of Domain Locality Information
- Error Handling
- Pattern for Domain Check Rewrite Rule
- Putting It All Together

Rewrite Rule Machinery

The MTA rewrite rule machinery checks a given string to see if it is a domain that needs to be locally handled. This is activated by a **\$V** or **\$Z** metacharacter. These metacharacters are syntactically similar to the **\$N**, **\$M**, **\$Q**, and **\$C** metacharacters, that is, they are followed by a pattern string. In the case of **\$N**, **\$M**, **\$Q**, and **\$C**, the pattern is matched against either the source or destination channel. In the case of **\$V** and **\$Z**, the pattern is a domain and the check is to see if it is local. **\$V** causes a rule failure for a non-local domain and **\$Z** causes a rule failure for a local domain.

The handling of these metacharacters is implemented as the following procedure:

1. Oracle Communications Messaging Server checks if the current domain matches a valid domain entry in the directory. Go to Step 3 if no entry exists.

2. If the domain has an entry in the directory, the attribute specified by the **ldap_domain_attr_routing_hosts** MTA option (default **mailRoutingHosts**) is retrieved from the domain entry. If this attribute is present, it lists the set of hosts able to handle users in this domain. This list is compared against the host specified by the **hostname msconfig** option and the list of hosts specified by the **mta.ldap_host_alias_list msconfig** option. These options can be overridden by the **ldap_local_host** and **ldap_host_alias_list** MTA options, respectively. If there is a match or the attribute is not present on the domain, the domain is local. If no match occurs, the domain is non-local.

The handling of domains considered to be non-local because of the **mailRoutingHosts** attribute depends on the setting of the **route_to_routing_host** MTA option. If the option is set to 0 (the default), the address is simply treated as non-local and MTA rewrite rules are used to determine routing. If the option is set to 1, a source route consisting of the first value listed in the **ldap_domain_attr_routing_hosts** MTA option is prepended to the address.

3. If no domain entry can be found, remove a component from the left-hand side of the domain and go to Step 1. If no components remain continue to Step 4.

A consequence of this backtracking up the domain tree is that if **example.com** is recognized as local, any subdomain of **example.com** is recognized as local. Situations might arise where this is undesirable, so an MTA option, **domain_uplevel**, is provided to control this behavior. Specifically, bit 0 (value = 1) of **domain_uplevel**, if clear, disables retries with domain components removed. The default value of **domain_uplevel** is 0.

4. Vanity domain checking now needs to be performed. Vanity domains do not have domain entries, rather, they are specified by attaching special domain attributes to one or more user entries. The vanity domain check is done by instituting an LDAP search by using the LDAP URL specified by the **domain_match_url** MTA option. The value of this option should be set to:

```
ldap:///${B?msgVanityDomain?sub?(msgVanityDomain=${D)}
```

\$B substitutes the value of the **ugldapbasedn msconfig** option. This is the base of the user tree in the directory. The **ldap_user_root** MTA option can be used to override the value of this **msconfig** option specifically for the MTA.

The actual return value from this search does not matter. What matters is if there is a value to return. If there is a return value the domain is considered to be local, if not it is considered to be non-local.

Domain Map Determination of Domain Locality

The following steps are performed to find valid domain entries in the directory. These steps are schema-level specific. For Oracle LDAP Schema 1, the steps are the following:

1. Convert the domain to a base DN in the domain tree.

This is done by converting the domain into a series of dc components and then adding a domain root suffix. The default suffix is obtained from the **dcroot msconfig** option. The default suffix is **o=internet**. So a domain of the form **a.b.c.d** would typically be converted into **dc=a,dc=b,dc=c,dc=d,o=internet**. The **dcroot** option can be overridden by setting the **ldap_domain_root** MTA option.

2. Look for an entry with the base DN found in Step 1 and an object class of either **inetDomain** or **inetDomainAlias**.

The search filter used for this purpose can be overridden by setting the `ldap_domain_filter_schema1` MTA option, which defaults to `(!(objectclass=inetDomain)(objectclass=inetdomainalias))`.

3. Exit with a failure if nothing is found.
4. If the object class of the entry found is `inetDomain`, check to make sure the entry has an `inetDomainBaseDn` attribute associated with the domain entry.

If it is present, it is saved for use in subsequent searches for user entries and processing terminates. If it is not present, the entry is assumed to be a domain alias and processing continues with Step 5. The MTA option `ldap_domain_attr_basedn` can be used to override the use of `inetDomainBaseDN`.

5. The entry must be a domain alias. Look up the new entry referenced by the `aliasedObjectName` attribute and return to Step 4.

Processing terminates with a failure if the no `aliasedObjectName` attribute is present. An alternative to the use of `aliasedObjectName` attribute can be specified with the MTA option `ldap_domain_attr_alias`. Processing can return to Step 4 at most once. Domain aliases pointing at domain aliases are not allowed.

In Schema 2, the action taken is much simpler: The directory is searched for an entry with the object class `sunManagedOrganization`, where the domain appears as a value of either the `sunPreferredDomain` or `associatedDomain` attribute. If need be, the use of the `sunPreferredDomain` and `associatedDomain` attributes for this purpose can be overridden with the respective MTA options `ldap_attr_domain1_schema2` and `ldap_attr_domain2_schema2`. The search is done under the root specified by the `dcrootmsconfig` option. The `dcroot` option can be overridden by setting the `ldap_domain_root` MTA option. Additionally, domain entries in Schema 2 are not required to have `inetDomainBaseDn` attributes. If they do not, the base of the user tree is assumed to be the domain entry itself.

Two MTA options support more efficient domain lookups from user base domain names. They are `ldap_basedn_filter_schema1`, which is a string specifying a filter used to identify Schema 1 domains when performing user base domain name searches. The default is the value of `ldap_domain_filter_schema1` if that MTA option is specified. If neither option is specified, the default is `(objectclass=inetDomain)`. `ldap_domain_filter_schema2` is a string specifying additional filter elements used to identify Schema 2 domains when performing user base domain name searches. The default is the value of `ldap_domain_filter_schema2`, if that MTA option is specified. If neither option is specified, the default is an empty string.

Caching of Domain Locality Information

Due to the frequency with which domain rewrite operations are performed and the expense of the directory queries (especially the vanity domain check) both negative and positive indications about domains need to be cached. This is implemented with an in-memory open-chained dynamically-expanded hash table. The maximum size of the cache is set by the `domain_match_cache_size` MTA option (default 100000) and the timeout for entries in the cache is set by the `domain_match_cache_timeout` MTA option (default 600 seconds).

Error Handling

Temporary server failures during this process have to be handled carefully, because when they occur, it is impossible to know if a given domain is local. Two outcomes are possible in such a case:

1. Return a temporary (4_xx_) error to the client telling it to try the address again later.
2. Accept the address but queue it to the reprocessing channel so it can be retried locally later.

Neither of these options is appropriate in all cases. For example, outcome 1 is appropriate when talking to a remote SMTP relay. But outcome 2 is appropriate when dealing with an SMTP submission from a local user.

While it would be possible in theory to handle temporary failures by using multiple rules with the same pattern, the overhead of repeating such queries, even with a cache in place, is unacceptable. For these reasons, the simple success/fall-through-to-the-next-rule matching model of domain rewriting is inadequate. Instead, a special template, specified by the MTA option **domain_failure**, is used in the event of a domain lookup failure. When a **\$V** operation fails, this template replaces the remainder of the current rewrite rule template being processed.

Pattern for Domain Check Rewrite Rule

This domain check needs to be performed before other rewrite rules have a chance to operate. This ordering is insured by using the special **\$*** on the left-hand side in the rule. The **\$*** pattern is checked prior to any other rules.

Putting It All Together

Taking all the machinery described so far into account, the configuration needs the following rewrite rule:

```
$*      $A$E$F$U%$H$V$H@&/IMTA_HOST/
```

and the **domain_failure** MTA option needs the following value:

```
reprocess-daemon$Mtcp_local$1M$1~-error$4000000?Temporary lookup failure
```

In this rewrite rule, **IMTA_HOST** is the host name associated with the local channel. The value of the **domain_failure** option shown here is the default value so it does not need to appear in the configuration under normal circumstances.

The ordering here is especially tricky. The MTA checks **\$V** after the address is rebuilt but before the route is added. This allows the MTA to change the route in the event of a temporary lookup failure. Pending channel match checks are applied any time the insertion point changes, so the **@** after the second **\$H** invokes the check. If the check succeeds, the remainder of the template applies and rewrite processing concludes. If the check fails, the rewrite fails and rewriting continues with the next applicable rewrite rule. If the check cannot be performed due to a temporary failure, template processing continues with the value specified by the **DOMAIN_FAILURE** MTA option. The value of this template first sets the routing host to **reprocess-daemon**. Then the template checks to see whether or not the MTA is dealing with a reprocessing channel of some sort or **tcp_local**. If the MTA is dealing with such a channel, the rule continues, making the routing host illegal and specifying a temporary failure as the outcome. If the MTA is not dealing with such a channel, the rule is truncated and successfully terminated, thereby rewriting the address to the reprocess channel.

Alias Expansion of Local Addresses

Once an address has been determined to be associated with the local channel it automatically undergoes alias expansion. The alias expansion process examines a number of sources of information, including:

1. The alias file (part of the compiled configuration)
2. The alias database
3. Alias URLs

The exact alias sources that are checked and the order in which they are checked depends on the setting of the **alias_magic** MTA option. For direct LDAP, set the option to 8764. This means that the URL specified by the **alias_url0** MTA option is checked first, then the URL specified by the **alias_url1** MTA option, then the URL specific by the **alias_url2** MTA option, and finally, the alias file. The alias database is not checked when this setting is active.

The following sections further describe alias expansion:

- [Alias Checking with LDAP URLs](#)
- [The \\$V Metacharacter](#)
- [Calling a Mapping From a URL](#)
- [The \\$R Metacharacter](#)
- [Determining the Attributes to Fetch](#)
- [Handling LDAP Errors](#)
- [Sanity Checks on the LDAP Result](#)
- [Support for Vanity Domains](#)
- [Support for Catchall Addresses](#)

Alias Checking with LDAP URLs

Checking of aliases in LDAP is implemented by specifying two special LDAP URLs as alias URLs. The first of these handles regular users and groups; vanity domains are handled by subsequent alias URLs. The first URL is specified as **alias_url0**:

```
alias_url0=ldap:/// $V?*?sub?$R
```

The \$V Metacharacter

Metacharacter expansion occurs prior to URL lookup. The two metacharacters used in the **alias_url0** value are **\$V** and **\$R**.

The **\$V** metacharacter converts the domain part of the address into a base DN. This is similar to the initial steps performed by the **\$V** rewrite rule metacharacter described previously in the section entitled "[Rewrite Rule Machinery](#)." **\$V** processing consists of the following steps:

1. Get the base DN for user entries in the current domain.
2. Get the canonical domain associated with the current domain. In Oracle LDAP Schema 1, the canonical domain name is given by the **inetCanonicalDomainName** attribute of the domain entry if the attribute is present. If the attribute is absent the canonical name is the name constructed in the obvious way from the DN of the actual domain entry. This differs from the current domain when the current domain is an alias. The name attribute used to store the canonical name may be overridden with the **ldap_domain_attr_canonical** MTA option.

In Schema 2, the canonical name is simply the value of the **SunPreferredDomain** attribute.

A utility for verifying canonical domain settings for domains with overlapping user entries is provided. See "[imsimta test -domain](#)."

3. If the base DN exists, substitute it into the URL in place of the **\$V**.
4. Any applicable hosted domain for this entry is now determined. This is done by comparing either the canonical domain (if bit 2 (value = 4) of **domain_uplevel** is clear) or the current domain (if bit 2 (value = 4) of **domain_uplevel** is set) with the **base.defaultdomain msconfig** option. If they do not match, the entry is a member of a hosted domain. The **base.defaultdomain** option can be overridden by setting the **ldap_default_domain** MTA option.
5. If the base DN determination fails, remove a component from the left-hand side of the domain and go to Step 1. The substitution fails if no components remain.

\$V also accepts an optional numeric argument. If it is set to 1 (for example, **\$1V**), a failure to resolve the domain in the domain tree is ignored and the base of the user tree specified by the **ugldapbasedn msconfig** option is returned.

If the attempt to retrieve the domain's base DN succeeds, the MTA also retrieves several useful domain attributes, which are needed later. The names of the attributes retrieved are set by the following MTA options:

- **ldap_domain_attr_uid_separator** (default **domainUidSeparator**)
- **ldap_domain_attr_smarthost** (default **mailRoutingSmartHost**)
- **ldap_domain_attr_catchall_address** (default **mailDomainCatchallAddress**)
- **ldap_domain_attr_catchall_mapping** (no default value)
- **ldap_domain_attr_blocklimit** (default **mailDomainMsgMaxBlocks**)
- **ldap_domain_attr_report_address** (default **mailDomainReportAddress**)
- **ldap_domain_attr_status** (default **inetDomainStatus**)
- **ldap_domain_attr_mail_status** (default **mailDomainStatus**)
- **ldap_domain_attr_conversion_tag** (default **mailDomainConversionTag**)
- **ldap_domain_attr_filter** (default **mailDomainSieveRuleSource**)
- **ldap_domain_attr_disk_quota** (no default)
- **ldap_domain_attr_message_quota** (no default)
- **ldap_domain_attr_autoreply_timeout** (no default)
- **ldap_domain_attr_nosolicit** (no default)
- **ldap_domain_attr_optin** (no default)
- **ldap_domain_attr_recipientlimit** (no default)
- **ldap_domain_attr_recipientcutoff** (no default)
- **ldap_domain_attr_sourceblocklimit** (no default)

Calling a Mapping From a URL

There might be unusual cases where the mapping from domain to base DN is done some other way. To accommodate such setups, the URL resolution process has the ability to call an MTA mapping. This is done with a metacharacter sequence of the general form:

\$|/mapping-name/mapping-argument|

The double quotation (") initiates and terminates the callout. The character immediately following the \$ is the separator between the mapping name and argument; a character should be chosen that does not collide with the expected character values used in either the mapping name or argument.

The \$R Metacharacter

The \$R metacharacter provides an appropriate filter for the URL. The intent is to produce a filter that searches all the attributes that might contain an email address for a particular user or group. The list of attributes to search comes from the **msconfig** option **mta.mailaliases.ldap_mail_aliases**. If this option is not set, the **mta.ldap_schematag msconfig** option is examined, and depending on its value, an appropriate set of default attributes is chosen as follows:

sims401 mail, rfc822mailalias

nms41 mail, mailAlternateAddress

ims50 mail, mailAlternateAddress, mailEquivalentAddress

The value of **mta.ldap_schematag** can be a comma-separated list. If more than one schema is supported, the combined list of attributes with duplicates eliminated is used. The **ldap_schematag** MTA option can be used to override the setting of **mta.ldap_schematag** specifically for the MTA.

Additionally, the filter searches not only for the address that was originally supplied, but also for an address with the same local part but the domain that was actually found in the domain tree, which was saved in the second step in "[The \\$V Metacharacter](#)." The iterative nature of the domain tree lookup means the two addresses may be different. This additional check is controlled by bit 1 (value = 2) of the **domain_uplevel** MTA option. Setting the bit enables the additional address check. The default value of **domain_uplevel** is 0.

For example, suppose that the domain **example.com** appears in the domain tree. Assuming Oracle LDAP Schema 1 is in force, a lookup of the address:

u@host1.example.com

The filter that results from the expansion of \$R and an **ims50 schematag** looks like:

```
( | (mail=u@example.com)
  (mail=u@host1.example.com)
  (mailAlternateAddress=u@example.com)
  (mailAlternateAddress=u@host1.example.com)
  (mailEquivalentAddress=u@example.com)
  (mailEquivalentAddress=u@host1.example.com) )
```

If, on the other hand, **domain_uplevel** was set to 1 rather than 3, the filter would be:

```
( | (mail=u@host1.example.com)
  (mailAlternateAddress=u@host1.example.com)
  (mailEquivalentAddress=u@host1.example.com) )
```

Determining the Attributes to Fetch

If the URL specifies an asterisk (*) for the list of attributes to return, the asterisk is replaced with the list of attributes the MTA is able to use. This list is dynamically generated from the various MTA option settings that specify the options the MTA consumes.

Handling LDAP Errors

At this point the resulting URL is used to perform an LDAP search. If an LDAP error of some kind occurs, processing terminates with a temporary failure indication (4xx error in SMTP). If the LDAP operation succeeds but fails to produce a result, the catchall address attribute for the domain retrieved from the `ldap_domain_attr_catchall_address` MTA option is checked. If it is set, its value replaces the current address.

If no catchall address attribute is set, the smarthost attribute for the domain retrieved from the `ldap_domain_attr_smarthost` MTA option is checked. If it is set, an address of the following form is created and alias processing terminates successfully with this result.

```
@<smarthost>:<user>@<domain>
```

Additionally, the conversion tag for the domain obtained from the `ldap_domain_attr_conversion_tag` MTA option is, if present, attached to the address so that conversions can be done prior to forwarding to the smarthost. If no catchall address or smarthost exists for the domain, processing of this alias URL terminates unsuccessfully.

Sanity Checks on the LDAP Result

After the LDAP search has returned a result, it is checked to verify that there is only one entry in it. If there are more than one, each entry is checked to see if it has the right object class for a user or a group, a non-deleted status, and for users, a UID. Entries that do not pass this check are ignored. If the list of multiple entries is reduced to one by this check, processing proceeds. If not, a duplicate or ambiguous directory error is returned.

Support for Vanity Domains

The `alias_url0` check is for a conventional user or a user in a hosted domain. If this fails a vanity domain check is also made. This is done with the following alias URL:

```
alias_url1=ldap:/// $B?*?sub? (&(msgVanityDomain=$D) $R)
```

Support for Catchall Addresses

Finally, a check for a catchall address of the form `@host` needs to be made in the `mailAlternateAddress` attribute. This form of wildcarding is allowed in both hosted and vanity domains, so the proper alias URL for it is:

```
alias_url2=ldap:/// $1V?*?sub? (mailAlternateAddress=@$D)
```

Note: The `+*` subaddress substitution mechanism has always worked with catch-all addresses in direct LDAP mode, but the string that was substituted was the subaddress only, not the entire local part. This has been changed so that the entire local part of the original address will be plugged into the catch-all address as a subaddress when this construct is used.

For example, given an address of the form `foo+bar@domain.com`, no local user `foo` in the `domain.com` domain, and a catch-all address for `domain.com` of `bletch+*@example.com`, the resulting address will now be `bletch+foo+bar@example.com`. It used to be `bletch+bar@example.com`.

Processing the LDAP Result

LDAP alias result processing is done in a number of order-dependent stages. These stages are described in the following sections:

- Object Class Check
- Entry Status Checks
- UID Check
- Message Capture
- Seeding the Reversal Cache
- Mail Host and Routing Address
- Miscellaneous Attribute Support
- Delivery Options Processing
- Additional Metacharacters for Use in Delivery Options
- Delivery Option Defaults
- Start and End Date Checks
- Optin and Presence Attributes
- Sieve Filter Handling
- Deferred Processing Control
- Group Expansion Attributes

Object Class Check

If the alias search succeeds, the object class of the entry is checked to make sure it contains an appropriate set of object classes for a user or a group. The possible sets of required object classes for users and groups is normally determined by what schemata are active. This is determined by the **mta.ldap_schematag** option.

Table 6–1 shows the user and group object classes that result from various **schematag** values.

Table 6–1 Object Classes Resulting from Various schematag Values

schematag	User Object Classes	Group Object Classes
sims40	inetMailRouting+inetmailuser	inetMailRouting+inetmailgroup
nms41	mailRecipient + nsMessagingServerUser	mailGroup
ims50	inetLocalMailRecipient+inetmailuser	inetLocalMailRecipient + inetmailgroup

The information in this table, like the rest of schema tag handling, is hard coded. However, there are also two MTA options, **ldap_user_object_classes** and **ldap_group_object_classes**, which can be set to specify different sets of object classes for users and groups respectively.

For example, a schema tag setting of **ims50,nms41** would be equivalent to the following option settings:

```
ldap_user_object_classes=inetLocalMailRecipient+inetmailuser,
mailRecipient+nsMessagingServerUser
```

```
ldap_group_object_classes=inetLocalMailRecipient+inetmailgroup, mailGroup
```

The LDAP result is simply ignored if it does not have a correct set of object classes appropriate for a user or a group. The MTA also determines if it is dealing with a user or a group and saves this information. This saved information will be used repeatedly later.

Note that the object class settings described here are also used to construct an actual LDAP search filter that can be used to check to see that an entry has the right object classes for a user or a group. This filter is accessible through the **\$K** metacharacter. It is also stored internally in the MTA's configuration for use by channel programs, as the **LDAP_UG_FILTER** option when the command **imsimta cnbuild -option** is used. This option is only written to the file. The MTA never reads it from the option file.

Entry Status Checks

Next the entry's status is checked. There are two status attributes, one for the entry in general and another specifically for mail service.

Table 6–2 shows the general and mail-specific user or group attributes in the schema tag entry to check against depending on what schema tag are in effect:

Table 6–2 Attributes to Check Against

schematag	Type	General	Mail-specific
sims40	users	inetsubscriberstatus	mailuserstatus
sims40	groups	none	inetmailgroupstatus
nms41	users	none	mailuserstatus
nms41	groups	none	none
Messaging Server 5.0	users	inetuserstatus	mailuserstatus
Messaging Server 5.0	groups	none	inetmailgroupstatus

If necessary the **ldap_user_status** and **ldap_group_status** MTA options can be used to select alternate general status attributes for users and groups respectively. The mail-specific user and group status attributes are controlled by the **ldap_user_mail_status** and **ldap_group_mail_status** MTA options.

Another factor that plays into this are the statuses for the domain itself, (**ldap_domain_attr_status** and **ldap_domain_attr_mail_status**). All in all there are four status attributes. They are combined by considering them in the following order:

1. Domain status
2. Domain mail status
3. User or group status
4. Mail user or group status

The first of these that specifies something other than "active" takes precedence over all the others. The other permissible status values are "inactive", "deleted", "removed", "disabled", "hold", and "overquota". "Hold", "disabled," and "removed" statuses may only be specified for mail domains, mail users, or mail groups. "Overquota" status can only be specified as a mail domain or mail user status.

All statuses default to "active" if a particular status attribute is not present. Unknown status values are interpreted as "inactive."

When the four statuses are combined, the following statuses for a user or group are possible: "active", "inactive", "deleted", "removed", "disabled", "hold", and "overquota."

Active status causes alias processing to continue. Inactive or overquota status results in immediate rejection of the address with a 4xx (temporary) error. Deleted, removed, and disabled statuses results in immediate rejection of the address with a 5xx (permanent) error. Hold status is treated as active as far as status handling is concerned but it sets an internal flag so that when delivery options are considered later on, any options that are there are overridden with an option list containing a single "hold" entry.

UID Check

The next step is to consider the entry's UID. UIDs are used for a variety of purposes and must be part of all user entries and may be included in group entries. A user entry without a UID is ignored and processing of this alias URL terminates unsuccessfully. UIDs for entries in a hosted domain can consist of the real UID, a separator character, and then a domain. The MTA only wants the real UID so the rest is removed if present using the domain separator character obtained with the **ldap_domain_attr_uid_separator** MTA option.

In the unlikely event that an attribute other than **uid** is used to store UIDs, the **ldap_uid** MTA option can be used to force use of a different attribute.

Message Capture

Next the LDAP attribute used to specify one or more message capture addresses is checked. The attribute used for this purpose must be specified with the **ldap_capture** MTA option. There is no default. Values of this attribute are treated as addresses and a special "capture" notification is generated and sent to these address containing the current message as an attachment. Additionally, the capture addresses are used to seed the address reversal cache in the likely event the address will subsequently appear as an envelope from: address.

Seeding the Reversal Cache

Next the primary address and any aliases attached to the user entry are considered. This information is used to seed the address reversal cache. It plays no part in the current address translation process. First, the primary address, personal name, recipient limit, recipient cutoff, and source block limit attributes are considered. The primary address is normally stored in the "**mail**" attribute; another attribute can be specified by setting the **ldap_primary_address** MTA option appropriately. (The primary address reverses to itself, of course.) There is no default attribute for any of the other attributes. If you want to use them, you must specify them with the **ldap_personal_name**, (see "[Vacation Autoreply Attributes](#)") **ldap_recipientlimit**, **ldap_recipientcutoff** (see the discussion on maximum allowed recipients or bad commands in the *Messaging Server Reference*), and **ldap_sourceblocklimit** (see the discussion on message size limits in the *Messaging Server Reference*) MTA options. The corresponding domain-level recipient limit, recipient cutoff, and source block limit attributes are also considered at this point. User-level settings completely override any domain-level setting.

Next, any secondary addresses are considered and a cache entry is made for each one. There are two sorts of secondary addresses: Those that undergo address reversal and those that do not. Both must be considered in order to properly seed the address reversal cache because of the need to check for message capture requests in all cases.

Secondary addresses that undergo reversal are normally stored in the **mailAlternateAddress** attribute. Another attribute can be specified by setting the **ldap_alias_addresses** MTA option. Secondary addresses that do not undergo reversal

are normally stored in the **mailEquivalentAddress** attribute. Another attribute can be specified with the **ldap_equivalence_addresses** MTA option.

Mail Host and Routing Address

It is now time to consider the **mailhost** and **mailRoutingAddress** attributes. The actual attributes considered can be overridden with the **ldap_mailhost** and **ldap_routing_address** MTA options, respectively. These attributes work together to determine whether or not the address should be acted on at this time or forwarded to another system.

The first step is to decide whether or not **mailhost** is meaningful for this entry. A preliminary check of the delivery options active for the entry is done to see if the entry is mailhost-specific. If it is not, **mailhost** checking is omitted. See the description of "Delivery Options Processing", and the # flag in particular, to understand how this check is done.

In the case of a user entry, the **mailhost** attribute must identify the local system in order to be acted on. The **mailhost** attribute is compared with the value of the **hostname msconfig** option and against the list of values specified by the **mta.ldap_host_alias_list msconfig** option. The **mailhost** attribute is deemed to identify the local host if any of these match.

A successful match means that the alias can be acted on locally and alias processing continues. An unsuccessful match means that the message needs to be forward to the mailhost to be acted on. A new address of the form:

```
@<mailhost>:<user>@<domain>
```

is constructed and becomes the result of the alias expansion operation.

The handling of a missing **mailhost** attribute is different depending on whether the entry is a user or a group. In the case of a user, a mailhost is essential, so if no **mailhost** attribute is present a new address of the form

```
@<smarthost>:<user>@<domain>
```

is constructed using the smart host for the domain determined by the **ldap_domain_attr_smarthost** MTA option. An error is reported if no smart host exists for the domain.

Groups, on the other hand, do not require a mailhost, so a missing mailhost is interpreted as meaning that the group can be expanded anywhere. So alias processing continues.

The **mailRoutingAddress** attribute adds one final wrinkle. Its presence causes processing to terminate with the **mailRoutingAddress** as the result.

Miscellaneous Attribute Support

Next the **mailMsgMaxBlocks** attribute is considered. First it is minimized with the domain block limit returned from the **ldap_domain_attr_blocklimit** MTA option. If the size of the current message is known to exceed the limit, alias processing terminates with a size-exceeded error. If the size is not known or does not exceed the limit, the limit is nevertheless stored and will be rechecked when the message itself is checked later. The use of **mailMsgMaxBlocks** can be overridden with the **ldap_blocklimit** MTA option.

Next a number of attributes are accessed and saved. Eventually these will be written into the queue file entry for use by the **ims_master** channel program, which will then

use them to update the store's user information cache. If the attributes are not found for individual users, domain-level attributes can be used to set defaults.

This step is skipped if the LDAP entry is for a group rather than a user or if the LDAP entry came from the alias cache and not from the LDAP directory. The logic behind the latter criteria is that frequent updates of this information are unnecessary and using the alias cache offers a reasonable criteria for when updates should be done. The names of the attributes retrieved are set by various MTA options.

Table 6–3 shows the MTA options which set the retrieved disk quota and message quota attributes.

Table 6–3 MTA Options Which Set the Retrieved Disk Quota and Message Quota Attributes

MTA Option	Attribute
ldap_disk_quota	mailQuota
ldap_message_quota	mailMsgQuota

Next a number of attributes are stored for possible use in conjunction with metacharacter substitutions later.

Table 6–4 shows the MTA options, the default attribute, and metacharacters.

Table 6–4 MTA Options, Default Attributes, and Metacharacters

MTA Option	Default Attribute	Metacharacters
ldap_program_info	mailProgramDeliveryInfo	\$P
ldap_delivery_file	mailDeliveryFileURL	\$F
ldap_spare_1	no default	\$1E \$1G \$E
ldap_spare_2	no default	\$2E \$2G \$G
ldap_spare_3	no default	\$3E \$3G
ldap_spare_4	no default	\$4E \$4G
ldap_spare_5	no default	\$5E \$5G

Spare slots for additional attributes are included so that you can use them to build customized address expansion facilities.

Next any values associated with the **mailconversiontag** attribute are added to the current set of conversion tags. The name of this attribute can be changed with the **ldap_conversion_tag** MTA option. If any values were associated with the domain's **mailDomainConversionTag** attribute, they are attached as well.

Delivery Options Processing

Next the **mailDeliveryOption** attribute is checked. The name of this attribute can be changed with the **ldap_delivery_option** MTA option. This is a multi-valued option and its values determine the addresses produced by the alias translation process. Additionally, the permissible values are different for users and groups. Common permissible values are **program**, **forward**, and **hold**. User-only values are **mailbox**, **native**, **unix**, and **autoreply**. The group-only values are **members**, **members_offline**, and **file**.

The conversion of the **mailDeliveryOption** attribute into appropriate addresses is controlled by the **delivery_options** MTA option. This option not only specifies what

addresses are produced by each permissible **mailDeliveryOption** value, but also what the permissible **mailDeliveryOption** values are and whether or not each one is applicable to users, groups, or both.

The value of this option consists of a comma-separated list of **deliveryoption=template** pairs, each pair with one or more optional single character prefixes.

The default value of the **delivery_options** option is:

```
DELIVERY_OPTIONS=*mailbox=$M%\$2I$_+$2S@ims-ms-daemon, \
    &members=*, \
    *native=$M@native-daemon, \
    /hold=@hold-daemon:$A, \
    *unix=$M@native-daemon, \
    &file=+$F@native-daemon, \
    &@members_offline=*, \
    program=$M%$P@pipe-daemon, \
    #forward=**, \
    *^!autoreply=$M+$D@bitbucket
```

Each delivery option corresponds to a possible **mailDeliveryOption** attribute value and the corresponding template specifies the resulting address using the same metacharacter substitution scheme used by URL processing.

Table 6–5 shows the single character prefixes available for the **delivery_options** options.

Table 6–5 Single-character Prefixes for Options in the DELIVERY_OPTIONS MTA Option

Character Prefix	Description
@	Sets a flag saying that the message needs to be redirected to the reprocess channel. Processing of the current user/group is abandoned. Flag ignored for messages originating from the reprocess channel.
*	Delivery option applies to users.
&	Delivery option applies to groups.
\$	Sets a flag saying expansion of this user or group is to be deferred.
^	Sets a flag saying that the vacation start and end times should be checked to see if this delivery option really is in effect.
#	Sets a flag saying expansion of this delivery option does not need to take place on the entry's designated mailhost. That is, the following entry is mailhost-independent. This lets the MTA check to see if all of a given user or group's delivery options are independent of the mailhost. If this condition is satisfied the MTA can act on the entry immediately rather than having to forward the message to the mailhost.
/	Sets a flag that causes all addresses produced by this delivery option to be held. Message files containing these recipient addresses will have a .HELD extension.
!	Sets a flag that says that autoreply operations should be handled internally by the MTA. It only makes sense to use this prefix on an autoreply delivery option. The option's value should direct the message to the bitbucket channel.

If neither ***** nor **&** are present, the delivery option is taken to apply to both users and groups.

Additional Metacharacters for Use in Delivery Options

Several additional metacharacters have been added to support this new use of the MTA's URL template facility. [Table 6–6](#) shows additional metacharacters and their descriptions for use in delivery options.

Table 6–6 Additional Metacharacters for Use in Delivery Options

Metacharacter	Description
\$	Force subsequent text to lower case.
\$^	Force subsequent text to upper case.
\$_	Perform no case conversion on subsequent text.
$\$nA$	Insert the $_n$ th character of the address. The first character is character 0. The entire address is substituted if n is omitted. This is intended to be used to construct autoreply directory paths.
$\$D$	Insert the domain part of the address.
$\$nE$	Insert the value of the $_n$ th spare attribute. If n is omitted the first attribute is used.
$\$F$	Insert the name of the delivery file (mailDeliveryFileURL attribute).
$\$nG$	Insert the value of the $_n$ th spare attribute. If n is omitted the second attribute is used.
$\$nH$	Insert the $_n$ th component of the domain from the original address counting from 0. The default is 0 if n is omitted.
$\$nI$	Insert hosted domain associated with alias. This metacharacter accepts an integer option n whose semantics are described in Table 6–7, "Integers Controlling Behavior Modification of the $\\$nI$ and $\\$nS$ Metacharacters" .
$\$nJ$	Insert the $_n$ th part of the host domain counting from 0. The default for n is 0.
$\$_nO$	Insert source route associated with the current address. This metacharacter accepts an integer option n whose semantics are described in Table 6–7, "Integers Controlling Behavior Modification of the $\\$nI$ and $\\$nS$ Metacharacters" .
$\$K$	Insert an LDAP filter that matches the object classes for a user or group. See the description of the LDAP_UG_FILTER output-only MTA option.
$\$L$	Insert the local part of the address.
$\$nM$	Insert the $_n$ th character of the UID. The first character is character 0. The entire UID is substituted if n is omitted.
$\$P$	Insert the program name (mailProgramDeliveryInfo attribute).
$\$nS$	Insert subaddress associated with the current address. This metacharacter accepts an integer option n whose semantics are described in Table 6–7, "Integers Controlling Behavior Modification of the $\\$nI$ and $\\$nS$ Metacharacters" .
$\$nU$	Insert the n th character of the dequoted form of the mailbox part of the current address. The first character is character 0. The entire dequoted mailbox is substituted if n is omitted.
$\$nX$	Insert the $_n$ th component of the mailhost. The entire mailhost is inserted if n is omitted.

[Table 6–7](#) shows how the integer option modifies the behavior of the $\$nI$ and $\$nS$ metacharacters.

Table 6–7 Integers Controlling Behavior Modification of the $\$nI$ and $\$nS$ Metacharacters

Integer	Description of Behavior
0	Fail if no value is available (default).

Table 6–7 (Cont.) Integers Controlling Behavior Modification of the \$nI and \$nS

Integer	Description of Behavior
1	Insert value if one is available. Insert nothing if not.
2	Insert value if one is available. Insert nothing and delete preceding character if one is not (this particular behavior is needed by the ims-ms channel).
3	Insert value if one is available. Insert nothing and ignore following character if one is not.

In addition to the metacharacters, [Table 6–8](#) shows two special template strings.

Table 6–8 Special Template Strings

Special Template String	Description
*	Perform group expansion. This value is not valid for user entries.
**	Expand the attribute named by the ldap_forwarding_address MTA option. This defaults to mailForwardingAddress .

With group expansion, for example, if a user's **mailDeliveryOption** value is set to **mailbox**, we form a new address consisting of the stripped UID, a percent sign followed by the hosted domain if one is applicable, a plus sign followed by the subaddress if one was specified, and finally **@ims-ms-daemon**.

Delivery Option Defaults

If the list of active delivery options is empty at this point, the first option on the list (usually mailbox) is activated for users and the second option on the list (usually members) is activated for groups.

Start and End Date Checks

Start and end dates are checked after the delivery option list has been read. There are two attributes whose names are controlled by the **ldap_start_date** (default **vacationStartDate**) and **ldap_end_date** (default **vacationEndDate**) MTA options, respectively. If one or more of the active delivery options specified the ^ prefix character, the values of these options are checked against the current date. If the current date is outside the range specified by these options, the delivery options with the ^ prefix are removed from the active set. For more information see "[Vacation Autoreply Attributes](#)."

Optin and Presence Attributes

The **ldap_optin1** through **ldap_optin8** MTA options specify LDAP attributes for per-user spam filter opt-in values based on destination addresses. If an option is specified and the attribute is present, it is appended to the current spam filter opt-in list. Any values set by the domain level attribute set by the **ldap_domain_attr_optin** MTA option are also appended to the list. **ldap_source_optin1** through **ldap_source_optin8** provide comparable originator-address-based per-user spam filter optins.

The **ldap_presence** MTA option can be used to specify a URL that can be resolved to return presence information about the user. If the option is specified and the attribute is present, its value is saved for possible use in conjunction with Sieve presence tests. The domain level attribute set by the **ldap_domain_attr_presence** MTA option is used as source for this URL if no value exists for the user entry.

Sieve Filter Handling

Next the **mailSieveRuleSource** attribute is checked for a Sieve filter that applies to this entry. If this attribute exists, it is parsed and stored at this point. The two possible forms for the value of this attribute are a single value that contains a complete Sieve script or multiple values where each value contains a piece of a Sieve script. The latter form is produced by the web filter construction interface. Special code is used to order the values and glue them together properly.

The use of the **mailSieveRuleSource** attribute specifically can be overridden by using the **ldap_filter** MTA option.

Deferred Processing Control

Next the **mailDeferProcessing** attribute is checked. This attribute can be changed by using the **ldap_reprocess** MTA option. Processing continues normally if this attribute exists and is set to **no**. But if this attribute is set to **yes** and the current source channel is not the reprocess channel, expansion of this entry is aborted and the original *user@domain* address is simply queued to the reprocess channel. If this attribute does not exist, the setting of the deferred processing character prefix associated with delivery options processing is checked. (See the section "[Delivery Options Processing](#).") The default for users is **no**. The default for groups is controlled by the MTA option **defer_group_processing**, which defaults to 1 (yes). Alias processing concludes at this point for user entries.

Group Expansion Attributes

A number of additional attributes are associated with group expansion and must be dealt with at this point. The names of these attributes are all configurable by using various MTA options.

[Table 6–9](#) lists the default attribute names, the MTA option to set the attribute name, and the way the attribute is processed by the MTA. The ordering of the elements in the table shows the order in which the various group attributes are processed. This ordering is essential for correct operation.

Table 6–9 Group Expansion Default Attributes and MTA Option to Set

Default Attribute	(MTA option to Set Attribute Name) How the Attribute is Processed
mgrpMsgRejectAction	(ldap_reject_action) Single-valued attribute that controls what happens if any of the subsequent access checks fail. Only one value is defined: TOMODERATOR , which if set instructs the MTA to redirect any access failures to the moderator specified by the mgrpModerator attribute. The default (and any other value of this attribute) causes an error to be reported and the message rejected.
mailRejectText	(ldap_reject_text) The first line of text stored in the first value of this attribute is saved. This text will be returned if any of the following authentication attributes cause the message to be rejected. This means the text can appear in SMTP responses so value has to be limited to US-ASCII to comply with current messaging standards.
mgrpBroadcasterPolicy	(ldap_auth_policy) Specifies level of authentication needed to send to the group. Possible tokens are SMTP_AUTH_REQUIRED or AUTH_REQ , both of which mean that the SMTP AUTH command must be used to identify the sender in order to send to the group; SMTP_AUTH_USED and AUTH_USED which are similar in effect to SMTP_AUTH_REQUIRED and AUTH_REQ , but do not require posters to authenticate; PASSWORD_REQUIRED , PASSWD_REQUIRED , or PASSWD_REQ , all of which mean the password to the list specified by the mgrpAuthPassword attribute must appear in an Approved: header field in the message; OR , which changes the or_clauses MTA option setting to 1 for this list; AND , which changes the or_clauses MTA option setting to 0 for this list; and NO_REQUIREMENTS , which is non-operational. Multiple values are allowed and each value can consist of a comma-separated list of tokens. If SMTP AUTH is called for it also implies that any subsequent authorization checks will be done against the email address provided by the SASL layer rather than the MAIL FROM address.
mgrpAllowedDomain	(ldap_auth_domain) Domains allowed to submit messages to this group. A match failure with the or_clauses MTA option set to 0 (the default) means access checking has failed and all subsequent tests are bypassed. A match failure with the or_clauses MTA option set to 1 sets a "failure pending" flag; some other access check must succeed in order for access checking to succeed. This check is bypassed if the submitter has already matched an ldap_auth_url . Can be multi-valued and glob-style wildcards are allowed.
mgrpDisallowedDomain	(ldap_cant_domain) Domains not allowed to submit messages to this group. A match means access checking has failed and all subsequent checks are bypassed. This check is bypassed if the submitter has already matched an ldap_auth_url . Can be multi-valued and glob-style wildcards are allowed.
mgrpAllowedBroadcaster	(ldap_auth_url) URL identifying mail addresses allowed to send mail to this group. Can be multivalued. Each URL is expanded into a list of addresses and each address is checked against the current envelope from: address. A match failure with the or_clauses MTA option set to 0 (the default) means access checking has failed and all subsequent tests are bypassed. A match failure with the or_clauses MTA option set to 1 sets a "failure pending" flag. Some other allowed access check must succeed in order for access checking to succeed. A match also disables subsequent domain access checks. The expansion that is performed is similar to an SMTP EXPN with all access control checks disabled. List expansion in the context of the mgrpallowedbroadcaster LDAP attribute now includes all the attributes used to store email addresses (normally mail , mailAlternateAddress , and mailEquivalentAddress). Previously only mail attributes were returned, making it impossible to send to lists restricted to their own members using alternate addresses.
mgrpDisallowedBroadcaster	(ldap_cant_url) URL identifying mail addresses not allowed to send mail to this group. Can be multivalued. Each URL is expanded into a list of addresses and each address is checked against the current envelope from: address. A match means access checking has failed and all subsequent checks are bypassed. The expansion that is performed is similar to an SMTP EXPN with all access control checks disabled.
mgrpMsgMaxSize	(ldap_maximum_message_size) Maximum message size in bytes that can be sent to the group. This attribute is obsolete but still supported for backwards compatibility; the new mailMsgMaxBlocks attribute should be used instead.

Table 6–9 (Cont.) Group Expansion Default Attributes and MTA Option to Set

Default Attribute	(MTA option to Set Attribute Name) How the Attribute is Processed
mgrpAuthPassword	(ldap_auth_password) Specifies a password needed to post to the list. The presence of a mgrpAuthPassword attribute forces a reprocessing pass. As the message is enqueued to the reprocessing channel, the password is taken from the header and placed in the envelope. Then, while reprocessing, the password is taken from the envelope and checked against this attribute. Additionally, only passwords that actually are used are removed from the header field. The or_clauses MTA option acts on this attribute in the same way it acts on the other access check attributes.
mgrpModerator	(ldap_moderator_url) The list of URLs given by this attribute to be expanded into a series of addresses. The interpretation of this address list depends on the setting of the ldap_reject_action MTA option. If ldap_reject_action is set to TOMODERATOR , this attribute specifies the moderator address(es) the message is to be sent to should any of the access checks fail. If ldap_reject_action is missing or has any other value, the address list is compared with the envelope from address. Processing continues if there is a match. If there is no match, the message is again sent to all of the addresses specified by this attribute. Expansion of this attribute is implemented by making the value of this attribute the list of URLs for the group. Any list of RFC822 addresses or DNs associated with the group is cleared, and the delivery options for the group are set to members . Finally, subsequent group attributes listed in this table are ignored.
mgrpDeliverTo	(ldap_group_url1) List of URLs which, when expanded, provides a list of mailing list member addresses.
memberURL	(ldap_group_url2) Another list of URLs which, when expanded, provides another list of mailing list member addresses.
uniqueMember	(ldap_group_dn) List of DNs of group members. DNs may specify an entire subtree. Unique member DNs are expanded by embedding them in an LDAP URL. The exact URL to use is specified by the group_dn_template MTA option. The default value for this option is: ldap:///SA??sub?mail=*\$A specifies the point where the uniqueMember DN is inserted.
mgrpRFC822MailMember	(ldap_group_rfc822) Mail addresses of members of this list.
rfc822MailMember	(ldap_group_rfc822) rfc822MailMember is supported for backwards compatibility. Either rfc822MailMember or mgrpRFC822MailMember , but not both, can be used in any given group.
mgrpErrorsTo	(ldap_errors_to) Sets the envelope originator (MAIL FROM) address to whatever the attribute specifies.
mgrpAddHeader	(ldap_add_header) Turns the headers specified in the attribute into header trimming ADD options.
mgrpRemoveHeader	(ldap_remove_header) Turns the headers specified into header trimming MAXLINES=-1 options.
mgrpMsgPrefixText	(ldap_prefix_text) Adds the specified text to the beginning of the message text, if any.
mgrpMsgSuffixText	(ldap_suffix_text) Adds the specified text to the ending of the message text, if any.
No Default	(ldap_add_tag) Checks the subject for the specified text; if it isn't present the text is added at the beginning of the subject field.

One final attribute is checked in the special case of group expansion as part of an SMTP EXPN command: **mgmanMemberVisibility** or **expandable**. The **ldap_expandable** MTA option can be used to select different attributes to check. Possible values are: **anyone**, which means that anyone can expand the group, **all** or **true**, which mean that the user has to successfully authenticate with SASL before expansion will be allowed, and **none**, which means that expansion is not allowed. Unrecognized values

are interpreted as **none**. If the attribute is not present, the **expandable_default** MTA option controls whether the expansion is allowed.

Alias entries are cached in a fashion similar to domain entries. The MTA options controlling the alias cache are **alias_entry_cache_size** (default 1000 entries) and **alias_entry_cache_timeout** (default 600 seconds). The entire LDAP return value for a given alias is retained in the cache.

Negative caching of alias entries is controlled by the **alias_entry_cache_negative** MTA option. A non-zero value enables caching of alias match failures. A zero value disables it. Negative caching of alias entries is disabled by default. The theory is that repeated specification of an invalid address is unlikely to occur very often in practice. In addition, negative caching may interfere with timely recognition of new users added to the directory. However, sites should consider re enabling negative caching of aliases in situations where vanity domains are heavily used. The search performed by the URL specified in **alias_url0** is less likely to be successful.

To Modify Group Membership Attribute Syntax

Support has been added for postprocessing LDAP expansion results with a mapping. The new **ldap_url_result_mapping** MTA option can be used to specify the name of a group attribute which in turn specifies the name of a mapping. This mapping will be applied to any results returned by expanding either a **mgrpDeliverTo** or **memberURL** attribute. The mapping probe will be of the form:

LDAP-URL | LDAP-result

If the mapping returns with **\$Y** set the mapping result string will replace the LDAP result for alias processing purposes. If the mapping returns with **\$N** set the result will be skipped.

This mechanism can be used to define groups based on attributes that don't contain proper email address. For example, suppose a company has placed pager numbers in all their user entries. Messages can be sent to these numbers via email by suffixing them with a particular domain. A group could then be defined as follows:

1. Define a new **mgrpURLResultMapping** attribute in the directory and set the **ldap_url_result_mapping** MTA option to this attribute's name.
2. Define a page-all group with the following attributes:

```
mgrpDeliverTo: ldap:///o=usergroup?pagerTelephoneNumber?sub
mgrpURLResultMapping: PAGER-NUMBER-TO-ADDRESS
```

3. Define the mapping:

```
PAGER-NUMBER-TO-ADDRESS
*|* "$1"@pagerdomain.com$Y
```

Even more interesting effects can be achieved by combining this mechanism with the **process_substitutions** mechanism described in the *Messaging Server Reference*. For example, it would be easy to create a metagroup where sending to an address of the form:

pager+user@domain.com

sends a page to the user named **user**.

Address Reversal

Address reversal with direct LDAP starts with a `USE_REVERSE_DATABASE` value of 4, which disables the use of any reverse database. You should also set `USE_TEXT_DATABASES` to read the `IMTA_TABLE:reverse.txt` file, as the `sleepycat` databases are being deprecated. It then builds on the routing facilities previously discussed. In particular, it begins with a reverse URL specification of the form:

```
REVERSE_URL=ldap:/// $V? $N?mail?sub? $R
```

Note: Changing `REVERSE_URL` for any reason is not recommended.

The `$V` metacharacter has already been described in the context of alias URLs. The list of attributes to search comes from the MTA option `ldap_mail_reverses`. If this option is not set, the `mta.ldap_schematag_mconfig` option is examined, and depending on its value, an appropriate set of default attributes is chosen.

Table 6–10 shows the `mta.ldap_schematag` values and the default attributes chosen.

Table 6–10 *mta.ldap_schematag Values and Attributes*

Schema Tag Value	Attributes
<code>sims40</code>	<code>mail,rfc822mailalias</code>
<code>nms41</code>	<code>mail,mailAlternateAddress</code>
<code>ims50</code>	<code>mail,mailAlternateAddress</code>

The `$R` metacharacter is used to specify the filter. See "[The \\$R Metacharacter](#)" for more information.

The `$N` metacharacter returns a list of the attributes of interest to address reversal.

The value of `$N` cannot exactly be controlled: the MTA constructs it from its own, hard-coded (and subject to change) list of the relevant attributes for address reversal purposes. If you use the various `ldap_*` global MTA options to change what the MTA thinks are the names of attributes of interest, you will, in fact, fetch different attributes from LDAP. But it is always whatever attributes that correspond semantically to the MTA's idea of relevant attributes. These are: `ldap_capture` (no default), `ldap_recipientlimit` (no default), `ldap_recipientcutoff` (no default), `ldap_sourceblocklimit` (no default), `ldap_source_channel` (no default), `ldap_personal_name` (no default), `ldap_source_conversion_tag` (no default), `ldap_primary_address` (`mail`), `ldap_alias_addresses` (`mailAlternateAddress`), `ldap_equivalence_addresses` (`mailEquivalentAddress`), plus the `ldap_spare_*` attributes.

As always, `mta.ldap_schematag` can be a comma-separated list. If more than one schema is supported, the combined list of attributes, with duplicates eliminated, is used.

Additionally, the filter searches not only for the address that was originally supplied, but also for an address with the same local part but the domain that was actually found in the domain tree (which was saved in Step 2 of "[Rewrite Rule Machinery](#)"). The iterative nature of the domain tree lookup means the two addresses may be different.

For example, suppose that the domain `example.com` appears in the domain tree and the MTA looks the address:

u@host1.example.com

The filter that results from the expansion of **\$R** and an **ims50** schema tag will be something like:

```
( | (mail=u@example.com)
  (mail=u@host1.example.com)
  (mailAlternateAddress=u@example.com)
  (mailAlternateAddress=u@host1.example.com)
  (mailEquivalentAddress=u@example.com)
  (mailEquivalentAddress=u@host1.example.com) )
```

Reverse lookup returns several attributes, and the MTA knows to use the mail attribute (more precisely, the attribute named by **ldap_primary_address**) as the one for address reversal. Note that the **mailEquivalentAddress** (more precisely, the attribute named by **ldap_equivalence_addresses**) is also permitted.

After the URL is constructed an LDAP search is performed. If the search is successful, LDAP returns multiple attributes in essentially arbitrary order. Unsuccessful searches or errors leave the original address unchanged.

Due to the frequency with which address reversal operations are performed, especially given the number of addresses that can appear in a message header, and the expense of the directory queries involved, both negative and positive results need to be cached. This is implemented with an in-memory, open-chained, dynamically-expanded hash table. The maximum size of the cache is set by the **reverse_address_cache_size** MTA option (default 100000) and the timeout for entries in the cache is set by the **reverse_address_cache_timeout** MTA option (default 600 seconds). The cache actually stores addresses themselves, not the LDAP URLs and LDAP results.

Asynchronous LDAP Operations

Asynchronous lookups avoid the need to store an entire large LDAP result in memory, which can cause performance problems in some cases. The MTA provides the ability to perform various types of lookups done by the MTA asynchronously.

Use of asynchronous LDAP lookups is controlled by the MTA option **ldap_use_async**. This option is a bit-encoded value. Each bit, if set, enables the use of asynchronous LDAP lookups in conjunction with a specific use of LDAP within the MTA.

Table 6–11 shows the bit and value settings for the **ldap_use_async** MTA option.

Table 6–11 Settings for the ldap_use_async MTA Option

Bit	Value	Specific Use of LDAP
0	1	ldap_group_url1 (mgrpDeliverTo) URLs
1	2	ldap_group_url2 (memberURL) URLs
2	4	ldap_group_dn (UniqueMember) DNs
3	8	auth_list , moderator_list , sasl_auth_list , and sasl_moderator_list nonpositional list option URLs
4	16	cant_list , sasl_cant_list nonpositional list option URLs
5	32	originator_reply nonpositional list option URLs
6	64	deferred_list , direct_list , hold_list , nohold_list nonpositional list option URLs
7	128	username_auth_list , username_moderator_list , username_cant_list nonpositional list option URLs

Table 6–11 (Cont.) Settings for the `ldap_use_async` MTA Option

Bit	Value	Specific Use of LDAP
8	256	alias file list URLs
9	512	alias database list URLs
10	1024	<code>ldap_cant_url</code> (<code>mgrpDisallowedBroadcaster</code>) outer level URLs
11	2048	<code>ldap_cant_url</code> inner level URLs
12	4096	<code>ldap_auth_url</code> (<code>mgrpAllowedBroadcaster</code>) outer level URLs
13	8192	<code>ldap_auth_url</code> inner level URLs
14	16384	<code>ldap_moderator_url</code> (<code>mgrpModerator</code>) URLs
15	32768	<code>ldap_jettison_url</code> (<code>mgrpJettisonBroadcasters</code>) URLs

The default value of the `ldap_use_async` MTA option is 0, which means that asynchronous LDAP lookups are disabled by default.

Settings Summary

To enable direct LDAP, the following MTA options need to be set:

```
alias_magic=8764
alias_url0=ldap:///V?*sub?R
use_reverse_database=4
use_domain_database=0
reverse_url=ldap:///V?N?mail?sub?R
```

Note: `ldap:///V?N?mail?sub?R` is the default setting for `reverse_url`. In most instances, there's no need to explicitly set it in your configuration.

If vanity domains are to be supported, the following additional options must be set:

```
domain_match_url=ldap:///B?msgVanityDomain?sub? (msgVanityDomain=$D)
alias_url1=ldap:///B?*sub? (&(msgVanityDomain=$D)R)
alias_url2=ldap:///1V?*sub? (mailAlternateAddress=@$D)
```

Note that the last of these options also handle the case of wild carded local parts in hosted as well as vanity domains. If wild carded local part support is desired but vanity domain support is not, the following option should be used instead:

```
alias_url1=ldap:///V?*sub?&(mailAlternateAddress=@$D)
```

The filter `ssrd:$A` clause needs to be removed from the `ims-ms` channel definition in the MTA configuration.

Processing Multiple Different LDAP Attributes with the Same Semantics

The MTA has the ability to process multiple different LDAP attributes with the same semantics. Note that this is not the same as processing of multiple values for the same attribute, which has always been supported. The handling attributes receive depends on the semantics of the attribute. The possible options are:

1. Multiple different attributes don't make sense and render the user entry invalid. This handling is the default for all attributes unless otherwise specified.
2. **ldap_autoreply_subject**, **ldap_autoreply_text**, and **ldap_autoreply_text_internal** all receive the handling described in "[Vacation Autoreply Attributes](#)." **ldap_spare_3** and **ldap_personal_name** attributes are also in this category.
3. Multiple different attributes do make sense and should all be acted on. This handling is currently in effect for **ldap_capture**, **ldap_alias_addresses**, **ldap_equivalence_addresses**, and **ldap_detourhost_optin**. Note that **ldap_detourhost_optin** attribute was first added to Messaging Server 6.3.

About MTA Services

This chapter describes general MTA services and configuration in Unified Configuration.

MTA Configuration Overview

In a legacy configuration, you manage the MTA configuration by editing various text files. In Unified Configuration, you manage the MTA configuration by using the **msconfig** command. The Unified Configuration stores the MTA configuration in a single file, **config.xml** (for the most part). The **msconfig** command performs syntax validation on option names and values as well as the Unified Configuration structure to a limited degree. For more information on the syntax and options for the **msconfig** command, see "[msconfig Command](#)."

Caution: Only edit your Unified Configuration by running the **msconfig** command. This saves old configurations and allows for rollback. Do not hand-edit any of the Unified Configuration files. Oracle Support may occasionally edit these files to work around any issues pertaining to **msconfig**.

Table 7-1 provides a comparison of how the MTA configuration is managed in legacy and Unified Configuration.

Table 7-1 MTA Configuration: Legacy Versus Unified Configuration

Legacy Configuration Files	Unified Configuration Method
mappings file	Running msconfig edit mappings loads mappings in legacy format in the administrator's chosen editor.
imta.cnf file	Running msconfig edit channels loads channel blocks in legacy format in the administrator's chosen editor. Running msconfig edit rewrite loads rewrite rules in legacy format in the administrator's chosen editor.
option.dat file	Get or set options directly by running msconfig .
job_controller.cnf file	Get or set options directly by running msconfig ; most options require job_controller prefix.
dispatcher.cnf file	Get or set dispatcher options; most require a dispatcher. prefix, service-specific settings are in a service:name group, such as service:SMTP.tcp_ports .
conversions file	Run msconfig edit conversions .

Table 7-1 (Cont.) MTA Configuration: Legacy Versus Unified Configuration

Legacy Configuration Files	Unified Configuration Method
aliases file	Run <code>msconfig edit aliases</code> .

Compiling the MTA Configuration

This section includes the following topics:

- [Changes for Compiled MTA Configurations](#)
- [Recompiling the MTA Configuration](#)

Changes for Compiled MTA Configurations

The following changes for compiled MTA configurations were introduced:

- The **configure** command no longer generates a compiled configuration by default, for both legacy and Unified Configuration.
- Unified Configuration provides the **-xmlfile=** switch to use test tools on a non-live configuration. (This was previously only possible with a compiled configuration.)
- The **imsimta version** command now shows if a compiled configuration is used or not.

Recompiling the MTA Configuration

Whenever you make a change to the MTA configuration, such as to mappings, rewrite rules, channel blocks, aliases, and so on, you must recompile the configuration, if you are using a compiled configuration. This compiles the configuration files into a single image in shared memory.

The compiled configuration has a static and dynamic reloadable part. If the dynamic part is changed, and you run the **imsimta reload** command, a running program reloads the dynamic data. The dynamic parts are mapping tables, aliases, and lookup tables.

Using a compiled configuration enables you to be able to test configuration changes more conveniently because the configuration files themselves are not "live" when a compiled configuration is in use.

Whenever a component of the MTA (such as a channel program) must read the configuration file, it first checks to see if a compiled configuration exists. If it does, the image is attached to the running program. If the image attach operation fails, the MTA falls back on reading the **config.xml** file.

If you make changes to the **reverse**, **forward**, or general databases, then issue the command **imsimta reload** to get the changes to take effect. If you make changes to rewrite rules, channel blocks, mappings, aliases, conversions, or job controller options, then you should issue an **imsimta cnbuild** command followed by an **imsimta restart smtp** command. If you make changes to dispatcher options, you need to do an **imsimta restart dispatcher** command. If you make changes to options that are included in the compiled configuration that affect the job controller, but not the SMTP server, in many cases you should issue the following commands: **imsimta cnbuild** and **imsimta restart job_controller**.

If you make changes to options that are included in the compiled configuration that affect both the SMTP server and the job controller, you should issue the following commands:

```
imsimta cnbuild
imsimta restart smtp
imsimta restart job_controller
```

(See "MTA Commands" for details on these commands.)

Other instances where you must restart the job controller:

- Changing the job controller configuration.
- Adding or changing use of the channel options **pool**, **maxjobs**, **master**, **slave**, **single**, **single_sys**, or **multiple**. Adding or changing a **threaddepth** channel option can be dealt with instead by running the **imsimta cache -change -thread_depth=...** command.
- If you want changes to master channel jobs to take effect immediately (rather than waiting for the controller to time-out existing channel jobs), then any relevant (which means almost all) changes to the MTA configuration or channels. (Changes to mappings or to MTA databases: (1) aren't usually relevant for outbound channel jobs, though they can matter to "intermediate" channels such as **conversion**, **process**, **reprocess**, and (2) if those intermediate channels are a concern, changes to mappings or to databases can often be handled by running the **imsimta reload** command, avoiding a restart of the Job Controller.) The desire to have changes take effect immediately needs to be balanced against the damage that restarting the Job Controller does. Also consider just how much longer a particular sort of job is going to run anyhow.

Any changes (such as additions or changes to options on channel definitions) also require running the **imsimta cnbuild** command. That's a basic, regardless of whether a Job Controller restart is needed.

Try to avoid restarting the Job Controller, especially at times of large numbers of messages in the queues, unless one of the preceding conditions necessitates a restart.

For more information on MTA services see the *Messaging Server Reference*.

Message Transfer Agent Command-line Utilities

The Message Transfer Agent (MTA) command-line utilities are used to perform various MTA maintenance, testing, and management tasks.

The MTA commands are also referred to as the **imsimta** commands. The **imsimta** script is located in the *MessagingServer_home/* directory.

MessagingServer_home represents the directory path in which you install the server.

MTA Commands

The commands are listed in [Table 8-1](#).

Table 8-1 Valid mail folder names.

Command	Description
<code>imsimta cache</code>	Performs operations on the queue cache.
<code>imsimta chbuild</code>	Compiles the MTA character set conversion tables.
<code>imsimta cnbuild</code>	Compiles the MTA configuration files.
<code>imsimta counters</code>	Performs operations on the channel counters.
<code>imsimta crdb</code>	Creates an MTA database.
<code>imsimta find</code>	Locates the precise filename of the specified version of an MTA log file.
<code>imsimta kill</code>	Terminates the specified process.
<code>imsimta process</code>	Lists currently running MTA jobs.
<code>imsimta program</code>	Manipulates the MTA program delivery options.
<code>imsimta purge</code>	Purges MTA log files.
<code>imsimta qclean</code>	Holds or deletes message files containing specific substrings in their envelope From: address, Subject: line, or content.
<code>imsimta qm</code>	Manages MTA message queues.
<code>imsimta qtop</code>	Displays the most frequently occurring envelope From: Subject:, or message content fields found in message files in the channel queues.
<code>imsimta refresh</code>	Combines the functionality of the imsimta cnbuild and imsimta restart utilities.
<code>imsimta reload</code>	Allows changes to certain configuration files to take effect without restarting the server.

Table 8–1 (Cont.) Valid mail folder names.

Command	Description
<code>imsimta renamedb</code>	Renames a MTA database.
<code>imsimta restart</code>	Restarts detached MTA processes.
<code>imsimta return</code>	Returns (bounces) a mail message to its originator.
<code>imsimta run</code>	Processes messages in a specified channel.
<code>imsimta shutdown</code>	Shuts down the MTA Job Controller and the MTA Dispatcher as well as individual processes running under the Dispatcher.
<code>imsimta start</code>	Starts the MTA Job Controller and Dispatcher.
<code>imsimta stop</code>	Shuts down the MTA Job Controller and the MTA Dispatcher as well as individual processes running under the Dispatcher.
<code>imsimta submit</code>	Processes messages in a specified channel.
<code>imsimta test</code>	Performs tests on mapping tables, wildcard patterns, address rewriting, and URLs.
<code>imsimta test -domain</code>	Verifies the validity of domain structures in the directory and displays domain-related information.
<code>imsimta version</code>	Prints the MTA version number.
<code>imsimta view</code>	Displays log files.

Command Descriptions

You need to be logged in as **root** (UNIX) or administrator (Windows NT) to run the MTA commands. Unless mentioned otherwise, all MTA commands should be run as **mailsrv** (the mail server user that is created at installation).

imsimta cache

The MTA maintains an in-memory cache of all the messages currently stored in its queues. This cache is called the queue cache. The purpose of the queue cache is to make dequeue operations perform more efficiently by relieving master programs from having to open every message file to find out which message to dequeue and in which order.

Syntax

```
imsimta cache -change -parameter | -dump | -sync[hronize] | -view <channel> | -walk
```

Options

The options for this command are:

Table 8–2 imsimta cache Command Options

Option	Description
<code>-change -global -debug=<i>integer</i></code>	Allows debugging for the job controller on the fly.
<code>-change -global -max_messages=<i>integer</i></code>	Allows setting the maximum number of messages for the job controller on the fly.
<code>-change -channel_template=<i>namemaster_job=command</i></code>	Allows changing the channel template for the job controller on the fly. Changing options for a channel template (e.g., tcp_*) changes that option for all channels derived from that template.

Table 8–2 (Cont.) imsimta cache Command Options

Option	Description
-change -channel_ template=nameslave_ job=command	Allows changing the channel template for the job controller on the fly. Changing options for a channel template (e.g., tcp_*) changes that option for all channels derived from that template.
-change -channel=namemaster_ job=command	Allows changing the master job for the job controller on the fly.
-change -channel=nameslave_ job=command	Allows changing the slave job for the job controller on the fly.
-change -channel=namethread_ depth=integer	Allows setting the thread depth for the job controller on the fly.
-change -channel=namejob_ limit=integer	Allows setting the job limit for the job controller on the fly.
-dump	Provides the same output as the -view switch with no channel specified.
-sync	Updates the active queue cache by updating it to reflect all non-held message files currently present in the <i>MessagingServer_home/imta/queue/</i> subdirectories. Note that the -sync option does not remove entries from the queue cache. The queue cache entries not corresponding to an actual queued message are silently discarded by channel master programs.
-view <channel>	Shows the current non-held entries in the MTA queue cache for a channel. <channel> is the name of the channel for which to show entries e.g. <i>tcp_local</i> This is a potentially expensive command if you have a large backlog of messages. If a <channel> is not specified, entries in all channels will be displayed.
-walk [-debug=xxx]	Used to diagnose potential problems with the job controller. Each time this command is run, the internal state of the message queues and job scheduling information is written to the job_controller.log file. In addition, the job controller debug mask is set to the value given. You should not run this command unless you are instructed to do so by support.

Examples

To synchronize the queue cache:

```
imsimta cache -sync
```

To view entries in the queue cache for the *tcp_local* channel, execute the command:

```
imsimta cache -view tcp_local
```

imsimta chbuild

The **imsimta chbuild** command compiles the character set conversion tables and loads the resulting image file into shared memory. The MTA ships with complete character set tables so you would not normally need to run this command. You would use **imsimta chbuild** only if you added or modified any character sets.

Syntax

```
imsimta chbuild [-image_file=<file_spec> | -noimage_file]
[-maximum | -nomaximum] [-option_file=<option_file>]
| -nooption_file] [-remove] [-sizes | -nosizes] [-statistics |
-nostatistics]
```

Options

The options for this command are:

Table 8–3 *imsimta chbuild* Command Options

Option	Description
-image_file = <i>file_spec</i> -noimage_file	By default, imsimta chbuild creates as output the image file named by the <code>IMTA_CHARSET_DATA</code> option of the MTA tailor file, <code>MessagingServer_home/config/imta_tailor</code> . With the -image_file option, an alternate file name may be specified. When the -noimage_file option is specified, imsimta chbuild does not produce an output image file. The -noimage_file option is used in conjunction with the -option_file option to produce as output an option file that specifies table sizes adequate to hold the tables required by the processed input files.
-maximum -nomaximum	The file <code>MessagingServer_home/config/maximum_charset.dat</code> is read in addition to the file named by the <code>IMTA_CHARSET_OPTION_FILE</code> option of the MTA tailor file, <code>MessagingServer_home/config/imta_tailor</code> , when -maximum is specified. This file specifies near <i>-maximum</i> table sizes but does not change any other settings. Use this option only if the current table sizes are inadequate. The -noimage and -option_file options should always be used in conjunction with this option - it makes no sense to output the enormous configuration that is produced by -maximum , but it does make sense to use -maximum to get past size restrictions in order to build a properly sized option file for use in building a manageable configuration with a subsequent imsimta chbuild invocation.
-option_file = <i>option_file</i> -nooption_file	imsimta chbuild can produce an option file that contains the correct table sizes to hold the conversion tables that were just processed (plus a little room for growth). The -option_file option causes this file to be output. By default, this file is the file named by the <code>IMTA_CHARSET_OPTION_FILE</code> option of the MTA tailor file, <code>MessagingServer_home/config/imta_tailor</code> . The value of the -option_file option may be used to specify an alternate file name. If the -nooption_file option is given, then no option file is output. imsimta chbuild always reads any option file (for example, the file named by the <code>IMTA_CHARSET_OPTION_FILE</code> option of the MTA tailor file) that is already present; use of this option does not alter this behavior. However, use of the -maximum option causes imsimta chbuild to read options from <code>maximum_charset.dat</code> in addition to <code>IMTA_CHARSET_OPTION_FILE</code> . This file specifies near-maximum table sizes. Use this option only if the current table sizes are inadequate, and only use it to create a new option file. The -noimage_file option should always be specified with -maximum , since a maximum-size image would be enormous and inefficient.
-remove	Removes any existing compiled character set conversion table. This is the file named by the <code>IMTA_CHARSET_DATA</code> option of the MTA tailor file, <code>MessagingServer_home/config/imta_tailor</code> .

Table 8-3 (Cont.) *imsimta chbuild* Command Options

Option	Description
-sizes -nosizes	The -sizes option instructs imsimta chbuild to output or suppress information on the sizes of the uncompiled conversion tables. The -nosizes option is the default.
-statistics -nostatistics	The -statistics option instructs imsimta chbuild to output or suppress information on the compiled conversion tables. This information gives a rough measurement of the efficiency of the compilation, and may indicate whether or not an additional rebuild with the -option_file option is needed. The -nostatistics option is the default.

Example

The standard command you use to compile character set conversion tables is:

```
imsimta chbuild
```

imsimta cnbuild

The **imsimta cnbuild** command compiles the textual configuration, option, mapping, conversion, circuit check and alias files, and loads the resulting image file into shared memory. The resulting image is saved to a file usually named **imta/lib/config_data** by the **IMTA_CONFIG_DATA** option of the MTA tailor file, *MessagingServer_home/config/imta_tailor*.

Whenever a component of the MTA (for example, a channel program) must read a compiled configuration component, it first checks to see whether the file named by the MTA tailor file option **IMTA_CONFIG_DATA** is loaded into shared memory; if this compiled image exists but is not loaded, the MTA loads it into shared memory. If the MTA finds (or not finding, is able to load) a compiled image in shared memory, the running program uses that image.

The reason for compiling configuration information is simple: performance. The only penalty paid for compilation is the need to recompile and reload the image any time the underlying configuration files are edited. Also, be sure to restart any programs or channels that load the configuration data only once when they start up—for example, the MTA multithreaded SMTP server.

It is necessary to recompile the configuration every time changes are made to any of the following files:

- MTA configuration file (or any files referenced by it)
- MTA system alias file
- MTA mapping file
- MTA option file
- MTA conversion file
- MTA security configuration file
- MTA circuit check configuration file
- MTA system wide filter file

Specifically, these are the files pointed at by the MTA tailor file options **IMTA_CONFIG_FILE**, **IMTA_ALIAS_FILE**, **IMTA_MAPPING_FILE**, **IMTA_OPTION_**

FILE, and IMTA_CONVERSION_FILE respectively, which usually point to the following files:

- *MessagingServer_home/config/imta.cnf*
- *MessagingServer_home/config/aliases*
- *MessagingServer_home/config/mappings*
- *MessagingServer_home/config/option.dat*
- *MessagingServer_home/config/conversions*

Note: Until the configuration is rebuilt, changes to any of these files are not visible to the running MTA system.

Syntax

```
imsimta cnbuild [-image_file=<file_spec> | -noimage_file]
[-maximum | -nomaximum] [-option_file=<option_file>]
| -nooption_file] [-remove] [-sizes | -nosizes] [-statistics |
-nostatistics]
```

Options

The options for this command are:

Table 8-4 *imsimta cnbuild Command Option*

Option	Description
-image_file= <i>file_spec</i> -noimage_file	By default, imsimta cnbuild creates as output the image file named by the IMTA_CONFIG_DATA option of the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> . With the -image_file option, an alternate filename can be specified. When the -noimage_file option is specified, imsimta cnbuild does not produce an output image file. This option is used in conjunction with the -option_file option to produce as output an option file which specifies table sizes adequate to hold the configuration required by the processed input files. The default value is -image_file=IMTA_CONFIG_DATA .
-maximum -nomaximum	<i>MessagingServer_home/config/maximum.dat</i> is read in addition to the file named by the IMTA_OPTION_FILE option in the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> . This file specifies near maximum table sizes but does not change any other option file parameter settings. Only use this option if the current table sizes are inadequate. The -noimage and -option_file options should always be used in conjunction with this qualifier; it makes no sense to output the enormous configuration that is produced by -maximum , but it does make sense to use -maximum to get past size restrictions in order to build a properly-sized option file so that a proportionately-sized configuration can be built with a subsequent imsimta cnbuild invocation. The default is -nomaximum .

Table 8–4 (Cont.) *imsimta cnbuild* Command Option

Option	Description
-option_file =[option_file] -nooption_file	imsimta cnbuild can optionally produce an option file that contains correct table sizes to hold the configuration that was just compiled (plus a little room for growth). The -option_file option causes this file to be output. By default, this file is the file named by the IMTA_OPTION_FILE option in the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> . The value on the -option_file option may be used to specify an alternate file name. If the -nooption_file option is given, then no option file will be output. imsimta cnbuild always reads any option file that is already present via the IMTA_OPTION_FILE option of the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> ; use of this option will not alter this behavior. However, use of the -maximum option causes imsimta cnbuild to read MTA options from the <i>MessagingServer_home/config/maximum.dat</i> file in addition to reading the file named by IMTA_OPTION_FILE . This file specifies near maximum table sizes. Use this option only if the current table sizes are inadequate, and only to create a new option file. The -noimage_file option should always be specified when -maximum is specified since a maximum-size image would be enormous and wasteful. The default value is -option_file=IMTA_OPTION_FILE .
-remove	Remove any existing compiled configuration; for example, remove the file named by the IMTA_CONFIG_DATA option of the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> .
-sizes -nosizes	The -sizes option instructs imsimta cnbuild to output information on the sizes of uncompiled MTA tables. The -nosizes option is the default.
-statistics -nostatistics	The -statistics option instructs imsimta cnbuild to output information table usage. This information gives a rough measurement of the efficiency of the compilation, and may indicate whether or not an additional rebuild with the -resize_tables option is needed. The -nostatistics option is the default.

Examples

To regenerate a compiled configuration enter the following command:

```
imsimta cnbuild
```

After compiling the configuration, restart any programs that may need to reload the new configuration. For example, the SMTP server should be restarted:

```
imsimta restart dispatcher
```

Note: **imsimta cnbuild** is executed whenever the **imsimta refresh** command is invoked.

imsimta counters

The MTA accumulates message traffic counters for each of its active channels. These statistics, referred to as channel counters, are kept in shared memory. The **imsimta counters** command manipulates these counters.

Syntax

```

imsimta counters -clear

imsimta counters -create [-max_channels=<value>]

imsimta counters -delete

imsimta counters -show [-associations | -noassociations] [-channels |
-nochannels] [-headers | -noheaders] [-output=<file_spec>]

```

Options

The options for this command are:

Table 8–5 *imsimta counters Command Options*

Option	Description
-associations -noassociations	Specifies whether or not to show the in-memory cache of association counters. The -associations option is the default. This option is only used with the -show option.
-channels -nochannels	Specifies whether or not to show the in-memory cache or channel counters. The -channels option is the default. This option is only used with the -show option.
-clear	The -clear command clears the in-memory channel counters.
-create	Creates the in-memory channel counters. Counters are not created by default. You must create the counters if you wish to use them. You must create them after restarting the MTA as well.
-headers -noheaders	Controls whether or not a header line describing each column in the table of counters is output. The -headers option is the default. This option is only used with the -show option.
-max_channels= <i>value</i>	By default, the in-memory channel counters can hold information for CHANNEL_TABLE_SIZE channels. CHANNEL_TABLE_SIZE is the value specified by the MTA option file option of the same name. Use the -max_channels= <i>value</i> option to select a different size. This option is used only with the -create option.
-delete	Deletes the in-memory channel counters.
-show	Displays the in-memory channel counters.
-headers -noheaders	Controls whether or not a header line describing each column in the table of counters is output. The -headers option is the default. This option is only used with the -show option.
-output= <i>file_spec</i>	Directs the output to the specified file. By default, the output appears on your display. This option is only used with the -show option.

Examples

To display the counters for all channels:

```
imsimta counters -show
```

imsimta crdb

The **imsimta crdb** command creates and updates MTA database files. **imsimta crdb** converts a plain text file into MTA database records; from them, it either creates a new database or adds the records to an existing database.

In general, each line of the input file must consist of a left side and a right side. The two sides are separated by one or more spaces or tabs. The left side is limited to 32 characters in a short database (the default variety) and 80 characters in a long database. The right side is limited to 80 characters in a short database and 256 in a long database. Spaces and tabs may not appear in the left side unless the **-quoted** option is specified. Comment lines may be included in input files. A comment line is a line that begins with an exclamation mark (!) in column 1.

Syntax

```
imsimta crdb <input-file-spec> <output-database-spec> [-append | -noappend]
[-count | -nocount] [-duplicates | -noduplicates] [-long_records |
-nolong_records] [-quoted | -noquoted] [-remove | -noremove] [-statistics |
-nostatistics] [-strip_colons | -nostrip_colons]
```

Options

The options for this command are:

Table 8–6 *imsimta crdb* Command Options

Option	Description
<i>input-file-spec</i>	A text file containing the entries to be placed into the database. Each line of the text file must correspond to a single entry. This attribute is mandatory.
<i>output-database-spec</i>	The initial name string of the files to which to write the database (unless -dump is specified). The .db extension is appended to the file name. This attribute is mandatory.
-append -noappend	When the default, -noappend , option is in effect, a new database is created, overwriting any old database of that name. Use the -append option to instruct the MTA to instead add the new records to an existing database. The -noappend option is the default. In the event of a duplicate record, the newly appended record overwrites the old record when -noduplicates is specified.
-count -nocount	Controls whether or not a count is output after each group of 100 input lines are processed. The -count option is the default.
-duplicates -noduplicates	Controls whether or not duplicate records are allowed in the output files. Currently, duplicate records are of use only in the domain database (rewrite rules database) and databases associated with the directory channel. The -noduplicates option is the default.
-long_records -nolong_records	Controls the size of the output records. By default, left sides are limited to 32 characters and right sides are limited to 80 characters. If -long_records is specified, the limits are changed to 80 and 256, respectively. The -nolong_records option is the default.
-quoted -noquoted	Controls the handling of quotes. Normally imsimta crdb pays no attention to double quotes. If -quoted is specified, imsimta crdb matches up double quotes in the process of determining the break between the left and right hand sides of each input line. Spaces and tabs are then allowed in the left side if they are within a matching pair of quotes. This is useful for certain kinds of databases, where spaces may form a part of database keys. The quotes are not removed unless the -remove option is also specified. The -noquoted option is the default.

Table 8–6 (Cont.) *imsimta crdb* Command Options

Option	Description
-remove -noremove	Controls the removal of quotes. If imsimta crdb is instructed to pay attention to quotes, the quotes are normally retained. If -remove is specified, imsimta crdb removes the outermost set of quotes from the left hand side of each input line. Spaces and tabs are then allowed in the left side if they are within a matching pair of quotes. This is useful for certain kinds of databases, where spaces may form a part of database keys. -remove is ignored if -quoted is not in effect. The -noremove option is the default.
-statistics -nostatistics	Controls whether or not some simple statistics are output by imsimta crdb , including the number of entries (lines) converted, the number of exceptions (usually duplicate records) detected, and the number of entries that could not be converted because they were too long to fit in the output database. -nostatistics suppresses output of this information. The -statistics option is the default.
-strip_colons -nostrip_colons	Instructs imsimta crdb to strip a trailing colon from the right end of the left hand side of each line it reads from the input file. This is useful for turning alias file entries into an alias database. The -nostrip_colons is the default.

Example

The following commands create an alias database with "long" record entries. The creation is performed in a two-step process using a temporary database to minimize any window of time, such as during database generation, when the database would be locked and inaccessible to the MTA.

```
imsimta crdb -long_records aliases-tmp
imsimta renamedb aliases-tmp IMTA_ALIAS_DATABASE
```

imsimta crdb -dump

The **imsimta crdb -dump** command writes the entries in MTA databases to a flat ASCII file. In particular, this command may be used to write the contents of an old style database to a file from which a new style database may be built using the **imsimta crdb** command. The output begins with a comment line that displays a proper **imsimta crdb** command to use in order to return the ASCII output to a database.

Note: Make sure you are logged in as **mailsrv** (the mail server user) before performing this command.

Syntax

```
imsimta crdb -dump <input-database-spec> [<output-file-spec>]
```

Options

The options for this command are:

Table 8–7 *imsimta crdb -dump Command Options*

Option	Description
<i>input-database-spec</i>	Database from which to read entries. By default, the MTA looks for a current format database of the given name; if this does not exist, the MTA will look for an old format database of the given name. The special keywords IMTA_ALIAS_DATABASE , IMTA_REVERSE_DATABASE , and IMTA_GENERAL_DATABASE are supported; the use of such a special keyword tells the MTA to dump the database specified by the corresponding MTA tailor file option.
<i>output-file-spec</i>	ASCII file to which the entries stored in the database are written. This file should be in a directory where you have write permissions. If an output file is not specified, the output is written to stdout .

Examples

The following command can be used to dump the contents of an alias database to a file, and then to recreate the alias database from that file.

```
imsimta crdb -dump IMTA_ALIAS_DATABASE alias.txt
imsimta crdb alias.txt alias-tmp
imsimta renamedb alias-tmp IMTA_ALIAS_DATABASE
```

imsimta find

The **imsimta find** utility locates the precise filename of the specified version of an MTA log file. MTA log files have a **-uniqueid** appended to the filename to allow for the creation of multiple versions of the log file. On UNIX, the **-uniqueid** is appended to the very end of the filename (the end of the file extension), while on Windows NT, the **-uniqueid** is appended to the end of the name part of the filename, before the file extension. The **imsimta find** utility understands these unique ids and can find the particular filename corresponding to the requested version of the file.

Syntax

```
imsimta find <file-pattern> [-f=<offset-from-first>] [-l=<offset-from-last>]
```

Options

The options for this command are:

Table 8–8 *imsimta find Command Options*

Option	Description
<i>-f=offset-from-first</i>	Finds the specified version of the file (starting from 0). For example, to find the earliest (oldest) version of the file, specify -f=0 . By default, imsimta find finds the most recent version of the file.
<i>-l=offset-from-last</i>	Finds the last version of the specified file. For example, to find the most recent (newest) version of the file, specify -l=0 . By default, imsimta find finds the most recent version of the file.
<i>file-pattern</i>	Specifies a filename pattern for which the log file to find.

Examples

The following command prints out the filename of the **tcp_local_slave.log-uniqueid** file most recently created:

```
imsimta find MessagingServer_home/data/log/tcp_local_slave.log
```

The following command displays the filename of the oldest **tcp_bitnet_master.log-uniqueid** file:

```
imsimta find MessagingServer_home/data/log/tcp_bitnet_master.log -f=0
```

imsimta kill

The **imsimta kill** utility immediately and indiscriminately terminates the specified process. This command is equivalent to the UNIX **kill -9** command. The process is terminated even if it is in the middle of transferring email. So use of the **imsimta shutdown** utility, which performs an orderly shutdown, is generally preferable.

Syntax

```
imsimta kill <component>
```

Note: You must have the same process id as the process to be killed, or be **root**. This utility is not available on Windows NT.

component is the MTA component to be killed. Valid values are **job_controller** and **dispatcher**.

imsimta process

This command displays the current MTA processes. Additional processes may be present if messages are currently being processed, or if certain additional MTA components are in use.

Syntax

```
imsimta process
```

Example

The following command shows current MTA processes:

```
imsimta process
imsimta process
```

USER	PID	S	VSZ	RSS	STIME	TIME	COMMAND
mailsrv	15334	S	21368	9048	17:32:44	0:01	imta/bin/dispatcher
mailsrv	15337	S	21088	10968	17:32:45	0:01	imta/bin/tcp_smtp_server
mailsrv	15338	S	21080	11064	17:32:45	0:01	imta/bin/tcp_smtp_server
mailsrv	15349	S	21176	10224	17:33:02	0:02	imta/bin/job_controller

imsimta program

The **imsimta program** command is used to manipulate the program delivery options.

This command can be executed as **root** or **mailsrv**. **mailsrv** is the default user for Oracle Communications Messaging Server, but could be whatever the specified user name for the Messaging Server is when Messaging Server is installed.

The program is passed the entire message, unparsed from **stdin**. This includes the From line (without the colon) as the first line, followed by the headers and the message body. This may include any MIME attachments that are part of the message.

Syntax

```
imsimta program -a -m <method> -p <program>
[-g <argument_list>] [-e <exec_permission>]
imsimta program -d -m <method>
```

```
imsimta program -c -m <method> -p <program>
| -g <argument_list> | -e <exec_permission>
```

Options

The options for this command are:

Table 8–9 imsimta program Command Options

Option	Description
-a	Add a method to the set of program delivery methods. This option cannot be used with the -d , -c , -l , or -u options.
-c	Change the arguments to a program that has already been entered.
-m <i>method</i>	Name given by the administrator to a particular method. This will be the name by which the method will be advertised to users. Method names must not contain spaces, tabs, or equal signs (=). The method name cannot be none or local. The method name is restricted to U.S. ASCII. This option is required with the -a , -d , -c , and -u options.
-p <i>program</i>	Actual name of the executable for a particular method. The executable should exist in the programs directory (<i>MessagingServer_home/data/site-programs</i>) for the add to be successful. It can be a symbolic link to an executable in some other directory. This option is required with the -a option.
-g <i>argument_list</i>	Argument list to be used while executing the program. If this option is not specified during an add, no arguments will be used. Each argument must be separated by a space and the entire argument list must be given within double quotes. If the %s tag is used in the argument list, it will be substituted with the user's username for programs executed by the users and with <i>username+programlabel</i> for programs executed by inetmail. <i>programlabel</i> is a unique string to identify that program. This option can be used with the -a and -c options.
-e <i>exec_permission</i>	<i>exec_permission</i> can be user or postmaster . If it is specified as user , the program is executed as the user. By default, execute permission for all programs are set to postmaster. Programs with <i>exec_permission</i> set to user can be accessed by users with UNIX accounts only. This option can be used with the -a and -c options. The directory from where this program is run as postmaster is the postmaster's home directory. If specified as user , then the user's home directory is the environment where the program is run as the user.
-d	Delete a method from the list of supported program delivery methods. This option cannot be used with the -a , -c , -l , or -u options.
-h	Help for this command.
-l	List all methods.

Table 8–9 (Cont.) imsimta program Command Options

Option	Description
-u	List all users using the method specified with the -m option.

Examples

To add a method **procmail1** that executes the program **procmail** with the arguments **-d username** and executes as the user, enter the following:

```
imsimta program -a -m procmail1 -p procmail -g "-d %s" -e user
```

imsimta purge

The **imsimta purge** command deletes older versions of MTA log files. **imsimta purge** can determine the age of log files from the uniqueid strings terminating MTA log file names.

Syntax

```
imsimta purge [<file-pattern>] -day=<dvalue> -hour=<hvalue> -num=<nvalue>
```

Options

The options for this command are:

Table 8–10 imsimta purge Command Options

Option	Description
<i>file-pattern</i>	If specified, the <i>file-pattern</i> option is a file name pattern that establishes which MTA log files to purge. The default pattern, if none is specified, is log/imta/log .
-day=dvalue	Purges all but the last <i>dvalue</i> days worth of log files.
-hour=hvalue	Purges all but the last <i>hvalue</i> hours worth of log files.
-num=nvalue	Purges all but the last <i>nvalue</i> log files. The default is 5.

Example

To purge all but the last five versions of each type of log file in the **log/imta** directory:

```
imsimta purge
```

imsimta qclean

The **imsimta qclean** utility holds or deletes message files containing specific substrings in their envelope From:address, Subject: line, or content.

Syntax

```
imsimta qclean
  [-content=<substring>] [-from=<substring>] [-subject=<substring>]
  [-to=<substring>] [-domain_to=<substring>] [-database] [-delete | -hold]
  [-directory_tree] [-ignore_zz] [-match=<keyword>] [-min_length=<n>]
  [-threads | -nothreads] [-verbose | -noverbose] [<channel>]
```

Options

The options for this command are:

Table 8–11 *imsimta qclean* Command Options

Option	Description
-content=substring-from=s ubstring-subject=substring -to=substring-domain_ to=substring	Specifies the substrings for which to search. Any combination of -content , -from , -subject , -to , and -domain_to may be specified. However, only one of each may be used. When a combination of such options is used, the -match option controls whether the options are interpreted as further restrictions (-match=AND) or as alternatives (-match=OR). The -domain_to option scans for frequently occurring envelope To: addresses. Identical to the -to option, except -domain_to looks at only the host.domain portion of the envelope To: address.
-database	Specifies that only message files identified by the queue cache is searched.
-delete	Deletes matching message files.
-hold	Holds matching message files.
-directory_tree	Searches all message files that are actually present in the channel queue directory tree.
-ignore_zz	Ignores queued message files with file names beginning with "ZZ". This option may be used to scan only those message files which represent queued messages which have failed at least one delivery attempt.
-match=keyword	Controls whether a message file must contain all (-match=AND) or only one of (-match=OR) the specified substrings in order to be held or deleted. The default is -match=AND .
-min_length=n	Specifies the minimum length of the substring for which to search. By default, each substring must be at least 24 bytes long. Use the -min_length option to override this limit.
-threads=n -nothreads	Accelerates the searching on multiprocessor systems by dividing the work amongst multiple, simultaneous running threads. To run <i>n</i> simultaneous searching threads, specify -threads=n . The value <i>n</i> must be an integer between 1 and 8. The default is -nothreads .
-verbose -noverbose	Requests that the utility displays operation information (-verbose). The default is -noverbose .
<i>channel</i>	Specifies an MTA channel area to be searched for matching messages. The * or ? wildcard characters may be used in the channel specification.

imsimta qm

The **imsimta qm** utility inspects and manipulates the channel queues and the messages contained in the queues. **imsimta qm** contains some functionality overlap with the **imsimta cache** and **imsimta counters** commands.

For example, some of the information returned by **imsimta cache -view** is also available through the **imsimta qm directory** command. However, **imsimta qm**, does not completely replace **imsimta cache** or **imsimta queue**.

Note that **imsimta qm -dir** looks through the entire directory tree by design as opposed to **imsimta qm -sum** which scans only well-defined channel queue directories.

You must be **root** or **mailsrv** to run **imsimta qm**.

imsimta qm can be run in an interactive or non-interactive mode. To run **imsimta qm** in the interactive mode, enter:

```
imsimta qm
```

You can then enter the sub-commands that are available for use in the interactive mode. To exit out of the interactive mode, enter **exit** or **quit**.

To run **imsimta qm** in the non-interactive mode, enter:

```
imsimta qm <sub-commands> [<options>]
```

Note that some of the sub-commands available in the interactive mode are not available in the non-interactive mode, and vice versa. Refer to the following sub-commands for availability in each mode.

Sub-Commands

clean

The **clean** sub-command holds or deletes message files containing specific substrings in their envelope From: address, Subject: line, or content.

Available in both interactive and non-interactive modes.

```
clean [-content=<substring>] [-from=<substring>] [-subject=<substring>]
      [-to=<substring>] [-domain_to=<substring>]
      [-database | -directory_tree] [-delete | -hold] [-ignore_zz]
      [-match=<keyword>] [-min_length=<n>] [-threads=<n> | -nothreads]
      [-verbose | -noverbose] [<channel>]
```

The options for this sub-command are:

Table 8–12 clean Sub-Command Options

Option	Description
-content=substring-from=substring-subject=substring-to=substring-domain_to=substring	Specifies the substrings for which to search. Any combination of each option may be used. However, only one of each may only be used. When a combination of such options is used, the -match option controls whether the options are interpreted as further restrictions (-match=AND), or as alternatives (-match=OR).The -domain_to option scans for frequently occurring envelope To: addresses. Identical to the -to option, except -domain_to looks at only the host.domain portion of the envelope To: address.The -from option can take an empty address using, for example, imsimta qm clean -from=\<\> .
-database -directory_tree	Controls whether the message files searched are only those with entries in the queue cache (-database) or all message files actually present in the channel queue directory tree (-directory_tree). When neither -database nor -directory_tree is specified, then the view selected with the view sub-command will be used. If no view sub-command has been issued, then -directory_tree is assumed.
-delete -hold	Specifies whether matching message files are held (-hold) or deleted (-delete). The -hold option is the default.
-ignore_zz	Ignores queued message files with file names beginning with "ZZ". This option may be used to scan only those message files which represent queued messages which have failed at least one delivery attempt.

Table 8–12 (Cont.) clean Sub-Command Options

Option	Description
<code>-match=keyword</code>	Controls whether a message file must contain all (<code>-match=AND</code>) or only one of (<code>-match=OR</code>) the specified substrings in order to be held or deleted. The substrings are specified by the <code>-content</code> , <code>-env_from</code> , and <code>-subject</code> options. The default is <code>-match=AND</code> .
<code>-min_length=n</code>	Overrides the length limit for each substring to be searched. By default, the limit is 24 bytes (<code>-min_length=24</code>).
<code>-threads=n</code> <code>-nothreads</code>	Accelerates the searching on multiprocessor systems by dividing the work amongst multiple, simultaneous running threads. To run <i>n</i> simultaneous searching threads, specify <code>-threads=n</code> . The value <i>n</i> must be an integer between 1 and 8. The default is <code>-nothreads</code> .
<code>-verbose</code> <code>-noverbose</code>	Requests that the utility displays operation information (<code>-verbose</code>). The default is <code>-noverbose</code> .
<code>channel</code>	Specifies a specific MTA channel area to be searched for matching messages. The * or ? wildcard characters may be used in the channel specification.

counters clear

The **counters clear** sub-command performs the following operations:

1. Creates the shared memory segment for the channel message and association counters if the segment does not already exist.
2. Sets all counter values to zero.
3. When `-channels` is specified, sets the counts of stored messages, recipients, and volume from the queue cache database.

Available for both interactive and non-interactive modes.

```
counters clear [-channels] [-associations]
```

The options for this sub-command are:

Table 8–13 counters clear Sub-Command Options

Option	Description
<code>-channels</code>	Clears the message counters
<code>-associations</code>	Clears the association counters

When neither option is specified, both are assumed. When `-associations` is specified and `-channels` is not specified, step (3) is not performed.

counters create

The **counters create** sub-command performs the following operations:

1. Creates the shared memory segment for the channel message and association counters if the segment does not already exist.
2. Sets the counts of stored messages, recipients, and volume from the queue cache database.

Available for both interactive and non-interactive modes.

```
counters create [-max_channels=<n>]
```

The option for this sub-command is:

Table 8–14 counters create Sub-Command Options

Option	Description
-max_channels=<i>n</i>	Tells the MTA how many channels to allow for in the memory segment. If this option is omitted, then the MTA looks at the imta.cnf file and determines a value on its own.

counters delete

The **counters delete** sub-command deletes the shared memory segment used for channel message and association counters. Note that active MTA server processes and channels will likely recreate the memory segment.

Available for both interactive and non-interactive modes.

```
counters delete
```

counters show

Use the **counters show** sub-command to display channel message counters. When the optional *channel-name* option is omitted, * (wildcard) is assumed and the message counters for all channels are displayed. The *channel-name* option may contain the * and ? wildcard characters.

The **counters show** sub-command performs the following operations:

1. Creates the shared memory segment for the channel message and associated counters if the segment does not already exist.
2. Sets the counts of stored messages, recipients, and volume from the queue cache database.
3. Displays the message counters for the specified channels.

Available for both interactive and non-interactive modes.

```
counters show [-headers] [-noheaders] [-output=<file-spec>] [<channel-name>]
```

The options for this sub-command are:

Table 8–15 counters show Sub-Command Options

Option	Description
-headers or -noheaders	Controls whether or not a heading is displayed. The -headers option is the default.
-output=<i>file_spec</i>	Causes the output to be written to a file. Any existing file with the same name as the output file is overwritten.

date

Displays the current time and date in RFC 822, 1123 format.

Available for both interactive and non-interactive modes.

```
date
```

delete

Deletes the specified messages displayed in the most recently generated message queue listing.

```
delete [-channel=<name> [-all]] [-confirm | -noconfirm]
      [-log | -nolog] [<id>...]
```

The **id** option specifies the messages to be deleted.

See "imsimta qm Options" for information on using the **-channel**, **-all**, **-confirm**, and **-log** options.

Available only in interactive mode.

directory

Generates a listing of queued message files. By default, the **imta/queue** directory tree is used as the source of queued message information; this default may be changed with the **view** sub-command. The **-database** and **-directory_tree** options may be used to override the default.

Available for both interactive and non-interactive modes.

```
directory [-held | -noheld] [-database] [-directory_tree]
      [-envelope] [-owner=<username>] [-from=<address>] [-to=<address>]
      [-match=<bool>] [-file_info | -nofile_info] [-total | -nototal]
      [<channel-name>]
```

The options for this sub-command are:

Table 8–16 *directory Sub-Command Options*

Option	Description
-database	Obtains message information from the Job Controller.
-directory_tree	Selects the on-disk directory tree as the source of message information.
-envelope	Generates a listing which also contains envelope address information.
-total -nototal	Generates size and count totals across all selected channels.
-owner=username	Lists only those messages owned by a particular user. Messages enqueued by a local user will be owned by that user; most other messages will be owned by mailsrv . Use of the -owner option implies -database .
-from=address and -to=address and -match=bool	Lists only those messages with envelope From: or To: addresses matching the specified address. When both -from and -to are specified, a message is listed if either its envelope From: or To: addresses match the specified addresses. This corresponds to the -match=or option. Specify -match=and to list only messages matching both the specified From: and To: addresses. Use of -from or -to implies -envelope . <i>address</i> can include a wild card (*) that matches a sequence of characters or a % character that matches a single character.
-held -noheld	By default, active messages are listed. Specify -held to instead list messages which have been marked as held. Note that -held implies -directory_tree .

Table 8–16 (Cont.) directory Sub-Command Options

Option	Description
-file_info -nofile_info	When the directory tree is scanned, each message file is accessed to determine its size as measured in units of blocks (normally 1024 bytes). To suppress this behavior and speed up generation of the listing, specify -nofile_info . When the queue cache database is used, the -nofile_info option is ignored as the size information is stored in the database.
<i>channel-name</i>	Restricts the listing to one or more channels. If the <i>channel-name</i> option is omitted, a listing is made for all channels. The channel name option may contain the * and ? wildcard characters.

exit

Exits the **imsimta qm** utility. Synonymous with the **quit** sub-command.

Available for both interactive and non-interactive modes.

```
exit
```

held

Generates a listing of message files which have been marked as held. This listing is always generated from the **imta/queue/** directory tree.

Available for both interactive and non-interactive modes.

```
held [-envelope] [-file_info | -nofile_info] [-total | -nototal]
    [-from=<address>] [-to=<address>] [-match=<bool>] [<channel-name>]
```

The options for this sub-command are:

Table 8–17 held Sub-Command Options

Option	Description
-envelope	Generates a listing which also contains envelope address information.
-total -nototal	Generate size and count totals across all selected channels.
-from=address and -to=address and -match=bool	Lists only those messages with envelope From: or To: addresses matching the specified address. When both -from and -to are specified, a message is listed if either its envelope From: or To: addresses match the specified addresses. This corresponds to the -match=or option. Specify -match=and to list only messages matching both the specified From: and To: addresses. Use of -from or -to implies -envelope .
-file_info -nofile_info	When the directory tree is scanned, each message file is opened to determine its size as measured in units of blocks (normally 1024 bytes). To suppress this behavior and speed up generation of the listing, specify -nofile_info .
<i>channel-name</i>	Restricts the listing to one or more channels. If the <i>channel-name</i> option is omitted, a listing is made for all channels. The <i>channel-name</i> option may contain the * and ? wildcard characters.

history

Displays any delivery history information for the specified messages from the most recently generated message queue listing.

Available only in interactive mode.

```
history [-channel=<name> [-all] ] [-confirm | -noconfirm] [<id>...]
```

Use the **id** option to specify the messages whose history is displayed.

See "[imsimta qm Options](#)" for information on using the **-channel**, **-all**, and **-confirm** options.

hold

Marks as held the specified messages from the most recently generated message queue listing

Available only in interactive mode.

```
hold [-channel=<name> [-all]] [-confirm | -noconfirm]
      [-log | -nolog] [<id>...]
```

Use the **id** option to specify the messages to mark as held.

See "[imsimta qm Options](#)" for information on the **-channel**, **-all**, **-confirm**, and **-log** options.

jobs

The **imsimta qm jobs** utility displays what messages are being processed by what jobs for what channels.

```
jobs
```

Example of output:

```
channel <channel name>
  job <pid>
    host <host name>
    host <host name>
    <count of hosts> HOST BEING PROCESSED BY JOB <pid>
    message <subdir/message name>
    message <subdir/message name>
    processed messages: <# messages successfully dequeued>
    failed processing attempts: <#messages reenqueued>
    <count of messages> MESSAGES BEING PROCESSES BY JOB <pid>
    <count of jobs> JOBS ACTIVE FOR CHANNEL foo
    <count of active channels> ACTIVE CHANNELS
```

quit

Exits the **imsimta qm** utility. Synonymous with the *exit* sub-command.

Available in both interactive and non-interactive modes.

```
quit
```

read

Displays the specified messages from the most recently generated message queue listing.

Available only in interactive mode.

```
read [-content | -nocontent] [-channel=<name> [-all]]
      [-confirm | -noconfirm] [<id>...]
```

The options for this sub-command are:

Table 8–18 read Sub-Command Options

Option	Description
-content -nocontent	Displays (-content) or suppresses display (-nocontent) of message content along with the envelope and header information . {-nocontent is the default.
<i>id</i>	Specifies the messages to display.

See "[imsimta qm Options](#)" for information on using the **-channel**, **-all**, and **-confirm** options.

release

If the specified message file is marked as held, it is renamed to remove the hold mark. The Job Controller, if running, is informed that the message is to be processed immediately, ahead of all other messages.

Available only in interactive mode.

```
release [-channel=<name> [-all]] [-confirm | -noconfirm]
        [-log | -nolog] [<id>...]
```

Use the *id* option to specify the messages to release from *.HELD* status.

Note: Run the **dir** subcommand prior to running **release**. For example:

```
$ imsimta qm
qm maint> dir -held
qm maint> release 1
```

See "[imsimta qm Options](#)" for information on using the **-channel**, **-all**, **-confirm**, and **-log** options.

return

Returns as undelivered the specified messages shown in the most recently generated message queue listing.

Available only in interactive mode.

```
return [-channel=<name> [-all]] [-confirm | -noconfirm]
        [-log | -nolog] [<id>...]
```

Use the **id** option to specify the messages to return.

See "[imsimta qm Options](#)" for information on using the **-channel**, **-all**, **-confirm**, and **-log** options.

run

Processes, line-by-line, the commands specified in a file.

Available in both interactive and non-interactive modes.

```
run [-ignore | -noignore] [-log | -nolog] <file-spec>
```

Specifically, **file-spec** is opened and each line from it is read and executed.

The options for this sub-command are:

Table 8–19 *run Sub-Command Options*

Option	Description
-ignore -noignore	Unless -ignore is specified, command execution will be aborted should one of the sub-commands generate an error.
-log -nolog	By default, each command is echoed to the terminal before being executed (the -log option). Specify -nolog to suppress this echo.

start

Restart processing of messages enqueued for the specified channel. The Job Controller not only marks the channel as "okay" to process, but it additionally starts processing jobs for the channel. This command takes effect whether the Job Controller is running or not.

```
start <channel>
```

The *channel* option specifies the channel to restart.

Note: The command **imsimta qm start/stop channel** might fail if run simultaneously for many channels at the same time. The tool might have trouble updating the `hold_list` and could report: "QM-E-NOTSTOPPED, unable to stop the channel; cannot update the hold list." **imsimta qm start/stop channel** should only be used sequentially with a few seconds interval between each run. If you only want the channel to run between certain hours, use the following options in the channel definition section in the job controller configuration file:

```
urgent_delivery=08:00-20:00
normal_delivery=08:00-20:00
nonurgent_delivery=08:00-20:00
```

stop

Stops processing of messages enqueued for the specified channel. This command prevents you from having to stop the Job Controller and recompiling the configuration. The channel does not process messages until a *start* command is issued for that channel. This state persists across restarts of the Job Controller, the Messaging Server, and the host computer itself. This command takes effect whether the Job Controller is running or not.

```
stop <channel>
```

The *channel* option specifies the channel to stop.

Note: The command **imsimta qm start/stop channel** might fail if run simultaneously for many channels at the same time. The tool might have trouble updating the `hold_list` and could report: "QM-E-NOTSTOPPED, unable to stop the channel; cannot update the hold list." **imsimta qm start/stop channel** should only be used sequentially with a few seconds interval between each run. If you only want the channel to run between certain hours, use the following options in the channel definition section in the job controller configuration file:

```
urgent_delivery=08:00-20:00
normal_delivery=08:00-20:00
nonurgent_delivery=08:00-20:00
```

summarize

The **summarize** sub-command displays a summary listing of message files.

```
summarize [-database | -directory_tree] [-heading | -noheading]
          [-held | -noheld] [-trailing | -notrailing]
```

The options for this sub-command are:

Table 8–20 summarize Sub-Command Options

Option	Description
-database -directory_tree	Controls whether the information presented is obtained from the Job Controller (-database) or by looking at the actual directory tree containing the channel queues (-directory_tree). When neither -database nor -directory_tree is specified, then the "view" selected with the view sub-command will be used. If no view sub-command has been issued, then -directory_tree is assumed.
-heading -noheading	Controls whether or not a heading line describing each column of output is displayed at the start of the summary listing. The -heading option is the default.
-held -noheld	Controls whether or not to include counts of .HELD messages in the output. The -noheld option is the default.
-trailing -notrailing	Controls whether or not a trailing line with totals is displayed at the end of the summary. The -trailing option is the default.

top

The **top** sub-command displays the most frequently occurring envelope From:, Subject:, or message content fields found in message files in the channel queues. When used in conjunction with the **clean** sub-command, **top** may be used to locate unsolicited bulk email in the query and hold or delete it.

```
top [-content[=<range>]] [-from[=<range>]] [-subject[=<range>]]
    [-to[=<range>]] [-database | -directory_tree] [-domain_to[=<range>]]
    [-held] [-ignore_zz] [-min_count=<n>] [-threads=<n> | -nothreads]
    [-top=<n>] [-verbose | -noverbose] [<channel>]
```

The options for this sub-command are:

Table 8–21 top Sub-Command Options

Option	Description
-content [= <i>range</i>] -from [= <i>range</i>] -subject [= <i>range</i>] -to [= <i>range</i>] -domain_ to [= <i>range</i>]	The -content , -from , -subject , and -to options are used to specify which frequently occurring fields should be displayed. By default, only Subject: fields are shown (-subject). Use -from to display frequent envelope From: fields, -to to display frequent envelope To: fields, or -content to display frequent message contents. Use -domain_to to display frequently occurring envelope To: addresses. Identical to -to option, except -domain_to looks at only the host.domain portion of the envelope To: address. Any combination of -content , -from , -to , -domain_to , and -subject may be specified. However, only one of each may be used. The -content , -from , -to , -domain_to , and -subject options accept the optional parameters START=<i>n</i> and LENGTH=<i>n</i> . These options indicate the starting position and number of bytes in the field to consider. The defaults are -content=(START=1,LENGTH=256) , -from=(START=1,LENGTH=2147483647) , -to=(START=1,LENGTH=2147483647) , -subject=(START=1,LENGTH=2147483647) , and -domain_to=(START=1,LENGTH=214783647) . Use of these options is useful when, for example, trying to identify occurrences of a spam message which uses random text at the start of the Subject: line.
-database -directory_ tree	Controls whether the message files scanned are only those with entries in the queue cache database (-database) or all message files actually present in the channel queue directory tree (-directory_tree). When neither -database nor -directory_tree is specified, then the "view" selected with the view sub-command will be used. If no view sub-command has been issued, then -directory_tree is assumed.
-held	Lists only the files which have a .HELD extension.
-ignore_zz	Ignores queued message files with file names beginning with "ZZ." This option may be used to scan only those message files which represent queued messages which have failed at least one delivery attempt.
-min_count=<i>n</i>	Changes the minimum number of times that a string must occur in order to be displayed. The default is -min_count=2 .
-threads=<i>n</i> -nothreads	Accelerates searching on multiprocessor systems by dividing the work amongst multiple, simultaneously running threads. To run <i>n</i> simultaneous searching threads, specify -threads=<i>n</i> . The value <i>n</i> must be an integer between 1 and 8. The default is -nothreads .
-top=<i>n</i>	Changes the amount of most frequently occurring fields that are displayed. The default is -top=20 .
-verbose -noverbose	Requests that the utility displays operation information (-verbose). The default is -noverbose .
<i>channel</i>	Specifies an MTA channel area to be scanned for string frequencies. The * or ? wildcard characters may be used in the channel specification.

view

Specifies the source of queued message information for subsequent directory commands.

Available only in interactive mode.

```
view -database | -directory_tree
```

By default, queued message listings are generated by scanning the **imta/queue/** directory tree. This corresponds to the **-directory_tree** option. You can, alternatively, generate the listings from the MTA queue cache database by issuing the **-database** option.

Settings made with the **view** sub-command remain the default until either another view command is issued or the utility exits. The default may be overridden with the **-database** or **-directory_tree** options of the directory command.

Note that the directory tree is always used when listing held message files.

version

Diagnostic command which outputs the Messaging Server version information as well as the compiled date and time for the **qm** program being run.

Available only in interactive mode.

```
version
```

imsimta qm Options

The **delete**, **history**, **hold**, **read**, **release**, and **return** sub-commands all support the following options and parameter:

Table 8–22 *imsimta qm Command Options*

Option	Description
-channel= <i>name</i>	Operates on the specified channel.
-all	The -all option may be used to operate on all of the previously listed messages. When used in conjunction with the -channel option, only those previously listed messages for the specified channel are operated on. The -all option may not be used in conjunction with an id option. However, -all or at least one id option must be specified.
-confirm and -noconfirm	-confirm prompts you to confirm each message release operation when the id option is not used to explicitly select messages. This prevents accidental delete -all sub-commands from being executed. The -noconfirm option suppresses this prompt.
-log and -nolog	Controls whether or not the operation on each selected message is reported.
<i>id</i>	The identification number of a message shown in the most recent listing generated by either the directory or the held sub-command. The identification number for a message is the integer value displayed in the left-most column of the listing. The <i>id</i> can also be a range or comma-separated list.

These options identify the messages to which the command is applied. When none of the options are specified, at least one *id* option must be supplied.

For example, in the following listing the first message displayed has an identification number of 1 and the second 2:

```
qm.maint> directory tcp_local
Sat, 26 Dec 2009 13:31:27 +1100 (EST)
Data gathered from the queue directory tree
```

```

Channel: tcp_local                Size Queued since
-----
  1 ZZg0N2j0Pzpq0.00            1 26 Dec 2009 13:31:08
  2 ZZg0N2j0Pzqq1.00            1 26 Dec 2009 13:31:12
-----
Total size:                       2

Grand total size:                  2

```

These two messages can therefore be selected by either "1,2" or "1-2".

Examples

The following sections list the non-interactive and interactive examples.

Non-Interactive Mode

The following example generates a list of queued messages:

```

bash-3.00# ./imsimta qm directory
Sat, 26 Dec 2009 13:32:26 +1100 (EST)
Data gathered from the queue directory tree

Channel: ims-ms                    Size Queued since
-----
  1 ZZg0N2j0PzEq2.00            1 26 Dec 2009 13:32:08
  2 ZZg0N2j0PzHq3.00            1 26 Dec 2009 13:32:20
-----
Total size:                       2

Channel: tcp_local                Size Queued since
-----
  3 ZZg0N2j0Pzpq0.00            1 26 Dec 2009 13:31:08
-----
Total size:                       1

Grand total size:                  3

```

Interactive Mode

In the following interactive session, the *directory* sub-command is used to obtain a list of queued messages. The **delete** sub-command is then used to delete the first of the displayed messages. Finally, another **directory** sub-command is issued that displays that the deleted message is indeed gone.

```

bash-3.00# ./imsimta qm
qm.maint> directory
Sat, 26 Dec 2009 13:32:47 +1100 (EST)
Data gathered from the queue directory tree

Channel: ims-ms                    Size Queued since
-----
  1 ZZg0N2j0PzEq2.00            1 26 Dec 2009 13:32:08
  2 ZZg0N2j0PzHq3.00            1 26 Dec 2009 13:32:20
-----
Total size:                       2

Channel: tcp_local                Size Queued since
-----
  3 ZZg0N2j0Pzpq0.00            1 26 Dec 2009 13:31:08
-----

```

```

Total size:                               1

Grand total size:                          3
qm.maint> delete 1
%QM-I-DELETED, deleted the message file
/opt/sun/comms/messaging/data/queue/ims-ms/000/ZZg0N2j0PzEq2.00
qm.maint> directory
Sat, 26 Dec 2009 13:32:53 +1100 (EST)
Data gathered from the queue directory tree

Channel: ims-ms                            Size Queued since
-----
      1 ZZg0N2j0PzHq3.00                    1 26 Dec 2009 13:32:20
-----
Total size:                               1

Channel: tcp_local                          Size Queued since
-----
      2 ZZg0N2j0Pzpq0.00                    1 26 Dec 2009 13:31:08
-----
Total size:                               1

Grand total size:                          2

```

imsimta qtop

The **imsimta qtop** utility displays the most frequently occurring envelope From:, To:, Subject:, or message content fields found in message files in the channel queues.

Syntax

```

imsimta qtop [-content[=<range>]] [-from[=<range>]] [-subject[=<range>]]
  [-to[=<range>]] [-domain_to[=<range>]] [-database | -directory_tree]
  [-ignore_zz] [-min_count=<n>] [-threads=<n> | -nothreads] [-top=<n>]
  [-verbose | -noverbose] [<channel>]

```

Options

The options for this command are:

Table 8-23 *imsimta qtop* Command Options

Option	Description
-content[= <i>range</i>]	Specifies which frequently occurring fields should be displayed. By default, only Subject: fields are shown (-subject). Specify -from to display frequent envelope From: fields, -to to display frequent envelope To: fields, or -content to display frequent message contents. Specify -domain_to to display frequently occurring envelope To: fields. Identical to -to option, except -domain_to looks at only the host.domain portion of the envelope To: address. Any combination may be specified. However, only one of each may be used. These options accept the START= <i>n</i> and LENGTH= <i>n</i> arguments. These arguments indicate the starting offset and number of bytes in the field to consider. The defaults are -content=(START=1,LENGTH=256), -from=(START=1,LENGTH=2147483647), -subject=(START=1,LENGTH=2147483647), and -domain_to=(START=1,LENGTH=2147483647).
-from[= <i>range</i>]	
-subject[= <i>range</i>]	
-to[= <i>range</i>]	
-domain_to[= <i>range</i>]	

Table 8–23 (Cont.) *imsimta qtop* Command Options

Option	Description
-database	Specifies that only message files identified by the queue cache database is searched.
-directory_tree	Searches all message files actually present in the channel queue directory tree.
-ignore_zz	Ignores queued message files with file name beginning with "ZZ." This option may be used to scan only those message files which represent queued messages which have failed at least one delivery attempt. For example, the following command indicates to which domains the MTA has problems delivering messages: imsimta qtop -ignore_zz -domain_to
-min_count=<i>n</i>	Changes the minimum number of times that a string must occur in order to be displayed. The default is -min_count=2 .
-threads=<i>n</i> -nothreads	Accelerates searching on multiprocessor systems by dividing the work amongst multiple, simultaneously running threads. To run <i>n</i> simultaneous searching threads, specify -threads=<i>n</i> . The value <i>n</i> must be an integer between 1 and 8. The default is -nothreads .
-top=<i>n</i>	Changes the amount of most frequently occurring fields that are displayed. The default is -top=20 .
-verbose -noverbose	Requests that the utility displays operation information (-verbose). The default is -noverbose .
<i>channel</i>	Specifies a channel area to be scanned for string frequencies. The * and ? wildcard characters may be used in the channel specification.

imsimta refresh

The **imsimta refresh** utility performs the following functions:

- Recompiles the MTA configuration files.
- Stops any MTA Job Controller or MTA Service Dispatcher jobs that are currently running.
- Restarts the Job Controller and MTA Service Dispatcher.

Essentially, **imsimta refresh** combines the function of **imsimta cnbuild** and **imsimta restart**.

Note: You must be logged in as **root** to run **imsimta refresh**.

Note: This command has been deprecated. On a production system, the **imsimta refresh** command should only be used as a last resort. This is because the command does far more than refresh the running MTA services. Specifically, it shuts down all running MTA services, builds a new compiled configuration from the inactive, human readable configuration files, and then restarts all the MTA services with the new compiled configuration. In shutting down the MTA services, all heuristic queuing information is lost (e.g., redelivery schedules). More often than not, the command "**imsimta restart**" with the specific MTA service which needs to be restarted (for example, **smtp, dispatcher, job_controller**). When a configuration change has been made and an MTA service needs to pick up that change, then first use the "**imsimta cnbuild**" followed by the "**imsimta restart**" command. Also, strongly consider using the "**imsimta reload**" command instead of the "**imsimta restart**" command.

Syntax

```
imsimta refresh [job_controller | dispatcher]
```

Options

The options for this command are:

Table 8–24 imsimta refresh Command Options

Option	Description
job_controller	Restarts the Job Controller.
dispatcher	Restarts the MTA Service Dispatcher.

If no component name is specified, all active components are restarted.

imsimta reload

Some parts of the MTA configuration can be changed and have these changes activated without having to stop and start the system. The reloadable parts of the configuration are:

mappings

aliases

general, forward and reverse lookup tables

These can be changed, compiled, and the changes activated by issuing the commands:

imsimta cnbuild

imsimta reload

The **imsimta reload** command informs the dispatcher and job controller of the change, and they in turn inform the processes they started.

Syntax

```
imsimta reload
```

imsimta renamedb

The **imsimta renamedb** command renames an MTA database. Since the MTA may optionally reference several "live" databases, that is, databases whose presence triggers their use by the MTA, it is important, first, to ensure that the MTA does not see such a database while it is in a mixed state, and second, to minimize any period of time during which the database is inaccessible. The **imsimta crdb** command locks the database it is creating to avoid having it accessed in a mixed state.

It is therefore recommended that the MTA databases be created or updated in a two-step process:

1. Create or update a temporary database.
2. Rename the temporary database to the "live" name by using the **imsimta renamedb** command.

The **imsimta renamedb** command, which must delete any old database files and rename the new database files, locks the database during the renaming process to avoid presenting the database in a mixed state. In this way the database is never accessible while it is in a mixed state, yet any window of time during which the database is inaccessible is minimized. Renaming is generally quicker than database generation.

Syntax

```
imsimta renamedb <old-database-spec> <new-database-spec>
```

Options

The options for this command are:

Table 8–25 *imsimta renamedb Command Options*

Option	Description
<i>old-database-spec</i>	The name of the database that is being renamed.
<i>new-database-spec</i>	The new name of the database. This may either be an actual pathname, or one of the special names such as IMTA_ALIAS_DATABASE , IMTA_REVERSE_DATABASE , IMTA_GENERAL_DATABASE , or IMTA_DOMAIN_DATABASE , listed in the MTA tailor file and pointing to actual pathnames.

Example

The following command renames the database **tmpdb** to be the actual MTA alias database (usually *MessagingServer_home/data/db/aliasesdb*).

```
imsimta renamedb tmpdb IMTA_ALIAS_DATABASE
```

imsimta restart

The **imsimta restart** command stops and restarts the Job Controller and Service Dispatcher. This causes all MTA master and slave programs to be restarted. It can also restart SMTP, LMTP, and SMTP_SUBMIT.

Detached MTA processes should be restarted whenever the MTA configuration is altered. These processes load information from the configuration only once and need to be restarted in order for configuration changes to become visible to them. In addition to general MTA configuration files, such as the **imta.cnf** file, some components, such as the MTA Service Dispatcher, have their own specific

configuration files, for example, **dispatcher.cnf**, and should be restarted after changes to any of these files.

Note: You must be logged in as *root* to use this utility.

Syntax

```
imsimta restart [job_controller|dispatcher|smtp|lmtpl|smtp_submit]
```

Restarting the MTA Service Dispatcher effectively restarts all the service components it handles. If no component name is given, all active components are restarted.

Example

To restart the MTA Job Controller and channel master programs:

```
imsimta restart job_controller
```

imsimta return

The **imsimta return** command returns a message to the message's originator. The returned message is a single multipart message with two parts. The first part explains the reason why the message is being returned. The text of the reason is contained in the file *return_bounce.txt* located in the *MessagingServer_home/config/locale/C/LC_MESSAGES* directory. The second part of the returned message contains the original message.

Syntax

```
imsimta return <message-file>
```

message-file is the name of the message file to return. The name may include wildcards, but if so, the specification must be quoted.

Example

The following command causes the specified message to be returned to its originators.

```
imsimta return /imta/queue/1/ZZ0FRW00A03G2EUS.00
```

imsimta run

The **imsimta run** command processes the messages in the channel specified by the channel option. Output during processing is displayed at your terminal, which makes your terminal unavailable for the duration of the operation of the utility. Refer also to the "**imsimta submit**" command which, unlike **imsimta run**, does not monopolize your terminal.

Note that a channel delivery program that is run using this command, unlike the **imsimta submit** command, attempts to deliver messages before any pending backoff delay has expired.

Syntax

```
imsimta run <channel>
```

Options

The option for this command is:

Table 8–26 *imsimta run Command Option*

Option	Description
<i>channel</i>	Specifies the channel to be processed. This option is mandatory.

Example

Type the following command to process any messages in the `tcp_local` channel:

```
imsimta run tcp_local
```

imsimta shutdown

The **imsimta shutdown** command shuts down the MTA Job Controller and the MTA Dispatcher. Shutting down the MTA Dispatcher shuts down all services (for example, SMTP) being handled by the Dispatcher. It can also be used to stop the SMTP, LMTP, SMTP_SUBMIT servers. Note that you can only restart a Dispatcher service that is currently running. If you do **imsimta shutdown smtp**, you must restart the Dispatcher to start the SMTP service again.

Note: You must be logged in as **root** to use this utility.

Syntax

```
imsimta shutdown [dispatcher|job_controller|smtp|smtp_submit|lmtp]
```

Example

Use the following command to shut down the MTA jobs:

```
imsimta shutdown
```

imsimta start

The **imsimta start** command starts up detached MTA processes. If no component option is specified, then the MTA Job Controller and MTA Service Dispatcher are started. Starting the Service Dispatcher starts all services the Service Dispatcher is configured to handle, which usually includes the SMTP server.

The services handled by the MTA Service Dispatcher must be started by starting the MTA Service Dispatcher. Only services not being handled by the MTA Service Dispatcher can be individually started via the **imsimta start** command. The Service Dispatcher may be configured to handle various services, for example, the multithreaded SMTP server.

Note: You must be logged in as *root* to use this utility.

Syntax

```
imsimta start [<component>]
```

If a component option is specified, then only detached processes associated with that component are started. The standard component names are:

- **dispatcher** - Multithreaded Service Dispatcher.
- **job_controller** - Schedules deliveries (dequeues messages).

Example

Use the following command to start the MTA Job Controller and MTA Service Dispatcher:

```
imsimta start
```

imsimta stop

The **imsimta stop** command shuts down the MTA Job Controller and the MTA Dispatcher. Shutting down the MTA Dispatcher shuts down all services (for example, SMTP) being handled by the Dispatcher. It can also be used to stop the SMTP, LMTP, SMTP_SUBMIT servers. Note that you can only restart a Dispatcher service that is currently running. If you do **imsimta shutdown smtp**, you must restart the Dispatcher to start the SMTP service again.

Note: You must be logged in as *root* to use this utility.

Syntax

```
imsimta stop [dispatcher|job_controller|smtp|smtp_submit|lmtp]
```

Example

Use the following command to shut down the MTA jobs:

```
imsimta stop
```

imsimta submit

The **imsimta submit** command directs the Job Controller to fork a process to execute the messages queued to the channel specified by the channel option.

Syntax

```
imsimta submit [<channel>] [poll]
```

Options

The options for this command are:

Table 8–27 *imsimta submit Command Options*

Option	Description
<i>channel</i>	Specifies the channel to be processed. The default, if this option is not specified, is the local channel 1 .
poll	If poll is specified, the channel program runs even if there are no messages queued to the channel for processing.

Example

Use the following command to process any messages in the **tcp_local** channel:

```
imsimta submit tcp_local
```

imsimta test

The **imsimta test** utilities perform tests on various areas of functionality of the MTA.

The **imsimta test -domain** utility is an interactive command with sub-commands that differ from the options common to the other **imsimta test** utilities. For information about **imsimta test -domain**, see "[imsimta test -domain](#)."

imsimta test -mapping

imsimta test -mapping tests the behavior of a mapping table in the mapping file. The result of mapping an input string will be output along with information about any meta characters specified in the output string.

If an input string is supplied on the command line, then only the result of mapping that input string will be output. If no input string is specified, **imsimta test -mapping** will enter a loop, prompting for an input string, mapping that string, and prompting again for another input string. **imsimta test -mapping** will exit when a CTRL-D is entered.

imsimta test -match

imsimta test -match tests a mapping pattern in order to test wildcard and global matching.

imsimta test -match prompts for a pattern and then for a target string to compare against the pattern. The output indicates whether or not the target string matched. If a match was made, the characters in the target string that matched each wildcard of the pattern is displayed. The **imsimta test -match** utility loops, prompting for input until the utility is exited with a CTRL-D.

imsimta test -rewrite

imsimta test -rewrite provides a test facility for examining the MTA's address rewriting and channel mapping process without actually sending a message. Various qualifiers can be used to control whether **imsimta test -rewrite** uses the configuration text files or the compiled configuration (if present), the amount of output produced, and so on.

If a test address is specified on the command line, **imsimta test -rewrite** applies the MTA address rewriting to that address, reports the results, and exits. If no test address is specified, **imsimta test -rewrite** enters a loop, prompting for an address, rewriting it, and prompting again for another address. **imsimta test -rewrite** exits when CTRL-D is entered.

When testing an email address corresponding to a restricted distribution list, **imsimta test -rewrite** uses as the posting address the return address of the local postmaster, which is usually **postmaster@localhost** unless specified by the MTA option **RETURN_ADDRESS** in the MTA Option file.

imsimta test -url

imsimta test -url tests an LDAP query URL. Note that the LDAP server to query is controlled by the setting of the MTA option **LDAP_SERVER** in **local.conf**.

Syntax

```

imsimta test -rewrite [-alias_file=<filename>]
  [-channel | -nochannel]
  [-check_expansions | -nocheck_expansions]
  [-configuration_file=<filename> ] [-database=<database_list>]
  [-debug | -nodebug] [-delivery_receipt | -nodelivery_receipt]
  [-destination_channel=<channel>] [-expandlimit=<integer>]
  [-extra_local_channel=<channel>] [-filter | -nofilter]
  [-from=_address_ | -nofrom] [-image_file=<filename> | -noimage_file]
  [-input=<input-file>] [-local_alias=<value> | -nolocal_alias]
  [-mapping_file=<file> | -nomapping_file]
  [-option_file=<filename> | -nooption_file] [-output=<output-file>]
  [-read_receipt | -noread_receipt] [-noreprocess] [-restricted=<setting>]
  [-sender=_from_address_] [-source_channel=<channel>]
  [-soptin] [-spares]

imsimta test -mapping [<input_string>] [-debug | -nodebug]
  [-flags=<chars> | -noflags]
  [-image_file=<filename> | -noimage_file] [-mapping_file=<filename>]
  [-option_file=<filename> | -nooption_file] [-table=<table-name>] [<address>]

imsimta test -match

imsimta test -url [-debug | -nodebug] [<ldap_url>]

imsimta test -exp -mm -message=<message-file>[-block] [-input=<input-file>]
  [-output=<output-file>][-sender=<authenticated-sender>] [-username=<username>]

```

Options

The options for this command are:

Table 8–28 *imsimta test Command Options*

Option	Description
<i>address</i>	Specifies the test address to be rewritten. If this option is omitted, then the command prompts for an address. Used with the -rewrite option.
<i>input_string</i>	The string to be matched in the left side of a mapping table. Used with the -mapping option.
ldap_url	The LDAP URL that imsimta test -url attempts to resolve.
-alias_file=filename	Specifies an alternate alias file for imsimta test -rewrite to use. imsimta test -rewrite normally consults the default alias file named by the IMTA_ALIAS_FILE option of the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> , during the rewriting process. This option has no effect unless -noimage_file is specified or no compiled configuration exists; any compiled configuration precludes reading any sort of alias file.
-block	Treats the entire input as a single sieve script. The default is to treat each line as a separate script.
-channel -nochannel	Controls whether imsimta test -rewrite outputs detailed information regarding the channel an address matches (for example, channel flags).

Table 8–28 (Cont.) *imsimta test* Command Options

Option	Description
-check_expansions -nocheck_expansions	Controls checking of alias address expansion. Normally the MTA considers the expansion of an alias to have been successful if any of the addresses to which the alias expands are legal. The -check_expansions option causes a much stricter policy to be applied: imsimta test -rewrite -check_expansions checks each expanded address in detail and reports a list of any addresses, expanded or otherwise, that fail to rewrite properly.
-configuration_file= <i>file</i>	Specifies an alternate file to use in place of the file named by IMTA_CONFIG_FILE . Normally, imsimta test -rewrite consults the default configuration file named by the IMTA_CONFIG_FILE option of the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> , during the rewriting process. This option has no effect unless -noimage_file is specified or no compiled configuration exists; use of any compiled configuration precludes reading any sort of configuration file.
-database= <i>database-list</i>	Disables references to various databases or redirects the database paths to nonstandard locations. imsimta test -rewrite normally consults the usual MTA databases during its operation. The allowed list items are alias , noalias , domain , nodomain , general , nogeneral , reverse , and noreverse . The list items beginning with "no" disable use of the corresponding database. The remaining items require an associated value, which is taken to be the name of that database.
-debug -nodebug	Enables the production of the additional, detailed explanations of the rewriting process. This option is disabled by default.
-delivery_receipt -nodelivery_receipt	Sets the corresponding receipt request flags. These options can be useful when testing the handling of sent or received receipt requests when rewriting forwarded addresses or mailing lists.
-destination_channel= <i>channel</i>	Controls to which destination or target channel imsimta test -rewrite rewrites addresses. Some address rewriting is destination channel specific; imsimta test -rewrite normally pretends that its channel destination is the local channel 1.
-envid	Specifies the envelope identifier. For example, a@b .

Table 8–28 (Cont.) *imsimta test* Command Options

Option	Description
-exp	<p>imsimta test -exp tests Sieve language statements against a specified RFC2822 message and sends the results of the filter to standard output. The syntax is as follows: <code>imsimta test -exp -mm -block -input=Sieve_language_scriptfile -message=rfc2822_message_filewhere</code>, <code>-block</code> treats the entire input as a single Sieve script. The default is to treat each line as a separate script and to evaluate it separately. The Sieve will only be evaluated once the end of file is reached.</p> <p><code>-input=Sieve_file</code> is a file containing the Sieve script. The default is to read the test script lines or script block from stdin.</p> <p><code>-message=message_file</code> is a text file containing the RFC 2822 message you want to test your Sieve script against. This has to be an RFC 2822 message only. It cannot be a queue file (not a <code>zz*.00</code> file). Once activated, this command reads script information, evaluates it in the context of the test message, and writes out the result. The result shows what actions would be taken as well as the result of evaluating the final statement in the script. Additional useful qualifiers are: <code>-from=address</code> specifies the envelope from address to be used in envelope tests. The default is to use the value specified by the <code>RETURN_ADDRESS</code> MTA option. <code>-output=file</code> writes results to file. The default is to write the results of script evaluation to stdout.</p> <p><code>-sender=<authenticated-sender></code> sets the authenticated sender address.</p> <p><code>-username=<username></code> sets the authentication username. Note that <code>-sender</code> and <code>-username</code> qualifiers simply simulate the setting of these fields. They do not actually perform any sort of authentication operation.</p>
-expandlimit=integer	The integer value is used to set the internal alias expansion limit within the MTA. The switch is intended to be used for testing expansion limit interactions with mailing lists and other MTA facilities.
-extra_local_channel=channel	(New in MS7u2) Tell the rewriting machinery that "channel" should have local channel semantics. This is useful in testing rewriting associated with process and conversion channel hops.
-filter -nofilter	Outputs any filters that are applied for the specified address.
-from=address -nofrom	Controls what envelope From: address is used for access control probes when the -from option is specified. If <i>address</i> is omitted, the postmaster return address is used for such probes. If the -nofrom option is specified, the MTA uses an empty envelope From: address for access probes.
-flags=chars -noflags	Specifies particular flags to be set during the mapping test when the -flags option is specified. For example, <i>chars</i> can be E (envelope), B (header/body), or I (message id) when testing a REVERSE mapping. Used with the -mapping option only.
-image_file=[filename] -noimage_file	The -noimage_file option instructs the command to unconditionally ignore any previously compiled configuration and to read configuration information from the various text files instead. When the -image_file option is specified without an optional file name, the compiled configuration is loaded from the file named by the <code>IMTA_CONFIG_DATA</code> option into the MTA tailor file, <code>MessagingServer_home/config/imta_tailor</code> , which is usually <code>MessagingServer_home/config/imta.cnf</code> . If, instead, a file name is specified, then the compiled configuration is loaded from the specified file.

Table 8–28 (Cont.) *imsimta test* Command Options

Option	Description
-input = <i>input-file</i>	Specifies a source for input. By default, imsimta test takes input from stdin .
-local_alias = <i>value</i> -nolocal_alias	Controls the setting of an alias for the local host. The MTA supports multiple "identities" for the local host; the local host may have a different identity on each channel. This option may be used to set the local host alias to the specified value; appearances of the local host in rewritten addresses are replaced by this value.
-mapping_file = <i>file</i> -nomapping_file	Instructs the command to use the specified mapping file rather than the default mapping file named by the IMTA_MAPPING_FILE option in the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> , which is usually the file named <i>MessagingServer_home/config/mappings</i> . This option has no effect unless -noimage_file is specified or no compiled configuration exists; use of any compiled configuration precludes reading the mappings file. Use of the -nomapping_file option will prevent the IMTA_MAPPING_FILE file from being read in when there is no compiled configuration.
-message = <i>message-file</i>	Specifies the text file containing the message that is tested. The <i>message-file</i> must be an RFC 822 message only; it cannot be a queue file.
-mm	Tells imsimta test -exp to load the sieve-specific extensions to the expression interpreter. This includes all the sieve tests and actions such as header , address , envelope , discard , fileinto , and keep . Without -mm you cannot test sieves. The command to test sieves against a message is: imsimta test -expression -mm -message=message
-noreprocess	Turns off the internal reprocessing flag that would otherwise be set. This option is useful for simulating the behavior of other components that operate without the reprocessing flag being set. This can be thought of as controlling whether or not rewrite_test acts as if it were the reprocessing channel. The biggest effect is that it turns off deferred list processing. Normally it should be done so this switch defaults on; use -noreprocessing to disable expansion.
-option_file = <i>filename</i> -nooption_file	Instructs the command to use the specified option file rather than the default option file named by the IMTA_OPTION_FILE option in the MTA tailor file, <i>MessagingServer_home/config/imta_tailor</i> , which is usually the file <i>MessagingServer_home/config/options.dat</i> . This option has no effect unless -noimage_file is specified or no compiled configuration exists; use of any compiled configuration precludes reading any sort of option file. Use of the -nooption_file option prevents the IMTA_OPTION_FILE file from being read in when there is no compiled configuration.
-output = <i>output-file</i>	Directs the output of imsimta test . By default, imsimta test writes output to stdout . This option only works if the mailsrv account has write access to the current working directory.
-read_receipt -noread_receipt	Sets the corresponding receipt request flags. This option can be useful when testing the handling of receipt requests at the time of rewriting forwarded addresses or mailing lists.

Table 8–28 (Cont.) *imsimta test* Command Options

Option	Description
-restricted = <i>setting</i>	Controls the setting of the restricted flag. By default, this flag has value 0 . When set to 1 , -restricted=1 , the restricted flag is set on and addresses are rewritten using the restricted mailbox encoding format recommended by RFC 1137. This flag is used to force rewriting of address mailbox names in accordance with the RFC 1137 specifications.
-sender = <i>from_address</i>	A value used to set the "authenticated sender" (final field) of FROM_ACCESS mapping table probes. That is, one received as a result of SASL authentication. This allows test -rewrite to be used to test these mappings.
-sopin	(New in MS7u2) Show any optins associated with the recipient address.
-source_channel = <i>channel</i>	Controls which source channel is performing the rewriting. Some address rewriting is source channel-specific; imsimta test -rewrite normally assumes that the channel source for which it is rewriting is the local channel <i>l</i> .
-spares	(New in MS7u2) Show any spare attributes associated with the recipient address.
-table = <i>table-name</i>	Specifies the name of the mapping table to test. If this option is not specified, then imsimta test -mapping prompts for the name of the table to use.

Example

This example shows typical output generated by **imsimta test -rewrite**. The most important piece of information generated by **imsimta test -rewrite** is displayed on the last few lines of the output, which shows the channel to which **imsimta test -rewrite** would submit a message with the specified test address and the form in which the test address would be rewritten for that channel. This output is invaluable when debugging configuration problems.

```
imsimta test -rewrite

Address: joe.blue
channel = 1
channel description =
channel description =
channel flags #1 = BIDIRECTIONAL MULTIPLE IMMNONURGENT NOSERVICEALL
channel flags #2 = NOSMTP POSTHEADBODY HEADERINC NOEXPROUTE
channel flags #3 = LOGGING NOGREY NORESTRICTED
channel flags #4 = EIGHTNEGOTIATE NOHEADERTRIM NOHEADERREAD RULES
channel flags #5 =
channel flags #6 = LOCALUSER NOX_ENV_TO RECEIPTHEADER
channel flags #7 = ALLOWSWITCHCHANNEL NOREMOTEHOST DATEFOUR DAYOFWEEK
channel flags #8 = NODEFRAGMENT EXQUOTA REVERSE NOCONVERT_OCTET_STREAM
channel flags #9 = NOTHURMAN INTERPRETENCODING

text/plain charset def = (7) US-ASCII 5 (8) ISO-8859-1 51
channel envelope address type = SOURCEROUTE
channel header address type = SOURCEROUTE
channel official host = mailserver.eng.alpha.com
channel local alias =
channel queue name =
channel after param =
channel daemon name =
```

```

channel user name      =
notices                =
channel group ids     =
header To: address    = joe.blue@mailserver.eng.alpha.com
header From: address  = joe.blue@mailserver.eng.alpha.com
envelope To: address  = joe.blue@mailserver.eng.alpha.com
(route (mailserver.eng.alpha.com,mailserver.eng.alpha.com))
envelope From: address = joe.blue@mailserver.eng.alpha.com
name                   =
mbox                   = joe.blue
Extracted address action list: joe.blue@mailserver.eng.alpha.com
Extracted 733 address action list: joe.blue@mailserver.eng.alpha.com
Expanded address:
  joe.blue@mailserver.eng.alpha.com
Submitted address list:
  ims-ms
joe.blue@ims-ms-daemon (sims-ms-daemon) *NOTIFY FAILURES* *NOTIFY DELAYS*
Submitted notifications list:
Address:
#

```

In the following example, the sample **PAGER** mapping is tested. The **-mapping_file** option is used to select the mapping file **pager_table.sample** instead of the default mapping file.

```

imsimta test -mapping -noimage_file \
  -mapping_file=MessagingServer_home/config/pager_table.sample

```

In the following example, the sample mapping pattern **\${ax1}@.xyz.com** is tested for several sample target strings:

```

imsimta test -match

Pattern: ${ax1}*@*.xyz.com
[ 1S] cglob [1ax]
[ 2] "@"
[ 3S] glob, req 46, reps 2
[ 4] "."
[ 5] "x"
[ 6] "y"
[ 7] "z"
[ 8] "."
[ 9] "c"
[ 10] "o"
[ 11] "m"
Target: xx11aa@sys1.xyz.com
Match.
0 - xx11aa
1 - sys1
Pattern: ${ax1}*@*.xyz.com
Target: 12a@node.xyz.com
No match.
Pattern: ${ax1}*@*.xyz.com
Target: 1xa@node.acme.com
Match.
0 - 1xa
1 - node
Pattern: ^D
%

```

imsimta test -domain

The **imsimta test -domain** utility verifies the validity of domain structures in the LDAP directory, lists all domains in the directory, displays domains' locations (DNs), displays the values of domain attributes, and performs other domain-related tests.

You must be **root** or **mailsrv** to run **imsimta test -domain**.

Syntax

The **imsimta test -domain** utility must be run in interactive mode. To run the utility, type:

```
imsimta test -domain
```

The utility displays the following prompt:

```
DOMAIN_MAP>
```

At the **DOMAIN_MAP>** prompt, you can enter the commands shown in [Table 8–29](#), "**imsimta test -domain Command Options**".

To exit the interactive mode, enter **exit** or **quit**.

imsimta test -domain Commands

The commands for this utility are:

Table 8–29 imsimta test -domain Command Options

Command	Description
enumerate	Displays a list of all available domains in the LDAP directory.
exit	Exits the interactive-mode utility.
help	Displays help (usage) for the commands. Note: The help command displays usage for the canonicalize and user commands, which are used internally by Messaging Server and have no effect on domains in the directory. Do not use these commands.
locate domain <i>domain name</i>	Selects (binds to) the domain entry specified with its domain name (<i>domain name</i>). Once you select a domain entry, you can display information about the domain and its attributes with the show and queryattribute commands.
locate basedn <i>baseDN name</i>	Selects (binds to) the domain entry specified with its base DN (<i>baseDN name</i>). Enclose the specified base DN in single quotes (') or double quotes ("). Once you select a domain entry, you can display information about the domain and its attributes with the show and queryattribute commands.
queryattribute	Displays the name and value of the domain attribute specified with the variable <i>attribute</i> . The specified attribute must be present in the domain currently selected for display.
quit	Exits the interactive-mode utility.
release	Releases the currently selected domain entry. Once you release a domain entry, no information is displayed with the show and query attribute commands.
show	Displays information about the currently selected domain, including the domain name, canonical domain name, base DN, and domain DN.
verify	Verifies domain-level Schema 1 and 2 information in the directory.

Examples

To display a list of all available domains in the directory:

```
imsimta test -domain
DOMAIN_MAP> enumerate
```

```
example.org
example.edu
example.com
```

To verify the domain structure of all available domains:

```
imsimta test -domain
DOMAIN_MAP> verify
```

```
%DMAP-E-CANONICAL, Overlapping domains 'example.com' and
'west.example.com' defined by entries 'o=example.com,o=rootsuffix'
and 'o=west.example.com,o=example.com,o=rootsuffix' have different
canonical domains 'example.com' and 'west.example.com'.
%DMAP-E-NODOMAINNAME, Domain entry with DN 'o=mycompany,o=rootsuffix'
does not have a domain name.
%DMAP-E-DOMAININVALID, Domain name 'domain_tst.com'
defined/referenced by domain entry with
DN 'o=domain_tst.com,o=rootsuffix' is syntactically invalid.
```

To select a domain entry for display by specifying the domain name:

```
imsimta test -domain
DOMAIN_MAP> locate domain example.org
```

```
entry located
```

To select a domain entry for display by specifying the base DN:

```
imsimta test -domain
DOMAIN_MAP> locate basedn "o=example.org,o=rootsuffix"
```

```
entry located
```

To display information about the domain entry:

```
imsimta test -domain
DOMAIN_MAP> show
```

```
Domain name: example.org
Canonical name: example.org
Lower case canonical name: example.org
Base DN: o=example.org,o=rootsuffix
Domain DN: o=example.org,o=rootsuffix
```

To show the value of the **sunPreferredDomain** attribute for the selected domain:

```
imsimta test -domain
DOMAIN_MAP> query sunPreferredDomain
```

```
Attribute value(s):
 [0] "example.org"
```

To verify domain-level Schema 1 and 2 information in the directory:

```
imsimta test -domain
DOMAIN_MAP> verify
```

Various checks are done by this utility, but the most important by far is verification of canonical domain settings for domains with overlapping user entries.

The verification utility can return the following fatal errors:

```
%DMAP-F-CANTGETDN, Cannot obtain DN of domain entry, directory error
%DMAP-F-INTDEFERROR, Internal defined flag error on domain '%.s', aborting
%DMAP-F-INTHASHERROR, Internal hash error, aborting
%DMAP-F-INTTREETSTRUCTERROR, Internal tree structure error, aborting
```

These are all indicative of an internal error in the verification code and should never occur.

The following domain errors can be reported:

```
%DMAP-E-ALIASTOOLONG, Domain alias '%s' in entry with DN '%s' is too long
%DMAP-E-BASEDNTOOLONG, Base DN pointer '%s' in entry for domain '%.s' is
too
                                long
%DMAP-E-CANONICAL, Overlapping domains '%.s' and '%.s' defined by entries
                        '%.s' and '%.s' have different canonical domains '%.s'
                        and '%.s'
%DMAP-E-CANONICALINVALID, Canonical domain '%.s' defined/referenced by
                        domain entry with DN '%.s' is syntactically
                        invalid
%DMAP-E-CANONICALTOOLONG, Canonical name '%s' in entry for domain '%.s'
                        is too long
%DMAP-E-CANTCONVDCDN, Cannot convert DN '%s' in DC tree to domain name
%DMAP-E-CANTEXTALIAS, Empty alias pointer attribute in '%.s' domain alias
                        entry
%DMAP-E-DOMAININVALID, Domain name '%.s' defined/referenced by domain entry
                        with DN '%.s' is syntactically invalid
%DMAP-E-DOMAINMULTDEF, Domain '%s' multiply defined by entries with DNs '%s'
                        and '%s'
%DMAP-E-DOMAINTOOLONG, Domain '%s' in entry with DN '%s' is too long
%DMAP-E-DOMAINUNDEF, Domain name '%.s' referenced by domain entry with DN
                        '%.s' never defined
%DMAP-E-EMPTYCANONICAL, Domain '%.s' has an empty canonical name
%DMAP-E-INVALIDBASEDN, Base DN pointer '%.s' in entry for domain '%.s'
                        is not a valid DN
%DMAP-E-MULTICANONICAL, Multivalued canonical name in entry for domain
                        '%.s', used value '%s' ignored '%s'
%DMAP-E-NOBASEDN, Domain '%.s' has no base DN
%DMAP-E-EMPTYBASEDN, Domain '%.s' has an empty base DN
%DMAP-E-NODOMAINNAME, Domain entry with DN '%s' does not have a domain
                        name
```

The following warnings can be reported:

```
%DMAP-W-DISALLOWEDATTR, Domain '%.s' has a disallowed attribute '%s'
                        with value '%s'
%DMAP-W-DNTOOLONG, Domain entry DN '%s' is too long
%DMAP-W-EMPAPPSTAT, Domain '%.s' has an empty application status
%DMAP-W-EMPDISALLOWED, Domain '%.s' has an empty disallowed attribute
                        '%s'
%DMAP-W-EMPDOMSTAT, Domain '%.s' has an empty domain status
%DMAP-W-EMPUIDSEP, Domain '%.s' has an empty UID separator
%DMAP-W-INVALIDAPPSTAT, Application status '%s' for domain '%.s' is
```

```

invalid
%DMAP-W-INVALIDDOMSTAT, Domain status '%s' for domain '%.*s' is invalid
%DMAP-W-INVALIDUIDSEP, UID separator '%s' for domain '%.*s' is invalid
%DMAP-W-MULTDOMAINNAMES, Domain entry with DN '%s' has multiple domain
names, used value '%s' ignored '%s'
%DMAP-W-MULTIAPPSTAT, Multivalued application status in entry for domain
'%. *s', used value '%s' ignored '%s'
%DMAP-W-MULTIBASEDN, Multivalued base DN pointer in entry for domain
'%. *s', used value '%s' ignored '%s'
%DMAP-W-MULTIDOMSTAT, Multivalued domain status in entry for domain
'%. *s', used value '%s' ignored '%s'
%DMAP-W-MULTIUIDSEP, Multivalued UID separator in entry for domain '%.*s',
used value '%s' ignored '%s'
%DMAP-W-MULTIVALIAS, Multivalued alias pointer in entry for domain alias
'%. *s', used value '%s' ignored '%s'
%DMAP-W-NOBASEDNODE, Base DN pointer '%.*s' in entry for domain '%.*s'
doesn't point at anything
%DMAP-W-NODOMAINNAME, Domain entry with DN '%s' has a blank domain alias
%DMAP-W-NOENTRIES, No domain entries found, aborting

```

imsimta version

The **imsimta version** command prints out the MTA version number, and displays the system's name, operating system release number and version, and hardware type.

Syntax

```
imsimta version
```

Example

To check the version of MTA you are running, execute the following command:

```
% imsimta version
```

imsimta view

The **imsimta view** utility displays log files.

Syntax

```
imsimta view <file-pattern> [-f <offset-from-first>] [-l <offset-from-last>]
```

Options

The options for this command are:

Table 8–30 *imsimta view* Command Options

Option	Description
<i>-f=offset-from-first</i>	Displays the specified version of the log file (starting from 0). For example, to find the earliest (oldest) version of the file, specify <i>-f=0</i> . By default, imsimta view finds the most recent version of the log file.
<i>-l=offset-from-last</i>	Displays the last version of the specified file. For example, to display the most recent (newest) version of the file, specify <i>-l=0</i> . By default, imsimta view finds the most recent version of the file.
<i>file-pattern</i>	Specifies a filename pattern to view.

Configuring Rewrite Rules

This information describes how to configure rewrite rules. Make sure you have read ["About MTA Services"](#) before using this information.

Oracle Communications Messaging Server's address rewriting facility is the primary facility for manipulating and changing the host or domain portion of addresses. Messaging Server provides other facilities for address manipulation, such as aliases, the address reversal database, and specialized mapping tables. For best performance, however, rewrite rules should be used whenever possible to perform address manipulations.

Before You Begin

When you make changes to rewrite rules, you must restart any programs or channels that load the configuration data only once when they start up, for example, the SMTP server, by using the **imsimta restart** command. If you are using a compiled configuration, you must recompile the configuration by running **imsimta cnbuild** and then restart.

For more information about compiling configuration information and starting programs, see ["Recompiling the MTA Configuration."](#)

Editing Rewrite Rules in Unified Configuration

To make rewrite rules changes:

- Use the **msconfig edit rewrites** command to bring up the rewrite rules in the editor specified by the **EDITOR** shell variable of the account that you are logged in as. For example:

```
msconfig edit rewrites
$* $A$E$F$U%$H$V$H@&/IMTA_HOST/
&/IMTA_HOST/ $U%$D@&/IMTA_HOST/
&/IMTA_DEFAULTDOMAIN/ $U%$D@&/IMTA_HOST/
. $U%$H$, $H@tcp_local-daemon
[] $E$R${INTERNAL_IP, $L}$U%[$L]@tcp_intranet-daemon
.&/IMTA_DEFAULTDOMAIN/ $U%$H.&/IMTA_DEFAULTDOMAIN/@tcp_intranet-daemon
* $U%$&0.&/IMTA_DEFAULTDOMAIN/
reprocess $U%reprocess.&/IMTA_HOST/@reprocess-daemon
reprocess.&/IMTA_HOST/ $U%reprocess.&/IMTA_HOST/@reprocess-daemon
process $U%process.&/IMTA_HOST/@process-daemon
process.&/IMTA_HOST/ $U%process.&/IMTA_HOST/@process-daemon
defragment $U%defragment.&/IMTA_HOST/@defragment-daemon
defragment.&/IMTA_HOST/ $U%defragment.&/IMTA_HOST/@defragment-daemon
conversion $U%conversion.&/IMTA_HOST/@conversion-daemon
```

```

conversion.&/IMTA_HOST/ $U%conversion.&/IMTA_HOST/@conversion-daemon
bitbucket $U%bitbucket.&/IMTA_HOST/@bitbucket-daemon
bitbucket.&/IMTA_HOST/ $U%bitbucket.&/IMTA_HOST/@bitbucket-daemon
hold-daemon $U%H@hold-daemon
~
~
~
~
~
~
"/var/tmp/manage.v_a474" 18 lines, 1226 characters

```

You can also add rewrite rules as regular option values by using the **set** command. For example, a separate queue for messages sent to a slowly responding domain could be accomplished with the following commands:

```

msconfig
msconfig> show rewrite.rule
role.rewrite.rule = $* $A$E$F$U%H$V$H@&/IMTA_HOST/
role.rewrite.rule = &/IMTA_HOST/ $U%D@&/IMTA_HOST/
role.rewrite.rule = &/IMTA_DEFAULTDOMAIN/ $U%D@&/IMTA_HOST/
role.rewrite.rule = . $U%H$, $H@tcp_local-daemon
role.rewrite.rule = [] $E$R$ {INTERNAL_IP, $L}$U@[ ]@tcp_intranet-daemon
role.rewrite.rule = .&/IMTA_DEFAULTDOMAIN/ $U%H.&/IMTA_DEFAULTDOMAIN/@tcp_
intranet-daemon
role.rewrite.rule = * $U%$&0.&/IMTA_DEFAULTDOMAIN/
role.rewrite.rule = reprocess $U%reprocess.&/IMTA_HOST/@reprocess-daemon
role.rewrite.rule = reprocess.&/IMTA_HOST/ $U%reprocess.&/IMTA_
HOST/@reprocess-daemon
role.rewrite.rule = process $U%process.&/IMTA_HOST/@process-daemon
role.rewrite.rule = process.&/IMTA_HOST/ $U%process.&/IMTA_
HOST/@process-daemon
role.rewrite.rule = defragment $U%defragment.&/IMTA_HOST/@defragment-daemon
defragment.&/IMTA_HOST/ $U%defragment.&/IMTA_HOST/@defragment-daemon
conversion $U%conversion.&/IMTA_HOST/@conversion-daemon
conversion.&/IMTA_HOST/ $U%conversion.&/IMTA_HOST/@conversion-daemon
bitbucket $U%bitbucket.&/IMTA_HOST/@bitbucket-daemon
bitbucket.&/IMTA_HOST/$U%bitbucket.&/IMTA_HOST/@bitbucket-daemon
hold-daemon $U%H@hold-daemon
msconfig> set rewrite.rule slow.com $U%D@TCP-SLOW
msconfig# set rewrite.rule .slow.com $U%H$D@TCP-SLOW
msconfig# set channel:tcp_slow.official_host_name TCP-SLOW
msconfig# set channel:tcp_slow.smtp
msconfig# set channel:tcp_slow.mx
msconfig# set channel:tcp_slow.pool SMTP_POOL
msconfig# set mapping:ORIG_SEND_ACCESS.rule "tcp_local|*|tcp_slow|*"
"$N$D30|Relaying not allowed without prior authentication"
msconfig# diff
< role.channel:tcp_slow.mx (novalue)
< role.channel:tcp_slow.official_host_name = TCP-SLOW
< role.channel:tcp_slow.pool = SMTP_POOL
< role.channel:tcp_slow.smtp (novalue)
< role.mapping:ORIG_SEND_ACCESS.rule = tcp_local|*|tcp_slow|* $N$D30|Relaying
not allowed without prior authentication
< role.rewrite.rule = .slow.com $U%H$D@TCP-SLOW
< role.rewrite.rule = slow.com $U%D@TCP-SLOW
msconfig# write -remark="Add separate queue for slow domain slow.com"
msconfig>

```

Note: The `set rewrite.rule` command always adds rules at the end of the list, so order-specific rules must be added in the desired order. Also, it is not presently practical to remove rewrite rules with the `unset` command.

Rewrite Rule Structure

The following example configuration shows how rewrite rules are used to route messages to the proper channel. No domain names are used to keep things as simple as possible. In the legacy configuration, the rewrite rules appear in the upper half followed by the channel definitions in the lower half.

```
! An example configuration file.!
! This is only an example of a configuration file. It serves
! no useful purpose and should not be used in a real system.
!
! Part I: Rewrite rules
a $U@a-daemon
b $U@b-daemon
c $U%c@b-daemon
d $U@d@a-daemon
!
! Begin channel definitions
```

Rewrite rules consist of two parts: a pattern, followed by an equivalence string or *template*. The two parts must be separated by spaces, although spaces are not allowed within the parts themselves. When you view mappings by using the `msconfig edit rewrite` command, you see something that resembles the following:

```
<pattern> <template>
<pattern> <template>
<pattern> <template>
```

pattern

Indicates the string to search for in the domain name. In [Table 9–3, "Extracted Addresses and Host Names"](#), the patterns are **a.com**, **b.org**, **c.edu**, and **d.com**.

If the pattern matches the domain part of the address, the rewrite rule is applied to the address. A blank space must separate the pattern from the template. For more information about pattern syntax, see "[Rewrite Rule Patterns and Tags](#)."

template

is one of the following:

```
<UserTemplate>%<DomainTemplate>@<ChannelTag> [<controls>]
<UserTemplate>@<ChannelTag> [<controls>]
<UserTemplate>%<DomainTemplate> [<controls>]
<UserTemplate>@<DomainTemplate>@<ChannelTag> [<controls>]
<UserTemplate>@<DomainTemplate>@<SourceRoute>@<ChannelTag> [<controls>]
```

where:

UserTemplate specifies how the user part of the address is rewritten. Substitution sequences can be used to represent parts of the original address or the results of a database lookup. The substitution sequences are replaced with what they represent to construct the rewritten address. In [Table 9–4, "Summary of Rewrite Rule Template Substitutions and Control Sequences"](#), the `$U` substitution sequence is used. For more

information, see ["Template Substitutions and Rewrite Rule Control Sequences."](#)

DomainTemplate specifies how the domain part of the address is rewritten. Like the *UserTemplate*, the *DomainTemplate* can contain substitution sequences.

ChannelTag indicates the channel to which this message is sent. (All channel definitions must include a channel tag as well as a channel name. The channel tag typically appears in rewrite rules, as well as in its channel definition.)

controls limits the applicability of a rule. Some control sequences must appear at the beginning of the rule. Other controls must appear at the end of the rule. For more information about controls, see ["Template Substitutions and Rewrite Rule Control Sequences."](#)

For more information about template syntax, see ["Rewrite Rule Templates."](#)

Rewrite Rule Patterns and Tags

This section consists of the following subsections:

- [A Rule to Match Percent Hacks](#)
- [A Rule to Match Bang-Style \(UUCP\) Addresses](#)
- [A Rule to Match Any Address](#)
- [Tagged Rewrite Rule Sets](#)

Most rewrite rule patterns consist either of a specific host name that will match only that host or of a subdomain pattern that will match any host or domain in the entire subdomain.

For example, the following rewrite rule pattern contains a specific host name that will match the specified host only:

```
host.example.org
```

The next rewrite rule pattern contains a subdomain pattern that will match any host or domain in the entire subdomain:

```
.example.org
```

This pattern will not, however, match the exact host name **example.org**. To match the exact host name **example.org**, a separate **example.org** pattern would be needed.

The MTA attempts to rewrite host and domain names starting from the specific host name and then incrementally generalizing the name to make it less specific. This means that a more specific rewrite rule pattern will be preferentially used over more general rewrite rule patterns. For example, assume the following rewrite rule patterns are present in the configuration:

```
hosta.subnet.example.org  
.subnet.example.org  
.example.org
```

Based on the rewrite rule patterns, an address of **jd@hosta.subnet.example.org** matches the **hosta.subnet.example.org** rewrite rule pattern. An address of **jd@hostb.subnet.example.org** matches the **.subnet.example.org** rewrite rule pattern. And an address of **jd@hostc.example.org** matches the **.example.org** rewrite rule pattern.

In particular, the use of rewrite rules incorporating subdomain rewrite rule patterns is common for sites on the Internet. Such a site will typically have a number of rewrite

rules for its own internal hosts and subnets, and then include rewrite rules for the top-level Internet domains into its configuration.

IP domain literals follow a similar hierarchical matching pattern, though with right-to-left (rather than left-to-right) matching. For example, the following pattern matches only and exactly the IP literal `[1.2.3.4]`:

```
[1.2.3.4]
```

The next pattern matches anything in the `[1.2.3.0]` subnet:

```
[1.2.3.]
```

In addition to the more common sorts of host or subdomain rewrite rule patterns already discussed, rewrite rules may also make use of several special patterns, summarized in [Table 9–1](#), and discussed in the following subsections.

Table 9–1 Summary of Special Patterns for Rewrite Rules

Pattern	Description/Usage
<code>\$*</code>	Matches any address. This rule, if specified, is tried first regardless of its position in the file.
<code>\$%</code>	Percent Hack Rule. Matches any host/domain specification of the form <code>A%B</code>
<code>\$!</code>	Bang-style Rule. Matches any host/domain specification of the form <code>B!A</code>
<code>[]</code>	IP literal match-all rule. Matches any IP domain literal.
<code>.</code>	Matches any host/domain specification. For example, <code>joe@[129.165.12.11]</code>

In addition to these special patterns, Messaging Server also has the concept of *tags*, which may appear in rewrite rule patterns. These tags are used in situations where an address may be rewritten several times and, based upon previous rewrites, distinctions must be made in subsequent rewrites by controlling which rewrite rules match the address. For more information, see [Tagged Rewrite Rule Sets](#).

The contents of the `internet.rules` file have been absorbed into the main configuration and the individual entries for top-level domains have been removed and replaced by the following single rewrite rule:

```
$U%$H$, $H@TCP-DAEMON
```

This rewrite rule replaces the reference to the `internet.rules` file that appears in legacy configurations as the following:

```
<IMTA_TABLE:internet.rules
```

To ensure that messages to Internet destinations (other than to the internal host destinations handled by using more specific rewrite rules) are properly rewritten and routed to an outgoing TCP/IP channel, the configuration must contain the following information:

- Rewrite rules with patterns that match the top-level Internet domains
- Templates that rewrite addresses matching such patterns to an outgoing TCP/IP channel

When additional Top Level Domains (TLDs) are defined by the IANA, perform the following based on whether or not you are using a Unified Configuration or a legacy configuration:

To Update TLDs

1. Obtain a new copy of the **tlds.txt** file from the Internet Assigned Numbers Authority (IANA).
2. Put the **tlds.txt** file in the *MessagingServer_homes***config** directory.
3. Run the following command:

```
imsimta chbuild
```

A Rule to Match Percent Hacks

If the MTA tries to rewrite an address of the form **A%B** and fails, it tries one extra rule before falling through and treating this address form as **A%B@localhost**. (For more information about these address forms, see "[Rewrite Rule Templates](#).") This rule is only activated when a local part containing a percent sign has failed to rewrite any other way (including the match all rule described later in this information).

The percent hack rule is useful for assigning some special, internal meaning to percent hack addresses.

A Rule to Match Bang-Style (UUCP) Addresses

If the MTA tries to rewrite an address of the form **B!A** and fails, it tries one extra rule before falling through and treating this address form as **B!A@localhost**. This extra rule is the bang-style rule. The pattern is **!**. The pattern never changes. This rule is only activated when a local part containing an exclamation point has failed to rewrite any other way (including the default rule described later).

The bang-style rule can be used to force UUCP style addresses to be routed to a system with comprehensive knowledge of UUCP systems and routing.

A Rule to Match Any Address

The special pattern **."** (a single period) matches any host/domain specification if no other rule matches and the host/domain specification cannot be found anywhere in the channel table. In other words, the **."** rule is used as a last resort when address rewriting would fail otherwise.

Note: Regarding substitution sequences, when the match-all rule matches and its template is expanded, **\$H** expands to the full host name and **\$D** expands to a single dot ("**."**"). Thus, **\$D** in a match-all rule template! is of limited use.

Tagged Rewrite Rule Sets

As the rewrite process proceeds it might be appropriate to bring different sets of rules into play. This is accomplished by the use of the rewrite rule tag. The current tag is prepended to each pattern before looking it up in the configuration file or domain database. The tag can be changed by any rewrite rule that matches by using the **\$T** substitution string in the rewrite rule template (described later).

Tags are somewhat sticky. Once set, they continue to apply to all hosts that are extracted from a single address. This means that care must be taken to provide alternate rules that begin with the proper tag values once any tags are used. In practice this is rarely a problem since tags are usually used in only very specialized applications. Once the rewriting of the address is finished the tag is reset to the default tag, which is an empty string.

By convention all tag values end in a vertical bar |. This character is not used in normal addresses and thus is free to delineate tags from the rest of the pattern.

Rewrite Rule Templates

The following sections describe in more detail template formats for rewrite rules. [Table 9–2](#) summarizes the template formats.

Table 9–2 Summary of Template Formats for Rewrite Rules

Template	Usage
A%B	A becomes the new user/mailbox name, B becomes the new host/domain specification, rewrite again. See "Repeated Rewrites Template, A%B."
A@B	Treated as A%B@B. See "Ordinary Rewriting Templates, A%B@C or A@B."
A%B@C	A becomes the new user/mailbox name, B becomes the new host/domain specification, route to the channel associated with the host C. See "Ordinary Rewriting Templates, A%B@C or A@B."
A (B)C	Treated as A (B)C@C. See "Specified Route Rewriting Templates, A (B)C@D or A (B)C."
A (B)C@D	A becomes the new user/mailbox name, B becomes the new host/domain specification, insert C as a source route, route to the channel associated with the host D. See "Specified Route Rewriting Templates, A (B)C@D or A (B)C."

Ordinary Rewriting Templates, A%B@C or A@B

The following template is the most common form of template. The rule is applied to the user part of the address and to the domain part of the address. The new address is then used to route the message to a specific channel (indicated by *ChannelTag*).

```
<UserTemplate>%<DomainTemplate>@<ChannelTag> [<controls>]
```

The next form of template is identical in application to the most common form of template. However, this form of template is possible only if *DomainTemplate* and *ChannelTag* are identical.

```
<UserTemplate>@<ChannelTag> [<controls>]
```

Repeated Rewrites Template, A%B

The following template format is used for meta-rules that require additional rewriting after application of the rule. After the rule is applied, the entire rewriting process is repeated on the resulting new address. (All other rewrite rule formats cause the rewriting process to terminate after the rule has been applied).

```
<UserTemplate>%<DomainTemplate> [<controls>]
```

While the special **A%B** form does cause rewriting of the current domain to restart, it is actually just a continuation of the current rewriting process. It does not rewrite the entire process from the beginning. It does not perform the **\$*** pattern when it goes through the second time.

For example, the following rule has the effect of removing all occurrences of the **.removable** domain from the ends of addresses:

```
.removable $U%$H
```

Extreme care must be taken when using these repeating rules as careless use can create a "rules loop". For this reason meta-rules should only be used when absolutely necessary. Be sure to test meta-rules with the **imsimta test -rewrite** command. For more information on the **test -rewrite** command, see the *Messaging Server Reference*.

Specified Route Rewriting Templates, A (B)C@D or A (B)C

The following template format works in the same way as the more common template **UserTemplate%DomainTemplate@ChannelTag** (note the difference in the first separator character), except that *ChannelTag* is inserted into the address as a source route. The message is then routed to *ChannelTag*:

```
<UserTemplate>@<DomainTemplate>@<Source-Route>@<ChannelTag> [<controls>]
```

The rewritten address becomes **@route:user@domain**. The following template is also valid:

```
<UserTemplate>@<DomainTemplate>@<ChannelTag> [<controls>]
```

For example, the following rule rewrites the address **jd@com1** into the source-routed address **@example.org:jd@com1**. The channel tag becomes **example.org**:

```
com1 $U@com1@example.org
```

Case Sensitivity in Rewrite Rule Templates

Unlike the patterns in rewrite rules, character case in templates is preserved. This is necessary when using rewrite rules to provide an interface to a mail system that is sensitive to character case. Substitution sequences like **\$U** and **\$D** that substitute material extracted from addresses also preserve the original case of characters.

When it is desirable to force substituted material to use a particular case, for example, to force mailboxes to lowercase on UNIX systems, special substitution sequences can be used in templates to force substituted material to the desired case. Specifically, **\$** forces subsequent substituted material into lower case, **\$^** forces subsequent substituted material into upper case, and **\$_** says to use the original case.

For example, you can use the following rule to force mailboxes to lowercase for **unix.example.org** addresses:

```
unix.example.org $\$U$_%unix.example.org
```

How the MTA Applies Rewrite Rules to an Address

The following steps describe how the MTA applies rewrite rules to a given address:

1. The MTA extracts the first host or domain specification from an address. An address can specify more than one host or domain name as in the case: `jdoe%hostname@example.org`.
2. After identifying the first host or domain name, the MTA conducts a search that scans for a rewrite rule whose pattern matches the host or domain name.
3. When the matching rewrite rule is found, the MTA rewrites the address according to the template portion of that rule.
4. Finally, the MTA compares the channel tag with the host names that are associated with each channel. If a match is found, the MTA enqueues the message to the associated channel; otherwise, the rewrite process fails. If the matching channel is the local channel, some additional rewriting of the address may take place by looking up the alias database and alias file.

These steps are described in more detail in the subsections that follow.

Note: Using a channel tag that does not belong to any existing channel will cause messages whose addresses match this rule to be bounced. That is, it makes the matching messages nonroutable.

This section consists of the following subsections:

- [Step 1. Extract the First Host or Domain Specification](#)
- [Step 2. Scan the Rewrite Rules](#)
- [Step 3. Rewrite Address According to Template](#)
- [Step 4. Finish the Rewrite Process](#)
- [Rewrite Rule Failure](#)
- [Syntax Checks After Rewrite](#)
- [Handling Domain Literals](#)

Step 1. Extract the First Host or Domain Specification

The process of rewriting an address starts by extracting the first host or domain specification from the address. (Readers not familiar with RFC 822 address conventions are advised to read that standard to understand the following discussion.) The order in which host/domain specifications in the address are scanned is as follows:

1. Hosts in source routes (read from left to right)
2. Hosts appearing to the right of the "at" sign (@)
3. Hosts appearing to the right of the last single percent sign (%)
4. Hosts appearing to the left of the first exclamation point (!) The order of the last two items is switched if the **bangoverpercent** option is in effect on the channel that is doing the address rewriting. That is, if the channel attempting to enqueue the message is, itself, marked with the **bangoverpercent** option.

Some examples of addresses and the host names that could be extracted first are shown in [Table 9-3](#).

Table 9–3 *Extracted Addresses and Host Names*

Address	First Host Domain Specification	Comments
user@a	a	A "short-form" domain name.
user@a.b.c	a.b.c	A "fully qualified" domain name (FQDN).
user@[0.1.2.3]	[0.1.2.3]	A "domain literal".
@a:user@b.c.d	a	Source-routed address with a short-form domain name, the "route".
@a.b.c:user@d.e.f	a.b.c	Source-routed address; route part is fully qualified.
@[0.1.2.3]:user@d.e.f	[0.1.2.3]	Source-routed address; route part is a domain literal.
(a,)b,@c:user@d.e.f	a	Source-routed address with an a to b to c routing.
(a,)[0.1.2.3]:user@b	a	Source-routed address with a domain literal in the route part.
user%A@B	B	This nonstandard form of routing is called a "percent hack".
user%A	A	No comment.
user%A%B	B	No comment.
user%%A%B	B	No comment.
A!user	A	"Bang-style" addressing; commonly used for UUCP.
A!user@B	B	No comment.
A!user%B@C	C	No comment.
A!user%B	B	nobangoverpercent option active; the default.
A!user%B	A	bangoverpercent option active.

RFC 822 does not address the interpretation of exclamation points (!) and percent signs (%) in addresses. Percent signs are customarily interpreted in the same manner as "at" signs (@) if no at sign is present, so this convention is adopted by the Messaging Server MTA.

The special interpretation of repeated percent signs is used to allow percent signs as part of local user names; this might be useful in handling some foreign mail system addresses. The interpretation of exclamation points conforms to RFC 976's "bang-style" address conventions and makes it possible to use UUCP addresses with the Messaging Server MTA.

The order of these interpretations is not specified by either RFC 822 or RFC 976, so the **bangoverpercent** and **nobangoverpercent** options can be used to control the order in which they are applied by the channel doing the rewriting. The default is more "standard," although the alternate setting may be useful under some circumstances.

Note: The use of exclamation points (!) or percent signs (%) in addresses is not recommended.

Step 2. Scan the Rewrite Rules

Once the first host or domain specification has been extracted from the address, the MTA consults the rewrite rules to find out what to do with it. The host/domain specification is compared with the pattern part of each rule (that is, the left side of each

rule). The comparison is case insensitive. Case insensitivity is mandated by RFC 822. The MTA is insensitive to case but preserves it whenever possible.

If the host or domain specification does not match any pattern, in which case it is said to "not match any rule", the first part of the host or domain specification - the part before the first period, usually the host name - is removed and replaced with an asterisk (*) and another attempt is made to locate the resulting host or domain specification, but only in the configuration file rewrite rules (the domain database is not consulted).

If this fails, the first part is removed and the process is repeated. If this also fails the next part is removed (usually a subdomain) and the rewriter tries again, first with asterisks and then without. All probes that contain asterisks are done only in the configuration file rewrite rules table; the domain database is not checked. This process proceeds until either a match is found or the entire host or domain specification is exhausted. The effect of this procedure is to try to match the most specific domain first, working outward to less specific and more general domains.

A more algorithmic view of this matching procedure is:

- The host/domain specification is used as the initial value for the comparison strings **spec_1** and **spec_2**. (For example, **spec_1 = spec_2 = a.b.c**).
- The comparison string **spec_1** is compared with the pattern part of each rewrite rule in the configuration and then the domain database until a match is found. The matching procedure is exited if a match is found.
- If no match is found, then the left-most, non-asterisk part of **spec_2** is converted to an asterisk. For example, if **spec_2** is **a.b.c** then it is changed to **.b.c**; if **spec_2** is **b.c**, then it is changed to **..c**. The matching procedure is exited if a match is found.
- If no match is found then the first part, including any leading period, of the comparison string **spec_1** is removed. Where **spec_1** has only one part (for example, **.c** or **c**), the string is replaced with a single period, **.**. If the resulting string **spec_1** is of nonzero length, then you return to step 1. If the resulting string has zero length (for example, was previously **.**), then the lookup process has failed and you exit the matching procedure. For example, suppose the address **dan@sc.cs.example.edu** is to be rewritten. This causes the MTA to look for the following patterns in the given order:

```
sc.cs.example.edu
*.cs.example.edu
.cs.example.edu
*.*.example.edu
.example.edu
*.*.*.edu
.edu
*.*.*.*
.
```

Step 3. Rewrite Address According to Template

Once the host/domain specification matches a rewrite rule, it is rewritten using the template part of the rule. The template specifies three things:

1. A new user name for the address.
2. A new host/domain specification for the address.
3. A channel tag that identifies an existing MTA channel to which messages to this address should be sent.

Step 4. Finish the Rewrite Process

One of two things can happen once the host/domain specification is rewritten.

- If the channel tag is associated neither with the local channel nor a channel marked with the **routelocal** option, or there are no additional host/domain specifications in the address, the rewritten specification is substituted into the address replacing the original specification that was extracted for rewriting, and the rewriting process terminates.
- If the channel tag matches the local channel or a channel marked **routelocal** and there are additional host/domain specifications that appear in the address, the rewritten address is discarded, the original (initial) host/domain specification is removed from the address, a new host/domain specification is extracted from the address, and the entire process is repeated. Rewriting will continue until either all the host/domain specifications are gone or a route through a non-local, non-routelocal channel is found. This iterative mechanism is how the MTA provides support for source routing. In effect, superfluous routes through the local system and routelocal systems are removed from addresses by this process.

Rewrite Rule Failure

If a host/domain specification fails to match any rewrite rule and no default rule is present, the MTA uses the specification "as-is"; for example, the original specification becomes both the new specification and the routing system. If the address has a nonsensical host/domain specification it will be detected when the routing system does not match any system name associated with any channel and the message will be bounced.

Syntax Checks After Rewrite

No additional syntax checking is done after the rewrite rules have been applied to an address. This is deliberate - it makes it possible for rewrite rules to be used to convert addresses into formats that do not conform to RFC 822. However, this also means that mistakes in the configuration may result in messages leaving the MTA with incorrect or illegal addresses.

Handling Domain Literals

Domain literals are handled specially during the rewriting process. If a domain literal appearing in the domain portion of an address does not match a rewrite rule pattern as is, the literal is interpreted as a group of strings separated by periods and surrounded by square brackets. The right-most string is removed and the search is repeated. If this does not work, the next string is removed, and so on until only empty brackets are left. If the search for empty brackets fails, the entire domain literal is removed and rewriting proceeds with the next section of the domain address, if there is one. No asterisks are used in the internal processing of domain literals; when an entire domain literal is replaced by an asterisk, the number of asterisks corresponds to the number of elements in the domain literal.

Like normal domain or host specifications, domain literals are also tried in most specific to least specific order. The first rule whose pattern matches will be the one used to rewrite the host or domain specification. If there are two identical patterns in the rules list, the one which appears first will be used.

As an example, suppose the address **dan@[128.6.3.40]** is to be rewritten. The rewriter looks for **[128.6.3.40]**, then **[128.6.3.]**, then **[128.6.]**, then **[128.]**, then **[]**, then **[...]**, and finally the match-all rule **"."**.

Template Substitutions and Rewrite Rule Control Sequences

Substitutions are used to rewrite user names or addresses by inserting a character string into the rewritten address, the value of which is determined by the particular substitution sequence used.

This section consists of the following subsections:

- Username and Subaddress Substitution, \$U, \$OU, \$1U
- Host/Domain and IP Literal Substitutions, \$D, \$H, \$nD, \$nH, \$L
- Literal Character Substitutions, \$\$, \$%, @\$
- LDAP Query URL Substitutions, \$]...[
- General Database Substitutions, \$(...)
- Apply Specified Mapping, \${...}
- Customer-supplied Routine Substitutions, \$[...]
- Single Field Substitutions, \$&, \$!, \$*, \$#
- Unique String Substitutions
- Source-Channel-Specific Rewrite Rules (\$M, \$N)
- Destination-Channel-Specific Rewrite Rules (\$C, \$Q)
- Direction-and-Location-Specific Rewrite Rules (\$B, \$E, \$F, \$R)
- Host-Location-Specific Rewrites (\$A, \$P, \$S, \$X)
- Changing the Current Tag Value, \$T
- Controlling Error Messages Associated with Rewriting (\$?)

For example, in the following template, the **\$U** is a substitution sequence. It causes the *username* portion of the address being rewritten to be substituted into the output of the template. Thus, if **jd@mailhost.example.org** was being rewritten by this template, the resulting output would be **jd@example.org**, the **\$U** substituting in the *username* portion, **jd**, of the original address:

```
$U@example.org
```

Control sequences impose additional conditions to the applicability of a given rewrite rule. Not only must the pattern portion of the rewrite rule match the host or domain specification being examined, but other aspects of the address being rewritten must meet conditions set by the control sequence or sequences. For example, the **\$E** control sequence requires that the address being rewritten be an envelope address, while the **\$F** control sequence requires that it be a forward pointing address. The following rewrite rule only applies to (rewrite) envelope To: addresses of the form **user@example.org**:

```
example.org $U@mail.example.org$E$F
```

If a domain or host specification matches the pattern portion of a rewrite rule but doesn't meet all of the criteria imposed by a control sequences in the rule's template, then the rewrite rule fails and the rewriter continues to look for other applicable rules.

Table 9–4 summarizes the template substitutions and control sequences.

Table 9–4 Summary of Rewrite Rule Template Substitutions and Control Sequences

Substitution Sequence	Substitutes
\$D	Portion of domain specification that matched.
\$H	Unmatched portion of host/domain specification; left of dot in pattern.
\$L	Unmatched portion of domain literal; right of dot in pattern literal.
\$U	User name from original address.
\$nA	Inserts the <i>n</i> th left character of the current address starting from position 0. The entire address is inserted if <i>n</i> is omitted.
\$nX	Inserts the <i>n</i> th left component of the mailhost starting from 0. The entire mailhost is inserted if <i>n</i> is omitted.
\$0U	Local part (username) from original address, minus any subaddress.
\$1U	Subaddress, if any, from local part (username) of original address.
\$\$	Inserts a literal dollar sign (\$).
\$%	Inserts a literal percent sign (%).
\$@	Inserts a literal at sign (@).
\$	Forces material to lowercase.
\$\$	Forces material to uppercase.
\$_	Uses original case.
\$=	Forces subsequent substituted characters to undergo quoting appropriate for insertion into LDAP search filters.
\$W	Substitutes in a random, unique string.
\$]...[LDAP search URL lookup.
\$.	Establish a string which will be processed as the mapping entry result in the event of a temporary LDAP lookup failure.
\$(text)	General database substitution; rule fails if lookup fails.
\$(...)	Applies specified mapping to supplied string.
\$[...]	Invoke customer supplied routine; substitute in result.
\$&n	The <i>n</i> th part of unmatched (or wildcarded) host, counting from left to right, starting from 0.
\$\$n	The <i>n</i> th part of unmatched (or wildcarded) host, as counted from right to left, starting from 0.
\$*n	The <i>n</i> th part of matching pattern, counting from left to right, starting from 0.
\$\$n	The <i>n</i> th part of matching pattern, counted from right to left, starting from 0.
\$\$D	Portion of domain specification that matched, preserving from the <i>n</i> th leftmost part starting from 0
\$\$H	Portion of host/domain specification that didn't match, preserving from the <i>n</i> th leftmost part starting from 0
Control Sequence	Effect on Rewrite Rule
\$1M	Apply only if the channel is an internal reprocessing channel.

Table 9–4 (Cont.) Summary of Rewrite Rule Template Substitutions and Control

Substitution Sequence	Substitutes
\$IN	Apply only if the channel is not an internal reprocessing channel.
\$1~	Perform any pending channel match checks. If the checks fail, successfully terminate processing of the current rewrite rule template.
\$A	Apply if host is to the right of the at sign
\$B	Apply only to header/body addresses.
\$Cchannel	Fail if sending to <i>channel</i> .
\$E	Apply only to envelope addresses.
\$F	Apply only to forward-directed (e.g., To\:) addresses.
\$Mchannel	Apply only if <i>channel</i> is rewriting the address.
\$Nchannel	Fail if <i>channel</i> is rewriting the address.
\$P	Apply if host is to the right of a percent sign.
\$Qchannel	Apply if sending to <i>channel</i> .
\$R	Apply only to backwards-directed (e.g., From\:) addresses.
\$S	Apply if host is from a source route
\$Tnewtag	Set the rewrite rule tag to newtag
\$Vhost	Fail if the host name is not defined in the LDAP directory (either in the DC tree or as a virtual domain). If the LDAP search times out, the remainder of the rewrite pattern from directly after the character following the host name is replaced with the MTA option string DOMAIN_FAILURE .
\$X	Apply if host is to the left of an exclamation point
\$Zhost	Fail if the host name is defined in the LDAP directory (either in the DC tree or as a virtual domain). If the LDAP search times out, the remainder of the rewrite pattern from directly after the character following the host name is replaced with the MTA option string DOMAIN_FAILURE .
\$nT	Overrides the default ALIAS_MAGIC setting, where <i>n</i> is an appropriate value for the ALIAS_MAGIC MTA option. Overrides the setting for the domain when the rule matches during alias expansion.
?\$errmsg	If rewriting fails, return <i>errmsg</i> instead of the default error message. The error message must be in US ASCII.
\$number?errmsg	If rewriting fails, return <i>errmsg</i> instead of the default error message, and set the SMTP extended error code to <i>a.b.c</i> : <ul style="list-style-type: none"> ▪ <i>a</i> is <i>number</i> / 1000000 (the first digit) ▪ <i>b</i> is (<i>number</i> / 1000) remainder 1000 (the value of the digits 2 through 4) ▪ <i>c</i> is <i>number</i> remainder 1000 (the value of the last three digits). The following example sets the error code to 3.45.89:\$3045089?the snark is a boojum

Username and Subaddress Substitution, \$U, \$OU, \$IU

Any occurrences of **\$U** in the template are replaced with the username (RFC 822 "local-part") from the original address. Note that user names of the form a."b" will be

replaced by "a.b" as RFC2822 deprecates the former syntax from RFC 822 and it is expected that the latter usage will become mandatory in the future.

Any occurrences of **\$0U** in the template are replaced with the username from the original address, minus any subaddress and subaddress indication character. Any occurrences of **\$1U** in the template are replaced with the subaddress and subaddress indication character, if any, from the original address. So note that **\$0U** and **\$1U** are complementary pieces of the username, with **\$0U\$1U** being equivalent to a simple **\$U**.

Host/Domain and IP Literal Substitutions, **\$D**, **\$H**, **\$nD**, **\$nH**, **\$L**

Any occurrences of **\$H** are replaced with the portion of the host/domain specification that was not matched by the rule. Any occurrences of **\$D** are replaced by the portion of the host/domain specification that was matched by the rewrite rule. The **\$nH** and **\$nD** characters are variants that preserve the normal **\$H** or **\$D** portion from the *n*th leftmost part starting counting from 0. That is, **\$nH** and **\$nD** omit the leftmost *n* parts (starting counting from 1) of what would normally be a **\$H** or **\$D**, substitution, respectively. In particular, **\$0H** is equivalent to **\$H** and **\$0D** is equivalent to **\$D**.

For example, assume the address **jdoe@host.example.org** matches the following rewrite rule:

```
host.example.org $U$1D@tcp_local-daemon
```

The resulting address is **jdoe@example.org** with **tcp_local-daemon** used as the outgoing channel. Here where **\$D** would have substituted in the entire domain that matched, **host.example.org**, the **\$1D** instead substitutes in the portions of the match starting from part 1 (part 1 being **example**), so substitutes in **example.org**.

\$L substitutes the portion of a domain literal that was not matched by the rewrite rule.

Literal Character Substitutions, **\$\$**, **\$\$%**, **\$\$@**

The **\$**, **%**, and **@** characters are normally metacharacters in rewrite rule templates. To perform a literal insertion of such a character, quote it with a dollar character, **\$**. That is, **\$\$** expands to a single dollar sign, **\$**; **\$\$%** expands to a single percent, **%** (the percent is not interpreted as a template field separator in this case); and **\$\$@** expands to a single at sign, **@** (also not interpreted as a field separator).

LDAP Query URL Substitutions, **\$[...]**

A substitution of the form **\$[ldap-url]** is interpreted as an LDAP query URL and the result of the LDAP query is substituted. Standard LDAP URLs are used with the host and port omitted. The host and port are instead specified with the **ugldaphost** and **ugldapport** options.

That is, the LDAP URL should be specified as follows where the square bracket characters, **[]**, indicate optional portions of the URL:

```
ldap:///dn[?attributes[?scope?filter]]
```

The **dn** is required and is a distinguished name specifying the search base. The optional attributes, scope, and filter portions of the URL further refine what information to return. For a rewrite rule, the desired attributes to specify returning might be a **mailRoutingSystem** attribute (or some similar attribute). The scope may be any of base (the default), one, or sub. And the desired filter might be to request the return of the object whose **mailDomain** value matches the domain being rewritten.

If the LDAP directory schema includes attributes **mailRoutingSystem** and **mailDomain**, then a possible rewrite rule to determine to which system to route a given sort of address might appear as the following where here the LDAP URL substitution sequence **\$D** is used to substitute in the current domain name into the LDAP query constructed:

```
.example.org \  
$U%$H$D$]ldap:///o=example.org?mailRoutingSystem?sub? \  
(mailDomain=$D)
```

For ease in reading, the backslash character is used to continue the single logical rewrite rule line onto a second physical line. [Table 9-5](#) lists the LDAP URL Substitution Sequences.

Table 9-5 LDAP URL Substitution Sequences

Substitution	Description
\$\$	Literal \$ character
\$.	Establishes a string which will be processed as the mapping entry result in the event of a temporary LDAP lookup failure. By default a temporary failure string remains set only for the duration of the current rule. "\$.." can be used to return to the default state where no temporary failure string is set and temporary LDAP failures cause mapping entry or rewrite rule failure. Note that all errors other than failure to match an entry in the directory are considered to be temporary errors; in general it isn't possible to distinguish between errors caused by incorrect LDAP URLs and errors caused by directory server configuration problems.
\$~account	Home directory of user account
\$A	Address
\$D	Domain name
\$H	Host name (first portion of fully qualified domain name)
\$L	Username minus any special leading characters such as ~ or _
\$S	Subaddress
\$U	Username

The MTA caches URL results from lookups done in rewrite rules and mappings. This URL result cache is controlled by the **URL_RESULT_CACHE_SIZE** (default 10000 entries) and **URL_RESULT_CACHE_TIMEOUT** (default 600 seconds) options.

General Database Substitutions, \$(...)

A substitution of the form **\$(text)** is handled specially. The text part is used as a key to access the special general text database. This database consists of the file *MessagingServer_home/db/generaldb.db*.

If "text-string" is found in the database, the corresponding template from the database is substituted. If "text-string" does not match an entry in the database, the rewrite process fails; it is as if the rewrite rule never matched in the first place. If the substitution is successful, the template extracted from the database is re-scanned for additional substitutions. However, additional **\$(text)** substitutions from the extracted template are prohibited in order to prevent endless recursive references.

As an example, suppose that the address `jd@eng.example.org` matches the following rewrite rule:

```
.EXAMPLENET $(H)
```

Then, the text string `example` will be looked up in the general database and the result of the look up, if any, is used for the rewrite rule's template. Suppose that the result of looking up `example` is `$u%eng.example.org`. Then the output of the template will be `jd@eng.example.org` (that is, username = `jd`, host/domain specification = `eng.example.org`), and the routing system will be `example.net`.

If a general text database exists it should be world readable to insure that it operates properly. See the *Messaging Server Reference* for more information.

Note: If you are using a crdb, on-disk database (the `use_text_databases` option has not been set with the bit saying to use a "text" database for the general database), then `generaldb.db` is the specific file containing the general database, in the hard-coded location `MessagingServer_home/data/db`.

If instead you are using a general "database" that is constructed by the MTA as an in-memory database built from an underlying text file, then the underlying text file is `IMTA_TABLE:general.txt` and the relevant bit (bit 0/value 1) of the `use_text_databases` option is set. **In this case, no database file resides in the `MessagingServer_home/data/db` directory.**

All the MTAs' regular crdb databases (alias database, reverse database, domain database, forward database, and general database) are expected to reside in the `MessagingServer_home/data/db` directory using hard-coded names. However, the general database, reverse database, and forward database can optionally be constructed in-memory by the MTA, as controlled by the `use_text_databases` option. In this case, instead of having a true crdb file located in the `MessagingServer_home/data/db/` directory, there would be a text file in `IMTA_TABLE:`.

Apply Specified Mapping, `$(...)`

A substitution of the form `.EXAMPLENET $(H) $(mapping,argument)` is used to find and apply a mapping from the MTA configuration. The `mapping` field specifies the name of the mapping table to use while `argument` specifies the string to pass to the mapping. The mapping must exist and must set the `$Y` flag in its output if it is successful; if it doesn't exist or doesn't set `$Y` the rewrite will fail. If successful the result of the mapping is merged into the template at the current location and re-expanded.

This mechanism allows the MTA rewriting process to be extended in various complex ways. For example, the username part of an address can be selectively analyzed and modified, which normally isn't a feature the MTA rewriting process is capable of.

Customer-supplied Routine Substitutions, `$(...)`

A substitution of the form `$(image,routine,argument)` is used to find and call a customer-supplied routine. At run-time on UNIX, the MTA uses `dlopen` and `dlsym` to dynamically load and call the specified routine from the shared library `image`. The routine is then called as a function with the following argument list:

```
status := <routine> (<argument>, <arglength>, <result>, <reslength>)
```

argument and *result* are 252 byte long character string buffers. On UNIX, *argument* and *result* are passed as a pointer to a character string, (for example, in C, as `char*`.) *arglength* and *reslength* are signed, long integers passed by reference. On input, *argument* contains the argument string from the rewrite rule template, and *arglength* the length of that string. On return, the resultant string should be placed in *result* and its length in *reslength*. This resultant string will then replace the "\$[*image,routine,argument*]" in the rewrite rule template. The routine should return 0 if the rewrite rule should fail and -1 if the rewrite rule should succeed.

This mechanism allows the rewriting process to be extended in all sorts of complex ways. For example, a call to some type of name service could be performed and the result used to alter the address in some fashion. Directory service lookups for forward pointing addresses (that is, To: addresses) to the host **example.org** might be performed as follows with the following rewrite rule. The **\$F**, described in "Direction-and-Location-Specific Rewrite Rules (\$B, \$E, \$F, \$R)" causes this rule to be used only for forward pointing addresses:

```
example.org $F$[LOOKUP_IMAGE,LOOKUP,$U]
```

A forward pointing address **jdoh@example.org** will, when it matches this rewrite rule, cause **LOOKUP_IMAGE** (which is a shared library on UNIX) to be loaded into memory, and then cause the routine **LOOKUP** called with **jdoh** as the argument option. The routine **LOOKUP** might then return a different address, say, **John.Doe%eng.example.org** in the result option and the value **-1** to indicate that the rewrite rule succeeded. The percent sign in the result string (see "Repeated Rewrites Template, A%B") **John.Doe@eng.example.org** as the address to be rewritten.

On UNIX systems, the site-supplied shared library image should be world readable.

Single Field Substitutions, \$&, \$!, \$*, \$#

Single field substitutions extract a single subdomain part from the host/domain specification being rewritten. The available single field substitutions are shown in Table 9-6.

Table 9-6 Single Field Substitutions

Control Sequence	Usage
\$&n	Substitute the <i>n</i> th element, n=0,1,2,...,9, in the host specification (the part that did not match or matched a wildcard of some kind). Elements are separated by dots; the first element on the left is element zero. The rewrite fails if the requested element does not exist.
\$!n	Substitute the <i>n</i> th element, n=0,1,2,...,9, in the host specification (the part that did not match or matched a wildcard of some kind). Elements are separated by dots; the first element on the right is element zero. The rewrite fails if the requested element does not exist.
\$*n	Substitute the <i>n</i> th element, n=0,1,2,...,9, in the domain specification (the part that did match explicit text in the pattern). Elements are separated by dots; the first element on the left is element zero. The rewrite fails if the requested element does not exist.

Table 9–6 (Cont.) Single Field Substitutions

Control Sequence	Usage
<code>\$#n</code>	Substitute the <i>n</i> th element, $n=0,1,2,\dots,9$, in the domain specification (the part that did match explicit text in the pattern). Elements are separated by dots; the first element on the right is element zero. The rewrite fails if the requested element does not exist.

Suppose the address **jd@eng.example.org** matches the following rewrite rule:

```
*.example.org $U%$&0.example.org@mailhub.example.org
```

Then the result from the template will be **jd@eng.example.org** with **mailhub.example.org** used as the routing system.

Unique String Substitutions

Each use of the **\$W** control sequence inserts a text string composed of upper case letters and numbers that is designed to be unique and not repeatable. **\$W** is useful in situation where non-repeating address information must be constructed.

Source-Channel-Specific Rewrite Rules (**\$M**, **\$N**)

It is possible to have rewrite rules that act only in conjunction with specific source channels. This is useful when a short-form name has two meanings:

1. When it appears in a message arriving on one channel.
2. When it appears in a message arriving on a different channel.

Source-channel-specific rewriting is associated with the channel program in use and the channel options **rules** and **norules**. If **norules** is specified on the channel associated with an MTA component that is doing the rewriting, no channel-specific rewrite checking is done. If **rules** is specified on the channel, then channel-specific rule checks are enforced. The **rules** option is the default.

Source-channel-specific rewriting is not associated with the channel that matches a given address. It depends only on the MTA component doing the rewriting and that component's channel table entry.

Channel-specific rewrite checking is triggered by the presence of a **\$N** or **\$M** control sequence in the template part of a rule. The characters following the **\$N** or **\$M**, up until either an at sign (@), percent sign (%), or subsequent **\$N**, **\$M**, **\$Q**, **\$C**, **\$T**, or **\$?** are interpreted as a channel name.

For example, **\$Mchannel** causes the rule to fail if *channel* is not currently doing the rewriting. **\$Nchannel** causes the rule to fail if *channel* is doing the rewriting. Multiple **\$M** and **\$N** clauses may be specified. If any one of multiple **\$M** clauses matches, the rule succeeds. If any of multiple **\$N** clauses matches, the rules will fail.

Destination-Channel-Specific Rewrite Rules (**\$C**, **\$Q**)

It is possible to have rewrite rules whose application is dependent upon the channel to which a message is being enqueued. This is useful when there are two names for some host, one known to one group of hosts and one known to another. By using different channels to send mail to each group, addresses can be rewritten to refer to the host under the name known to each group.

Destination channel-specific rewriting is associated with the channel to which the message is to be dequeued and processed by, and the **rules** and **norules** options on that channel. If **norules** is specified on the destination channel, no channel-specific rewrite checking is done. If **rules** is specified on the destination channel, channel-specific rule checks are enforced. The **rules** option is the default.

Destination channel-specific rewriting is not associated with the channel matched by a given address. It depends only on the message's envelope **To:** address. When a message is enqueued, its **envelope To:** address is first rewritten to determine to which channel the message is enqueued. During the rewriting of the **envelope To:** address, any **\$C** and **\$Q** control sequences are ignored. After the **envelope To:** address is rewritten and the destination channel determined, then the **\$C** and **\$Q** control sequences are honored, as other addresses associated with the message are rewritten.

Destination-channel-specific rewrite checking is triggered by the presence of a **\$C** or **\$Q** control sequence in the template part of a rule. The characters following the **\$C** or **\$Q**, up until either an at sign (@), percent sign (%), or subsequent **\$N**, **\$M**, **\$C**, **\$Q**, **\$T**, or **\$?** are interpreted as a channel name.

For example, **\$Qchannel** causes the rule to fail if *channel* is not the destination. For another example, **\$Cchannel** causes the rule to fail if *channel* is the destination. Multiple **\$Q** and **\$C** clauses may be specified. If any one of multiple **\$Q** clauses matches, the rule succeeds. If any of multiple **\$C** clauses matches, the rule fails.

Direction-and-Location-Specific Rewrite Rules (**\$B**, **\$E**, **\$F**, **\$R**)

Sometimes you need to specify rewrite rules that apply only to envelope addresses or, alternately, only to header addresses. The control sequence **\$E** forces a rewrite to fail if the address being rewritten is not an envelope address. The control sequence **\$B** forces a rewrite to fail if the address being rewritten is not from the message header or body. These sequences have no other effects on the rewrite and may appear anywhere in the rewrite rule template.

Addresses may also be categorized by direction. A forward pointing address is one that originates on a **To:**, **Cc:**, **Resent-to:**, or other header or envelope line that refers to a destination. A backward pointing address is something like a **From:**, **Sender:**, or **Resent-From:**, that refers to a source. The control sequence **\$F** causes the rewrite to be applied if the address is forward pointing. The control sequence **\$R** causes the rewrite to be applied if the address is reverse pointing.

Host-Location-Specific Rewrites (**\$A**, **\$P**, **\$S**, **\$X**)

Circumstances occasionally require rewriting that is sensitive to the location where a host name appears in an address. Host names can appear in several different contexts in an address:

- In a source route
- To the right of the at sign (@)
- To the right of a percent sign (%) in the local-part
- To the left of an exclamation point in the local-part

Under normal circumstances, a host name should be handled in the same way, regardless of where it appears. Some situations might require specialized handling.

Four control sequences are used to control matching on the basis of the host's location in the address.

- **\$S** specifies that the rule can match a host extracted from a source route.

- **\$A** specifies that the rule can match a host found to the right of the @ sign.
- **\$P** specifies that the rule can match a host found to the right of a % sign.
- **\$X** specifies that the rule can match a host found to the left of an exclamation point (!).

The rule fails if the host is from a location other than the one specified. These sequences can be combined in a single rewrite rule. For example, if **\$S** and **\$A** are specified, the rule matches hosts specified in either a source route or to the right of the at sign. Specifying none of these sequences is equivalent to specifying all of them; the rule can match regardless of location.

Changing the Current Tag Value, \$T

The **\$T** control sequence is used to change the current rewrite rule tag. The rewrite rule tag is prepended to all rewrite rule patterns before they are looked up in the configuration file and domain database. Text following the **\$T**, up until either an at sign, percent sign, **\$N**, **\$M**, **\$Q**, **\$C**, **\$T**, or **\$?** is taken to be the new tag.

Tags are useful in handling special addressing forms where the entire nature of an address is changed when a certain component is encountered. For example, suppose that the special host name **internet**, when found in a source route, should be removed from the address and the resulting address forcibly matched against the **tcp_local-daemon** channel.

This could be implemented with rules like the following (**localhost** is assumed to be the official name of the local host):

```
internet $$SU@localhost$Tmtcp-force|  
mtcp-force|. $U%H@tcp_local-daemon
```

The first rule will match the special host name **internet** if it appears in the source route. It forcibly matches **internet** against the local channel, which insures that it will be removed from the address. A rewrite tag is then set. Rewriting proceeds, but no regular rule will match because of the tag. Finally, the default rule is tried with the tag, and the second rule of this set fires, forcibly matching the address against the **tcp_local-daemon** channel regardless of any other criteria.

Controlling Error Messages Associated with Rewriting (\$?)

The MTA provides default error messages when rewriting and channel matching fail. The ability to change these messages can be useful under certain circumstances. For example, if someone tries to send mail to an Ethernet router box, it may be considered more informative to say something like "our routers cannot accept mail" rather than the usual "illegal host/domain specified".

A special control sequence can be used to change the error message that is printed if the rule fails. The sequence **\$?** is used to specify an error message. Text following the **\$?**, up to either an at sign (@), percent sign (%), **\$N**, **\$M**, **\$Q**, **\$C**, **\\$T**, or **\$?** is taken to be the text of the error message to print if the result of this rewrite fails to match any channel. The setting of an error message is "sticky" and lasts through the rewriting process.

A rule that contains a **\$?** operates just like any other rule. The special case of a rule containing only a **\$?** and nothing else receives special attention. The rewriting process is terminated without changing the mailbox or host portions of the address and the host is looked up as-is in the channel table. This lookup is expected to fail and the error message will be returned as a result.

For example, assume the final rewrite rule in the MTA configuration file is as follows:

```
. $?Unrecognized address; contact postmaster@example.org
```

In this example, any unrecognized host or domain specifications that can fail will, in the process of failing, generate the error message: **Unrecognized address; contact postmaster@example.org**.

Handling Large Numbers of Rewrite Rules

The MTA always reads rewrite rules and stores them in memory in a hash table. Use of a compiled configuration bypasses the overhead associated with reading the configuration file each and every time the information is needed; a hash table is still used to store all of the rewrite rules in memory. This scheme is adequate for small to medium numbers of rewrite rules. However, some sites may require as many as 10,000 rewrite rules or more, which can consume prohibitive amounts of memory.

The MTA solves this problem by providing an optional facility for storing large numbers of rewrite rules in an ancillary indexed data file. Whenever the regular configuration file is read, the MTA checks for the existence of the domain database in the *MessagingServer_home/data/db/* directory. If this database exists, it is opened and consulted whenever an attempted match fails on the rules found in the configuration file. The domain database is only checked if a given rule is not found in the configuration file, so rules can always be added to the configuration file to override those in the database. By default, the domain database is used to store rewrite rules associated with hosted domains.

Testing Rewrite Rules

You can test rewrite rules with the **imsimta test -rewrite** command. The **-noimage** qualifier will allow you to test changes made to the configuration file prior to recompiling the new configuration.

You may find it helpful to rewrite a few addresses using this utility with the **-debug** qualifier. This will show you step-by-step how the address is rewritten. For example, issue the following command:

```
imsimta test -rewrite -debug joe@example.org
```

For a detailed description of the **imsimta test -rewrite** utility, see the *Messaging Server Reference*.

Rewrite Rules Example

The following example provides sample rewrite rules and how sample addresses would be rewritten by the rules.

Suppose the configuration file for the system **SC.CS.EXAMPLE.EDU** contained the rewrite rules shown in the following example:

```
sc $U@sc.cs.example.edu
sc1 $U@sc1.cs.example.edu
sc2 $U@sc2.cs.example.edu
* $U%$&0.cs.example.edu
*.cs $U%$&0.cs.example.edu
*.cs.example $U%$&0.cs.example.edu
*.cs.example.edu $U%$&0.cs.example.edu@ds.adm.example.edu
sc.cs.example.edu $U@$D
```

```

sc1.cs.example.edu $U@$D
sc2.cs.example.edu $U@$D
sd.cs.example.edu $U@sd.cs.example.edu
.example.edu $U%H.example.edu@cds.adm.example.edu
.edu $U@$H$D@gate.adm.example.edu
[] $U@[L]@gate.adm.example.edu
    
```

Table 9–7 shows some sample addresses and how they would be rewritten and routed according to the rewrite rules.

Table 9–7 Sample Addresses and Rewrites

Initial address	Rewritten as	Routed to
user@sc	user@sc.cs.example.edu	sc.cs.example.edu
user@sc1	user@sc1.cs.example.edu	sc1.cs.example.edu
user@sc2	user@sc2.cs.example.edu	sc2.cs.example.edu
user@sc.cs	user@sc.cs.example.edu	sc.cs.example.edu
user@sc1.cs	user@sc1.cs.example.edu	sc1.cs.example.edu
user@sc2.cs	user@sc2.cs.example.edu	sc2.cs.example.edu
user@sc.cs.example	user@sc.cs.example.edu	sc.cs.example.edu
user@sc1.cs.example	user@sc1.cs.example.edu	sc1.cs.example.edu
user@sc2.cs.example	user@sc2.cs.example.edu	sc2.cs.example.edu
user@sc.cs.example.edu	user@sc.cs.example.edu	sc.cs.example.edu
user@sc1.cs.example.edu	user@sc1.cs.example.edu	sc1.cs.example.edu
user@sc2.cs.example.edu	user@sc2.cs.example.edu	sc2.cs.example.edu
user@sd.cs.example.edu	user@sd.cs.example.edu	sd.cs.example.edu
user@aa.cs.example.edu	user@aa.cs.example.edu	ds.adm.example.edu
user@a.eng.example.edu	user@a.eng.example.edu	cds.adm.example.edu
user@a.cs.example.edu	user@a.cs.example.edu	gate.adm.example.edu - route inserted
user@b.cs.example.edu	user@b.cs.example.edu	gate.adm.example.edu - route inserted
user@[1.2.3.4]	user@[1.2.3.4]	gate.adm.example.edu - route inserted

Basically, what these rewrite rules say is: If the host name is one of our short-form names (**sc**, **sc1** or **sc2**) or if it is one of our full names (**sc.cs.example.edu**, and so on), expand it to our full name and route it to us. Append **cs.cmu.edu** to one part short-form names and try again. Convert one part followed by **.cs** to one part followed by **.cs.example.edu** and try again. Also convert **.cs.example** to **.cs.example.edu** and try again.

If the name is **sd.cs.example.edu** (some system we connect to directly, perhaps) rewrite and route it there. If the host name is anything else in the **.cs.example.edu** subdomain, route it to **ds.cs.example.edu** (the gateway for the **.cs.example.edu** subdomain). If the host name is anything else in the **.example.edu** subdomain route it to **cds.adm.example.edu** (the gateway for the **.example.edu** subdomain). If the host name is anything else in the **.edu** top-level domain route it to **gate.adm.example.edu** (which is presumably capable of routing the message to its proper destination). If a domain literal is used send it to **gate.adm.example.edu** as well.

Most applications of rewrite rules (like the previous example) will not change the username (or mailbox) part of the address in any way. The ability to change the username part of the address is used when the MTA is used to interface to mailers that do not conform to RFC 822 - mailers where it is necessary to stuff portions of the host/domain specification into the username part of the address. This capability should be used with great care if it is used at all.

Using Predefined Channels

When you first install Oracle Communications Messaging Server, several channels are already defined (see "Predefined Channels"). This information describes how to use predefined channel definitions in the MTA.

If you have not already read "About MTA Services," you should do so before reading this information. See "Configuring Rewrite Rules" for information about configuring the rewrite rules.

Predefined Channels

Table 10–1 lists some of the predefined channels.

Table 10–1 *Some Predefined Channels*

Channel	Definition
defaults	Used to specify which options are defaults for various channels.
l	UNIX only. Used to make routing decisions and for submitting mail using UNIX mail tools.
ims-ms	Performs final delivery of mail to the local store.
native	UNIX only. Delivers mail to /var/mail . (Messaging Server does not support /var/mail access. User must use UNIX tools to access mail from the /var/mail store.)
pipe	Used to perform delivery via a site-supplied program or script. Commands executed by the pipe channel are controlled by the administrator by using the imsimta program interface.
reprocessprocess	These channels are used for deferred, offline message processing. The reprocess channel is normally invisible as a source or destination channel. The process channel is visible like other MTA channels.
defragment	Provides the means to reassemble MIME fragmented messages.
conversion	Performs body-part-by-body-part conversions on messages flowing through the MTA.
bitbucket	Used for messages that need to be discarded.
inactive/deleted	Used to process messages for users who have been marked as inactive or deleted in the directory. Typically, bounces the message and returns custom bounce message to the sender of the message.

Table 10–1 (Cont.) Some Predefined Channels

Channel	Definition
hold	Used to hold messages for users, for example, when a user is migrated from one mail server to another.
sms	Provides support for one-way email to an SMS gateway.
tcp_localtcp_intranettcp_authtcp_submittcp_tas	Implements SMTP over TCP/IP. The multithreaded TCP SMTP channel includes a multithreaded SMTP server that runs under the control of the Dispatcher. Outgoing SMTP mail is processed by the channel program <code>tcp_smtp_client</code> , and runs as needed under the control of the Job Controller. tcp_local receives inbound messages from remote SMTP hosts. Depending on whether you use a <code>smarthost</code> /firewall configuration, either sends outbound messages directly to remote SMTP hosts or sends outbound messages to the <code>smarthost</code> /firewall system. Sometimes tcp_local gets mail from remote SMTP hosts via proxy or firewall. tcp_local is also sometimes used for internal relay activities. tcp_intranet receives and sends messages within the intranet. tcp_auth is used as a switch channel for <code>tcp_local</code> ; authenticated users switch to the tcp_auth channel to avoid relay-blocking restrictions. tcp_submit accepts message submissions, usually from user agents, on the reserved submission port 587 (see RFC 2476). tcp_tas is a special channel used by sites doing Unified Messaging.

To Deliver Messages to Programs Using the Pipe Channel

Users might want incoming mail passed to a program instead of to their mailbox. For example, users might want their incoming mail sent to a mail sorting program. The **pipe** channel performs delivery of messages using per-user, site-supplied programs.

To facilitate program delivery, you must first register programs as able to be invoked by the **pipe** channel. Do this by using the `imsimta program` utility. This utility gives a unique name to each command that you register as able to be invoked by the **pipe** channel. End users can then specify the method name as a value of their `mailprogramdeliveryinfo` LDAP attribute.

For example, to add a UNIX command `myprocmail` as a program that can be invoked by a user, you would first register the command by using the `imsimta program` utility as shown in the following example. This example registers a program called `myprocmail` that executes the program `procmail` with the arguments `-dusername` and executes as the user:

```
imsimta program -a -m myprocmail -p procmail -g "-d %s" -e user
```

Make sure the executable exists in the `programs` directory `MessagingServer_home/data/site-programs`. Make sure also that the execute permissions are set to "others."

To enable a user to access the program, the user's LDAP entry must contain the following attributes and values:

```
maildeliveryoption: program
mailprogramdeliveryinfo: myprocmail
```

Alternative delivery programs must conform to the following exit code and command-line argument restrictions:

Exit Code Restrictions. Delivery programs invoked by the **pipe** channel must return meaningful error codes so that the channel knows whether to dequeue the message, deliver for later processing, or return the message.

If the subprocess exits with an exit code of 0 (**EX_OK**), the message is presumed to have been delivered successfully and is removed from the MTA queues. If it exits with an exit code of 71, 74, 75, or 79 (**EX_OSERR**, **EX_IOERR**, **EX_TEMPFAIL**, or **EX_DB**), a temporary error is presumed to have occurred and delivery of the message is deferred. If any other exit code is returned, then the message will be returned to its originator as undeliverable. These exit codes are defined in the system header file **sysexits.h**.

Command Line Arguments. Delivery programs can have any number of fixed arguments as well as the variable argument, **%s**, representing the user name for programs executed by the user or **username+domain** for programs executed by the postmaster, "inetmail." For example, the following command line delivers a recipient's mail using the program **procmail**:

```
/usr/lib/procmail -d %s
```

To Configure the Native (/var/mail) Channel

An option file can be used to control various characteristics of the native channel. This native channel option file must be stored in the MTA configuration directory and named **native_option** (for example, *MessagingServer_homes/config/native_option*).

Option files consist of several lines. Each line contains the setting for one option. An option setting has the form:

```
<option>=<value>
```

The *value* can be either a string or an integer, depending on the option's requirements. [Table 10-2](#) describes some local channel options.

Table 10-2 Local Channel Options

Options	Descriptions
FORCE_CONTENT_LENGTH (0 or 1; UNIX only)	If FORCE_CONTENT_LENGTH=1 , then the MTA adds a Content-length: header line to messages delivered to the native channel, and causes the channel not to use the ">From" syntax when "From" is at the beginning of the line. This makes local UNIX mail compatible with Sun's newer mail tools, but potentially incompatible with other UNIX mail tools.
FORWARD_FORMAT (string)	Specifies the location of the users' .forward files. The string %u indicates that it is substituted in each user id. The string %h indicates that it is substituted in each user's home directory. The default behavior, if this option is not explicitly specified, corresponds to: FORWARD_FORMAT=%h/.forward

Table 10–2 (Cont.) Local Channel Options

Options	Descriptions
REPEAT_COUNT (integer) SLEEP_TIME (integer)	In case the user's new mail file is locked by another process when the MTA tries to deliver the new mail, these options provide a way to control the number and frequency of retries the native channel program should attempt. If the file cannot be opened after the number of retries specified, the messages remain in the native queue and the next run of the native channel attempts to deliver the new messages again. The REPEAT_COUNT option controls how many times the channel programs attempt to open the mail file before giving up. REPEAT_COUNT defaults to 30, (30 attempts). The SLEEP_TIME option controls how many seconds the channel program waits between attempts. SLEEP_TIME defaults to 2 (two seconds between retries).
SHELL_TIMEOUT (integer)	Controls the length of time in seconds the channel waits for a user's shell command in a .forward to complete. Upon such time-outs, the message are returned to the original sender with an error message resembling "Time-out waiting for <i>user's</i> shell command <i>command</i> to complete." The default is 600 (10 minutes).
SHELL_TMPDIR (directory-specific)	Controls the location where the local channel creates its temporary files when delivering to a shell command. By default, such temporary files are created in users' home directories. Using this option, the administrator may instead choose the temporary files to be created in another (single) directory. For example: SHELL_TMPDIR=/tmp

To Temporarily Hold Messages Using the Hold Channel

The hold channel is used to hold the messages of a recipient temporarily prevented from receiving new messages. Messages may be held because users' names are being changed or their mailboxes are being moved from one mailhost or domain to another. There may also be other reasons to temporarily hold messages.

When messages are to be held, they are directed to the hold channel, in the *MessagingServer_home/queue/hold* directory, using the same mechanism used to direct messages to the reprocess channel. In this way, the envelope To: addresses are unchanged. The messages are written to the hold channel queue, in the *MessagingServer_home/queue/hold* directory, as *ZZxxx.HELD* files. This prevents them from being seen by the job controller, and thus they are "held." Use the **imsimta qm dir -held** command to view a list of **.HELD** files. These messages can be selected and released by using the **imsimta qm release** command. Releasing them changes their name to *ZZxxx.00* and informs the job controller. The messages are then processed by the master program associated with the hold channel, **reprocess.exe**. Thus the message (and the To: addresses) are processed by using the normal rewriting machinery.

See "**imsimta qm**" for more information.

The Conversion Channel

The **conversion** channel enables you to perform arbitrary body part-by-body part processing on specified messages flowing through the MTA. (Note that a body part is different than a message in that a message can contain multiple body parts as, for instance, in an attachment. Also, body parts are specified and delineated by MIME headers.) This processing can be done by any site-supplied programs or command procedures and can do such things such as convert text or images from one format to another, virus scanning, language translation, and so forth. Various message types of

the MTA traffic are selected for conversion, and specific processes and programs can be specified for each type of message body part.

The prerequisite for using the **conversion** channel is understanding the concept of channels (see "Channels").

Implementing the conversion channel consists of the following high-level steps:

1. Selecting message traffic for processing
2. Specifying how different messages will be processed. These procedures are described later.

Note: A default conversion channel is automatically created in the MTA configuration. This channel can be used as is and requires no modification.

This section consists of the following sections:

- [MIME Overview](#)
- [Selecting Traffic for Conversion Processing](#)
- [To Control Conversion Processing](#)
- [To Bounce, Delete, Hold, Retry Messages Using the Conversion Channel Output](#)
- [Conversion Channel Example](#)
- [Automatic Arabic Character Set Detection](#)
- [To Automatically Detect Arabic Character Sets](#)

MIME Overview

The conversion channel makes extensive use of the MIME (Multipurpose Internet Mail Extensions) header lines. Knowledge of message construction and MIME header fields is required. For complete information on MIME, refer to <http://www.faqs.org/rfcs/>. A short overview of MIME is presented here for convenience.

Message Construction

A simple message consists of a header and a body. The header is at the top of the message and contains certain control information such as date, subject, sender, and recipient. The body is everything after the first blank line after the header. MIME specifies a way to construct more complex messages which can contain multiple body parts, and even body parts nested within body parts. Messages like these are called multi-part messages, and, as mentioned earlier, the conversion channel performs body part-by-body part processing of messages.

MIME Headers

The MIME specification defines a set of header lines for body parts. These include **MIME-Version**, **Content-type**, **Content-Transfer-Encoding**, **Content-ID**, and **Content-disposition**. The conversion channel uses the **Content-type** and **Content-disposition** headers most frequently. The following shows an example of some MIME header lines:

```
Content-type: APPLICATION/wordperfect5.1;name=Poem.wpc
Content-transfer-encoding: BASE64
Content-disposition: attachment; filename=Poem.wpc
```

Content-description: "Project documentation Draft1 wordperfect format"

Note: MIME header lines are not the same as general, non-MIME header lines such as **To:**, **Subject:** and **From:**. Basically, for Conversion channel discussion, MIME header lines start with the string **Content-**.

Content-type Header

The MIME **Content-Type** header describes the content of the body-part. The following shows an example of a **Content-Type** header format:

Content-type:*type/subtype; parameter=value; parameter=value...*

type describes the type of content of the body part. Examples of type are **Text**, **Multipart**, **Message**, **Application**, **Image**, **Audio**, and **Video**.

subtype further describes content type. Each **Content-type** has its own set of subtypes. For examples: **text/plain**, **application/octet-stream**, and **image/jpeg**. Content Subtypes for MIME mail are assigned and listed by the IANA (Internet Assigned Numbers Authority). A copy of the list is at <http://www.iana.org/assignments/media-types>.

parameter is specific to Content-type/subtype pairs. For example, the **charset** and the **name** options are shown below:

```
Content-type: text/plain; charset=us-ascii
Content-type: application/msword; name=temp.doc
```

The **charset** option specifies a character set for a textual message. The **name** option gives a suggested file name to be used if the data were to be written to a file.

Note: **Content-Type** values, **subtypes**, and option names are case-insensitive.

Content-disposition Header

The MIME **Content-disposition** header provides presentation information for the body-part. It is often added to attachments specifying whether the attachment body part should be displayed (**inline**) or presented as a file name to be copied (**attachment**). The **Content-disposition** header has the following format:

Content-disposition:*disposition_type; parameter=value;parameter=value...*

disposition_type is usually **inline** (display the body part) or **attachment** (present as file to save.) **Attachment** usually has the parameter **filename** with a value specifying the suggested name for the saved file.

For details on the **Content-disposition** header, refer to RFC2183.

Selecting Traffic for Conversion Processing

Unlike other MTA channels, the conversion channel is not normally specified in an address or MTA rewrite rule. Instead, messages are sent through the conversion channel if they meet the criteria specified in the **CONVERSIONS** mapping table. Entries to the table have the following format:

IN-CHAN=source-channel;OUT-CHAN=destination-channel;CONVERT Yes/No

As the MTA processes each message it probes the **CONVERSIONS** mapping table (if one is present). If the *source-channel* is the channel from which the message is coming

and *destination-channel* is the channel to which the message is going, then the action following **CONVERT** is taken. **Yes** means the MTA diverts the message through the conversion channel on its way to its *destination-channel*. If no match is found, or if **No** was specified, the message is queued to the regular destination channel.

If you route messages to the conversion channel by using conversion mappings, your other mappings should continue to work. However, if you use rewrite rules to route messages to the conversion channel, you might have to adjust your mappings to accommodate what you've done.

Note: An address of the form **user@conversion.localhostname** or **user@conversion** will be routed through the conversion channel, regardless of the **CONVERSIONS** mapping table.

The following example routes all non-internal messages--messages originating from, or destined to, the Internet--through the conversion channel.

CONVERSIONS

```
IN-CHAN=tcp_local;OUT-CHAN=*;CONVERT Yes
IN-CHAN=*;OUT-CHAN=tcp_local;CONVERT Yes
```

The first line specifies that messages coming from the **tcp_local** channel will be processed. The second line specifies that messages going to the **tcp_local** channel will also be processed. The **tcp_local** channel handles all messages going to and coming from the Internet. Since the default is to not go through the conversion channel, any other messages won't go through the conversion channel.

Note that this is a very basic table, and that it might not be sufficient for a site with a more customized configuration, for example, one using multiple outbound-to-the-Internet **tcp_*** channels, or using multiple inbound-from-the-Internet **tcp_*** channels.

To Control Conversion Processing

This section describes how to control conversion processing.

When a message is sent to the conversion channel, it is processed body part-by-body part. Processing is controlled by the **msconfig edit conversions** command.

Each entry consists of one or more lines containing one or more *name=value* option clauses. Where the name in the > option clauses is one of the options in [Table 9-5, "LDAP URL Substitution Sequences"](#). The values in the option clauses conform to MIME conventions. Every line except the last must end with a semicolon (;). A physical line in this file is limited to 252 characters. You can split a logical line into multiple physical lines using the back slash (\) continuation character. Entries are terminated either by a line that does not end in a semicolon, one or more blank lines, or both.

The following is a simple example of an **msconfig edit conversions** entry:

Conversions Entry

```
out-chan=ims-ms; in-type=application; in-subtype=wordperfect5.1;
out-type=application; out-subtype=msword; out-mode=block;
command="/usr/bin/convert -in=wordp -out=msword 'INPUT_FILE' 'OUTPUT_FILE' "
```

The clauses **out-chan=ims-ms; in-type=application; in-subtype=wordperfect5.1** qualify the body part. That is, they specify the type of part to be converted. The header of each part is read and its **Content-Type:** and other header information is extracted. The entries in **msconfig edit conversions** are then scanned in order from first to last; any **in-*** options present, and the **OUT-CHAN** option, if present, are checked. If all of these options match the corresponding information for the body part being processed, then the conversion specified by the **command=** or **delete=** clause is performed, and the * **out-*** options are set.

If no match occurs, then the part is matched against the next entry in **msconfig edit conversions**. Once all body parts have been scanned and processed (assuming there is a qualifying match), then the message is sent onwards to the next channel. If there are no matches, no processing occurs, and the message is sent to the next channel.

out-chan=ims-ms specifies that only message parts destined for the **ims-ms** channel will be converted. **in-type=application** and **in-subtype=wordperfect5.1** specifies that the MIME **Content-type** header for the message part must be **application/wordperfect5.1**.

Message parts can be further qualified with additional **in-*** options. (See [Table 10-7, "Conversion Options"](#).) The entry above will trigger conversion actions on a message part which has the following MIME header lines:

```
Content-type: APPLICATION/wordperfect5.1;name=Draft1.wpc
Content-transfer-encoding: BASE64
Content-disposition: attachment; filename=Draft1.wpc
Content-description: "Project documentation Draft1 wordperfect format"
```

After the three qualifying options in "[Conversions Entry](#)," the next two options, **out-type=application** and **out-subtype=mword**, specify replacement MIME header lines to be attached to the "processed" body part. **out-type=application** and **out-subtype=mword** specify that the MIME **Content-type/subtype** of the outgoing message be **application/mword**.

Note that since the **in-type** and **out-type** options are the same, **out-type=application** is not necessary since the conversion channel defaults to the original MIME labels for outgoing body parts. Additional MIME labels for outgoing body parts can be specified with additional output options.

out-mode=block ("[Conversions Entry](#)") specifies the file type that the site-supplied program will return. In other words, it specifies how the file will be stored and how the conversion channel should be read back in the returned file. For example, an html file is stored in text mode, while an **.exe** program or a zip file is stored in block/binary mode. Mode is a way of describing that the file being read is in a certain storage format.

The final option in "[Conversions Entry](#)" specifies the action to take on the body part:

```
command="/usr/bin/convert -in=wordp -out=mword 'INPUT_FILE' 'OUTPUT_FILE' "
```

The **command=** parameter specifies that a program will execute on the body part. **/usr/bin/convert** is the hypothetical command name; **-in=wordp** and **-out=mword** are hypothetical command line arguments specifying the format of the input text and output text; **INPUT_FILE** and **OUTPUT_FILE** are conversion channel environmental variables (see "[To Use Conversion Channel Environmental Variables](#)") program should store its converted body part.

Note: Envelope originator and recipient information is now provided as `x-envelope-from` and `x-envelope-to` fields respectively when a file containing the outer message header is requested by a regular conversion entry.

Instead of executing a command on the body part, the message part can simply be deleted by substituting `DELETE=1` in place of the **command** option.

Note: Whenever **conversions** is modified, you must recompile the configuration if running a compiled configuration (see "[Compiling the MTA Configuration](#)").

Conversion Channel Information Flow

The flow of information is as follows: a message containing body parts comes into the conversion channel. The conversion channel parses the message, and processes the parts one by one. The conversion channel then qualifies the body part, that is, it determines if it should be processed or not by comparing its MIME header lines to the *qualifying parameters* (see [Table 10-7, "Conversion Options"](#)). If the body part qualifies, the conversion processing commences.

If MIME or body part information is to be passed to the conversion script, it is stored in an environmental variable (see "[To Use Conversion Channel Environmental Variables](#)") as specified by *information passing parameters* (see [Table 10-7, "Conversion Options"](#)).

At this point, an action specified by an *action parameter*, (see [Table 10-7, "Conversion Options"](#)) is taken on the body part. Typically the action is that the body part be deleted or that it be passed to a program wrapped in a script. The script processes the body part and then sends it back to the conversion channel for reassembling into the post-processed message. The script can also send information to the conversion channel by using the conversion channel *output options* (see [Table 10-5, "Conversion Channel Output Options"](#)). This can be information such as new MIME header lines to add to the output body part, error text to be returned to the message sender, or special directives instructing the MTA to initiate some action such as bounce, delete, or hold a message.

Finally, the conversion channel replaces the header lines for the output body part as specified by the *output parameters* (see [Table 10-7, "Conversion Options"](#)).

To Use Conversion Channel Environmental Variables

When operating on message body parts, it is often useful to pass MIME header line information, or entire body parts, to and from the site-supplied program. For example, a program may require **Content-type** and **Content-disposition** header line information as well as a message body part. Typically a site-supplied program's main input is a message body part which is read from a file. After processing the body part, the program will need to write it to a file from which the conversion channel can read it. This type of information passing is done by using conversion channel environmental variables.

Environmental variables can be created in **conversions** using the **parameter-symbol-*** parameter or by using a set of pre-defined conversion channel environmental variables (see [Table 10-5, "Conversion Channel Output Options"](#)).

The following **conversions** entry and incoming header show how to pass MIME information to the site-supplied program using environment variables.

msconfig edit conversions entry:

```
in-channel=*; in-type=application; in-subtype=*;
parameter-symbol-0=NAME; parameter-copy-0=*;
dparameter-symbol-0=FILENAME; dparameter-copy-0=*;
message-header-file=2; original-header-file=1;
override-header-file=1; override-option-file=1;
command="/bin/viro-scan500.sh "INPUT_FILE' "OUTPUT_FILE' "
```

Incoming header:

```
Content-type: APPLICATION/msword; name=Draft1.doc
Content-transfer-encoding: BASE64
Content-disposition: attachment; filename=Draft1.doc
Content-description: "Project documentation Draft1 msword format"
```

in-channel=*; **in-type=application**; **in-subtype=*** specify that a message body part from any input channel of type **application** will be processed.

parameter-symbol-0=NAME specifies that value of the **Content-type** parameter **name**, if present (**Draft1.doc** in our example), be stored in an environment variable called **NAME**.

parameter-copy-0=* specifies that all **Content-type** parameters of the input body part be copied to the output body part.

dparameter-symbol-0=FILENAME specifies that the value of the **Content-disposition** parameter **filename** (**Draft1.doc** in our example), be stored in an environment variable called **FILENAME**.

dparameter-copy-0=* specifies that all **Content-disposition** parameters of the input body part be copied to the output body part.

message-header-file=2 specifies that the original header of the message as a whole (the outermost message header) be written to the file specified by the environment variable **MESSAGE_HEADERS**.

original-header-file=1 specifies that the original header of the enclosing **MESSAGE/RFC822** part are written to the file specified by the environment variable **INPUT_HEADERS**.

override-header-file=1 specifies that MIME headers are read from the file specified by environmental variable **OUTPUT_HEADERS**, overriding the original MIME header lines in the enclosing MIME part. **\$OUTPUT_HEADERS** is an on-the-fly temporary file created at the time conversion runs. A site-supplied program would use this file to store MIME header lines changed during the conversion process. The conversion channel would then read the MIME header lines from this file when it re-assembles the body part. Note that only MIME header lines can be modified. Other general, non-MIME header lines cannot be altered by the conversion channel.

override-option-file=1 specifies that the conversion channel read *conversion channel options* from the file named by the **OUTPUT_OPTIONS** environmental variable. See ["To Use Conversion Channel Output Options."](#)

command="MessagingServer_home/bin/viro-scan500.sh" specifies the command to execute on the message body part.

Table 10-3 describes conversion channel environment variables.

Table 10–3 Conversion Channel Environment Variables

Environment Variable	Description
ATTACHMENT_NUMBER	Attachment number for the current part. This has the same format as the ATTACHMENT-NUMBER conversion match option.
CONVERSION_TAG	The current list of active conversion tags. This corresponds to the TAG conversion match option.
INPUT_CHANNEL	The channel that enqueued the message to the conversion channel. This corresponds to the IN-CHANNEL conversion match option.
INPUT_ENCODING	Encoding originally present on the body part.
INPUT_FILE	Name of the file containing the original body part. The site-supplied program should read this file.
INPUT_HEADERS	Name of the file containing the original header lines for the body part. The site-supplied program should read this file.
INPUT_TYPE	MIME Content-type of the input message part.
INPUT_SUBTYPE	MIME content subtype of the input message part.
INPUT_DESCRIPTION	MIME content-description of the input message part.
INPUT_DISPOSITION	MIME content-disposition of the input message part.
MESSAGE_HEADERS	Name of the file containing the original outermost header for an enclosing message (not just the body part) or the header for the part's most immediately enclosing MESSAGE/RFC822 part. The site-supplied program should read this file.
OUTPUT_CHANNEL	The channel the message is headed for. This corresponds to the OUT-CHANNEL conversion match option.
OUTPUT_FILE	Name of the file where the site-supplied program should store its output. The site-supplied program should create and write this file.
OUTPUT_HEADERS	Name of the file where the site-supplied program should store MIME header lines for an enclosing part. The site-supplied program should create and write this file. Note that file should contain actual MIME header lines (not option=value lines) followed by a blank line as its final line. Note also that only MIME header lines can be modified. Other general, non-MIME header lines cannot be altered by the conversion channel.
OUTPUT_OPTIONS	Name of the file from which the site-supplied program should read conversion channel options. See " To Use Conversion Channel Output Options. "
PART_NUMBER	The part number for the current part. This has the same format as the PART-NUMBER conversion match option.
PART_SIZE	The size in bytes of the part being processed.

Mail Conversion Tags

Mail conversion tags are special tags which are associated with a particular recipient or sender. When a message is being delivered, the tag is visible to the conversion channel program, which may make use of it for special processing. Conversion tags are stored in the LDAP directory.

Mail conversion tags could be used as follows: the administrator can set up selected users with a mail conversion tag value of **harmonica**. The administrator then has a

conversion channel setup which, when processing that mail, will detect the presence of the tag and the value of **harmonica**. When that happens, the program will perform some arbitrary function.

Mail conversion tags can be set on a per user or a per domain basis. The recipient LDAP attribute at the domain level is **MailDomainConversionTag** (modifiable with the MTA option **LDAP_DOMAIN_ATTR_CONVERSION_TAG**). At the user level it is **MailConversionTag** (modifiable with the MTA option **LDAP_CONVERSION_TAG**). Both of these attributes can be multivalued with each value specifying a different tag. The set of tags associated with a given recipient is cumulative, that is, tags set at the domain level are combined with tags set at the user level.

Sender-based conversion tags can be set with the MTA options **LDAP_SOURCE_CONVERSION_TAG** and **LDAP_DOMAIN_ATTR_SOURCE_CONVERSION_TAG**, which specify user and domain level LDAP attributes respectively for conversion tags associated with these source address. There is no default attribute for either of these options.

Two new actions are available to system Sieves: **addconversiontag** and **setconversiontag**. Both accept a single argument: A string or list of conversion tags. **addconversiontag** adds the conversion tag(s) to the current list of tags while **setconversiontag** empties the existing list before adding the new ones. Note that these actions are performed very late in the game so **setconversiontag** can be used to undo all other conversion tag setting mechanisms. These allow you put conversion tags in the Sieves filters.

The Sieve envelope test accepts **conversiontag** as an envelope field specifier value. The test checks the current list of tags, one at a time. Note that the **:count** modifier, if specified, allows checking of the number of active conversion tags. This type of envelope test is restricted to system Sieves. Also note that this test only sees the set of tags that were present prior to Sieve processing--the effects of **setconversiontag** and **addconversiontag** actions are not visible.

Including Conversion Tag Information in Various Mapping Probes

A new MTA option, **INCLUDE_CONVERSIONTAG**, has been added to selectively enable the inclusion of conversion tag information in various mapping probes. This is a bit-encoded value. [Table 10-4](#) shows the bits assigned to the various mapping probes. In all cases the current set of tags appears in the probe as a comma separated list.

Table 10-4 Including Conversion Tag Information in Various Mapping Probes

Position	Value	Mapping
0	1	CHARSET_CONVERSION - added as ;TAG= field before ;CONVERT.
1	2	CONVERSION - added as ;TAG= field before ;CONVERT
2	4	FORWARD - added just before current address (delim)
3	8	ORIG_SEND_ACCESS - added at end of probe (delim)
4	16	SEND_ACCESS - added at end of probe (delim)
5	32	ORIG_MAIL_ACCESS - added at end of probe (delim)
6	64	MAIL_ACCESS - added at end of probe (delim)

To Use Conversion Channel Output Options

Conversion channel output options ([Table 10-5](#), "Conversion Channel Output Options") are dynamic variables used to pass information and special directives from

the conversion script to the conversion channel. For example, during body part processing the script may want to send a special directive asking the conversion channel to bounce the message and to add some error text to the returned message stating that the message contained a virus.

The output options are initiated by setting **OVERRIDE-OPTION-FILE=1** in the desired conversion entry. Output options are then set by the script as needed and stored in the environmental variable file, **OUTPUT_OPTIONS**. When the script is finished processing the body part, the conversion channel reads the options from the **OUTPUT_OPTIONS** file.

The **OUTPUT_OPTION** variable is the name of the file from which the conversion channel reads options. Typically it is used as an on-the-fly temporary file to pass information. The example below shows a script that uses output options to return an error message to a sender who mailed a virus.

```
/usr/local/bin/viro_screen2k $INPUT_FILE # run the virus screener

if [ $? -eq 1 ]; then
echo "OUTPUT_DIAGNOSTIC='Virus found and deleted.'" > $OUTPUT_OPTIONS
echo "STATUS=178029946" >> $OUTPUT_OPTIONS
else
cp $INPUT_FILE $OUTPUT_FILE # Message part is OK
fi
```

In this example, the system diagnostic message and status code are added to the file defined by **\$OUTPUT_OPTIONS**. If you read the **\$OUTPUT_OPTIONS** temporary file out you would see something like:

```
OUTPUT_DIAGNOSTIC="Virus found and deleted."
STATUS=178029946
```

The line **OUTPUT_DIAGNOSTIC='Virus found and deleted'** tells the conversion channel to add the text **Virus found and deleted** to the message.

178029946 is the **PMDF_FORCERETURN** status per the *pmdf_err.h* file which is found in the *_MessagingServer_home/include/deprecated/pmdf_err.h*. This status code directs the conversion channel to bounce the message back to the sender. (For more information on using special directives refer to "[To Bounce, Delete, Hold, Retry Messages Using the Conversion Channel Output.](#)")

A complete list of the output options is shown in [Table 10-5](#).

Table 10-5 Conversion Channel Output Options

Option	Description
OUTPUT_TYPE	MIME content type of the output message part.
OUTPUT_SUBTYPE	MIME content subtype of the output message part.
OUTPUT_DESCRIPTION	MIME content description of the output message part.
OUTPUT_DIAGNOSTIC	Text to include as part of the message sent to the sender if a message is forcibly bounced by the conversion channel.
OUTPUT_DISPOSITION	MIME content-disposition of the output message part.
OUTPUT_ENCODING	MIME content transfer encoding to use on the output message part.

Table 10-5 (Cont.) Conversion Channel Output Options

Option	Description
OUTPUT_MODE	MIME Mode with which the conversion channel should write the output message part, hence the mode with which recipients should read the output message part.
STATUS	Exit status for the converter. This is typically a special directive initiating some action by the conversion channel. A complete list of directives can be viewed in <i>MessagingServer_home/include/deprecated/pmdf_err.h</i>

Headers in an Enclosing MESSAGE/RFC822 Part

When performing conversions on a message part, the conversion channel has access to the header in an enclosing **MESSAGE/RFC822** part, or to the message header if there is no enclosing **MESSAGE/RFC822** part. Information in the header may be useful for the site-supplied program.

If an entry is selected that has **ORIGINAL-HEADER-FILE=1**, then all the original header lines of the enclosing **MESSAGE/RFC822** part are written to the file represented by the **ORIGINAL_HEADERS** environment variable. If **OVERRIDE-HEADER-FILE=1**, then the conversion channel will read and use as the header on that enclosing part the contents of the file represented by the **ORIGINAL_HEADERS** environment variable.

To Call Out to a Mapping Table from a Conversion Entry

out-parameter-* values may be stored and retrieved in an arbitrarily named mapping table. This feature is useful for renaming attachments sent by clients that send all attachments with a generic name like **att.dat** regardless of whether they are **postscript**, **mword**, **text**, or whatever. This is a generic way to relabel the part so that other clients (Outlook for example) are able to open the part by reading the extension.

The syntax for retrieving a parameter value from a mapping table is as follows:

"mapping-table-name:mapping-input[\$Y, \$N]"

\$Y returns a parameter value. If there is no match found or the match returns **\$N**, then that parameter in the conversions entry is ignored or treated as a blank string. Lack of a match or a **\$N** does not cause the conversion entry itself to be aborted.

Consider the following mapping table:

```
X-ATT-NAMES

postscript temp.PS$Y
wordperfect5.1 temp.WPC$Y
mword temp.DOC$Y
```

The following conversion entry for the above mapping table results in substituting generic file names in place of specific file names on attachments:

```
out-chan=tcp_local; in-type=application; in-subtype=*;
in-parameter-name-0=name; in-parameter-value-0=*;
out-type=application; out-subtype='INPUT-SUBTYPE';
out-parameter-name-0=name;
out-parameter-value-0="'X-ATT-NAMES:\\'INPUT_SUBTYPE\\'";
command="cp "INPUT_FILE" "OUTPUT_FILE"
```

In the example above, **out-chan=tcp_local; in-type=application; in-subtype=*** specifies that a message to be processed must come from the **tcp_local** channel with the **content-type** header of **application/*** (* specifies that any subtype would do).

in-parameter-name-0=name; in parameter-value-0=* additionally specifies that the message must have a content-type parameter called *** name=*** and that any value for that parameter will be accepted (again, * specifies that any parameter value would do.)

out-type=application; specifies that the MIME **Content-type** parameter for the post-processing message be **application**.

out-subtype='INPUT-SUBTYPE'; specifies that the MIME **subtype** parameter for the post-processing body part be the **INPUT-SUBTYPE** environmental variable, which is the original value of the input **subtype**. Thus, if you wanted change

Content-type: application/xxxx; name=foo.doc

to

Content-type: application/msword; name=foo.doc

then you would use

out-type=application; out-subtype=msword

out-parameter-name-0=name; specifies that the output body part will have a MIME **Content-type name=** parameter.

```
out-parameter-value-0='X-ATT-NAMES:\\' INPUT_SUBTYPE\\'';
```

says to take the value of the **INPUT_SUBTYPE** variable (that is, the original content-type header subtype value of the original body part) and search the mapping table **X-ATT-NAMES**. If a match is found, the content-type parameter specified by **out-parameter-name-0** (that is, **name**) receives the new value specified in the **X-ATT-NAMES** mapping table. Thus, if the original subtype was **msword**, the value of the name parameter will be **temp.DOC**.

To Bounce, Delete, Hold, Retry Messages Using the Conversion Channel Output

This section describes how to use the conversion channel options to bounce, delete, or hold messages. The basic procedure is as follows:

1. Set **OVERRIDE-OPTION-FILE=1** in the appropriate conversions entry. This tells the conversion channel to read the output options from the **OUTPUT_OPTIONS** file.
2. Use the conversion script to determine what action is required on a particular message body part.
3. In the script, specify the special directive for that action by writing the **STATUS=directive-code** option in the **OUTPUT_OPTIONS** file.

A complete listing of special directives can be found in *MessagingServer_home/include/deprecated/pmdf_err.h*. The ones commonly used by the conversion channel can be found in [Table 10-6](#).

Table 10-6 Special Directives Commonly Used by the Conversion Channel

NAME	Hex Value	Decimal Value
PMDF_FORCEHOLD	0x0A9C86AA	178030250
PMDF_FORCERETURN	0x0A9C857A	178029946
PMDF_FORCEDELETE	0x0A9C8662	178030178

Table 10–6 (Cont.) Special Directives Commonly Used by the Conversion Channel

NAME	Hex Value	Decimal Value
PMDF_ FORCEDISCARD	0x0A9C86B3	178030259
PMDF__AGN	0x0A9C809A	178028698

The functions of these directives will be explained by using examples.

To Bounce Messages

To bounce a message using the conversion channel set **OVERRIDE-OPTION-FILE=1** in the appropriate **conversions** file entry and add the following line to your conversion script:

```
echo "STATUS=178029946" >> $OUTPUT_OPTIONS
```

If you wish to add a short text string to the bounced message add the following line to the conversion script:

```
echo OUTPUT_DIAGNOSTIC=text-string>> $OUTPUT_OPTIONS
```

where text string is something like: **"The message sent from your machine contained a virus which has been removed. Be careful about executing email attachments."**

To Conditionally Delete a Message or Its Parts

It may be useful to delete parts conditionally, depending on what they contain. This can be done by using the output options. By contrast, the **DELETE=1** conversion parameter clause unconditionally deletes a message part.

To delete a message part using the output options, set **OVERRIDE-OPTION-FILE=1** in the appropriate conversions entry and add the following line to your conversion script:

```
echo "STATUS=178030178" >> $OUTPUT_OPTIONS
```

Similarly, to delete the entire message you could use:

```
echo "STATUS=178030259" >> $OUTPUT_OPTIONS
```

To Hold a Message

It may be useful to hold messages conditionally, depending on what they contain. To delete a message part using the output options, set **OVERRIDE-OPTION-FILE=1** in the appropriate conversions entry and add the following line to your conversion script:

```
echo "STATUS=178030250" >> $OUTPUT_OPTIONS
```

This requests that the conversion channel hold the message as a **.HELD** file in the conversion channel queue.

To Cause Messages to Be Reprocessed

When a converter script encounters a temporary resource problem (for example, the system can't connect to an external server, a needed file is locked, and so on), you can use **PMDF__AGN** to tell the conversion channel to consider processing messages that have encountered a temporary error. The MTA will record a "Q" status message in the **mail.log_current** file, retain the message in the conversion channel, and retry the processing later.

Add the following line to your conversion script:

```
echo "STATUS=178028698" >> $OUTPUT_OPTIONS
```

Conversion Channel Example

The CONVERSIONS mapping and set of conversion rules seen in the following examples cause GIF, JPEG, and BITMAP files sent to the hypothetical channel `tcp_docuprint` to be converted into PostScript automatically. Several of these conversions use the hypothetical `/usr/bin/ps-converter.sh` to make that transformation. An additional rule that converts WordPerfect 5.1 files into Microsoft Word files is included.

```
CONVERSIONS
```

```
IN-CHAN=*;OUT-CHAN=tcp_docuprint;CONVERT Yes
```

```
out-chan=ims-ms; in-type=application; in-subtype=wordperfect5.1;
out-type=application; out-subtype=microsoftword; out-mode=block;
command="/bin/doc-convert -in=wp -out=msw 'INPUT_FILE' 'OUTPUT_FILE'"
```

```
out-chan=tcp_docuprint; in-type=image; in-subtype=gif;
out-type=application; out-subtype=postscript; out-mode=text;
command="/bin/ps-convert -in=gif -out=ps 'INPUT_FILE' 'OUTPUT_FILE'"
```

```
out-chan=tcp_docuprint; in-type=image; in-subtype=jpeg;
out-type=application; out-subtype=postscript; out-mode=text;
command="/bin/ps-convert -in=jpeg -out=ps 'INPUT_FILE' 'OUTPUT_FILE'"
```

```
out-chan=tcp_docuprint; in-type=image; in-subtype=bitmap;
out-type=application; out-subtype=postscript; out-mode=text;
command="/bin/ps-convert -in=bmp -out=ps 'INPUT_FILE' 'OUTPUT_FILE'"
```

The conversion options are described in [Table 10-7](#):

Table 10-7 Conversion Options

Option	Description
<i>Part 1: Qualifying options (Specifies the options for which the message must match before it will be converted.)</i>	Not Applicable
OUT-CHAN,OUT-CHANNEL	Output channel to match for conversion (wildcards allowed). The conversion specified by this entry is performed only if the message is destined for this specified channel.
IN-CHAN,IN-CHANNEL	Input channel to match for conversion (wildcards allowed). The conversion specified by this entry is only performed if the message is coming from the specified channel.
IN-TYPE	Input MIME type to match for conversion (wildcards allowed). The conversion specified is performed only if this field matches the MIME type of the body part.
IN-SUBTYPE	Input MIME subtype to match for conversion (wildcards allowed). The conversion specified by this entry is performed only if this field matches the MIME subtype of the body part.

Table 10-7 (Cont.) Conversion Options

Option	Description
IN-PARAMETER-NAME <i>-n</i>	Specifies the name of the Input MIME Content-Type parameter that must match for conversion; The $n = 0, 1, 2, \dots$ is used to optionally pair the specified parameter name requirement with a value required by using IN-PARAMETER-VALUE-n with the same value of n .
IN-PARAMETER-VALUE <i>-n</i>	Specifies the value required of the input MIME Content-Type parameter whose name is specified in the corresponding IN-PARAMETER-NAME-n . The conversion specified by this entry is performed only if the input body part has the content-type parameter specified by the corresponding IN-PARAMETER-NAME-n and its value matches the value of this parameter. Wildcards allowed.
IN-PARAMETER-DEFAULT <i>-n</i>	Default value to use if the input MIME Content-Type parameter specified by the corresponding IN-PARAMETER-NAME-n is not present.
IN-DISPOSITION	Input MIME Content-Disposition to match for conversion.
IN-DPARAMETER-NAME <i>-n</i>	Specifies the name of the Input MIME Content-Disposition parameter that must match for conversion; The $n = 0, 1, 2, \dots$ is used to optionally pair the specified parameter name requirement with a value required by using IN-DPARAMETER-VALUE-n with the same value of n .
IN-DPARAMETER-VALUE <i>-n</i>	Specifies the value required of the input MIME Content-Disposition parameter whose name is specified in the corresponding IN-DPARAMETER-NAME-n . The conversion specified by this entry is performed only if the input body part has the Content-Disposition parameter specified by the corresponding IN-DPARAMETER-NAME-n and its value matches the value of this parameter. Wildcards allowed.
IN-DPARAMETER-DEFAULT <i>-n</i>	Default value to use if the input MIME Content-Disposition parameter specified by the corresponding IN-DPARAMETER-NAME-n is not present.
IN-DESCRIPTION	Input MIME Content-Description to match for conversion.
IN-SUBJECT	Input Subject from enclosing MESSAGE/RFC822 part.
TAG	Input tag, as set by a mail list CONVERSION_TAG parameter.
<i>Part 2: Output Parameters (Specify the body part's post-conversion output settings.)</i>	Not Applicable
OUT-TYPE	Output MIME type if it is different than the input type.
OUT-SUBTYPE	Output MIME subtype if it is different than the input subtype.
OUT-PARAMETER-NAME <i>-n</i>	Specifies the name of a content-type parameter which will be set on the output body part.
OUT-PARAMETER-VALUE <i>-n</i>	Output MIME Content-Type parameter value corresponding to OUT-PARAMETER-NAME-n .
PARAMETER-COPY <i>-n</i>	Specifies the name of a content-type parameter which should be copied from the input body part to the output body part.
OUT-DISPOSITION	Output MIME Content-Disposition if it is different than the input MIME Content-Disposition .
OUT-DPARAMETER-NAME <i>-n</i>	Output MIME Content-Disposition parameter name; $n=0, 1, 2, \dots$

Table 10-7 (Cont.) Conversion Options

Option	Description
OUT-DPARAMETER-VALUE-<i>n</i>	Output MIME Content-Disposition parameter value corresponding to OUT-DPARAMETER-NAME-<i>n</i> .
DPARAMETER-COPY-<i>n</i>	A list of the Content-Disposition : parameters to copy from the input body part's Content-Disposition : parameter list to the output body part's Content-Disposition : parameter list; <i>n</i> = 0, 1, 2,... Takes as argument the name of the MIME parameter to copy, as matched by an IN-PARAMETER-NAME-<i>n</i> clause. Wildcards may be used in the argument. In particular, an argument of * means to copy all the original Content-Disposition : parameters.
OUT-DESCRIPTION	Output MIME Content-Description if it is different than the input MIME Content-Description .
OUT-MODE	Mode in which to read and store the converted file. This should be BLOCK (binaries and executables) or TEXT .
OUT-ENCODING	Encoding to apply to the converted file when the message is reassembled.
<i>Part 3: Action Parameters (Specify an action to take on a message part.)</i>	Not Applicable
COMMAND	Command to execute to perform conversion. Command to execute to perform conversion. This parameter is required; if no command is specified, the entry is ignored. Use / to specify paths, not \. Example: command="D:/tmp/mybat.bat"
DELETE	0 or 1. If this flag is set, the message part is deleted. (If this is the only part in a message, then a single empty text part is substituted.)
RELABEL	RELABEL=1 will relabel the MIME label to whatever is specified by the Output parameters. Relabel=0 does nothing. Usually relabelling is done on mislabeled parts (example: from Content-type: application/octet-stream to Content-type: application/msword) so users can "doubleclick" to open a part, rather than having to save the part to a file and open it with a program.
SERVICE-COMMAND	SERVICE-COMMAND=command will execute a site-supplied procedure that will operate on entire MIME message (MIME headers and content body part). Also, unlike other CHARSET-CONVERSION operations or conversion channel operations, the service-command are expected to do their own MIME disassembly, decoding, re-encoding, and reassembly. Note that this flag causes an entry to be ignored during conversion channel processing; SERVICE-COMMAND entries are instead performed during character set conversion processing. Use / to specify paths, not \. Example: command="D:/tmp/mybat.bat"
<i>Part 4: Information Passing Parameters (Used to pass information to and from the site-supplied program.)</i>	Not Applicable
DPARAMETER-SYMBOL-<i>n</i>	Environment variable into which the Content-disposition parameter value, if present, will be stored; <i>n</i> = 0, 1, 2,... Each DPARAMETER-SYMBOL-<i>n</i> is extracted from the Content-Disposition : parameter list in order (<i>n</i> =0 is first parameter, <i>n</i> =2 second, etc.) and placed in the specified environment variable prior to executing the site-supplied program.

Table 10-7 (Cont.) Conversion Options

Option	Description
PARAMETER-SYMBOL-<i>n</i>	Specifies the name of a content-type parameter which, if present in the input body part, its value will be stored in an environment variable of the same name. If the parameter does not exist in the input body part, the environment variable will not exist in the process. For example, if you specify parameter-symbol-0=foo , and there's a content type parameter foo with value bar, you end up with an environment variable foo with value bar. Environment variable into which the Content-Type parameter value, if present, will be stored; <i>n</i> = 0, 1, 2... Each PARAMETER-SYMBOL-<i>n</i> is extracted from the Content-Type: parameter list in order (<i>n</i> =0 is first parameter, <i>n</i> =2 second, etc.) and placed in an environment variable of the same name prior to executing the site-supplied program. Takes as argument the variable name into which the MIME parameter to convert, as matched by an IN-PARAMETER-NAME-<i>n</i> clause.
MESSAGE-HEADER-FILE	If set to 1, the original header of the immediately enclosing body part are written to the file specified by the environmental variable MESSAGE_HEADERS . If set to 2, the original header of the message as a whole (the outermost message header) are written to the file.
ORIGINAL-HEADER-FILE	0 or 1. If set to 1, the original header of the enclosing MESSAGE/RFC822 part (not just the body part) are written to the file represented by the environmental variable ORIGINAL_HEADERS .
OVERRIDE-HEADER-FILE	0 or 1. If set to 1, then MIME header lines are read by the conversion channel from the environmental variable OUTPUT_HEADERS , overriding the original header lines in the enclosing MIME part.
OVERRIDE-OPTION-FILE	If OVERRIDE-OPTION-FILE=1 , the conversion channel reads options from the OUTPUT_OPTIONS environmental variable.
PART-NUMBER	Dotted integers: <i>a. b. c...</i> The part number of the MIME body part.

Automatic Arabic Character Set Detection

A new **auto_ef** program was added to automatically detect Arabic character sets.

You can call the **auto_ef** program from the conversion channel to automatically detect and label most unlabeled or incorrectly labeled text messages in Arabic character sets. These unlabeled or mislabeled messages are usually sent from Yahoo or Hotmail in Arabic.

Without the correct character set labeling, many mail clients cannot display the messages correctly.

If a message has MIME content-type headers, the **auto_ef** program examines and processes only those with text/plain content type. If the message is not labeled with a MIME content-type header, then **auto_ef** adds a text/plain content-type unconditionally.

To activate or enable this program, you must:

To Automatically Detect Arabic Character Sets

1. Edit your mappings using **msconfig edit mappings** to enable a conversion channel for the source and destination channel of your choosing. To enable a

conversion channel for all mail coming in from the Internet to your local users, add a section to your mappings similar to the following:

```
CONVERSIONS
```

```
IN-CHAN=tcp*;OUT-CHAN=ims-ms;CONVERT YES
```

Note that the **IN** and **OUT** channels depend on your configuration. If you are deploying on a relay MTA, you must modify the channels to fit your configuration. For example, **IN-CHAN=tcp*;OUT-CHAN=tcp*;CONVERT YES**. Or, you could turn it on for all channels as follows: **IN-CHAN=;OUT-CHAN=;CONVERT YES**

2. Create a conversions entry by running **msconfig edit conversions** that contains the following:

```
!
in-channel=*; out-channel=*;
in-type=text; in-subtype=*;
parameter-copy-0=*; dparameter-copy-0=*;
original-header-file=1; override-header-file=1;
command="_MessagingServer_home_
/lib/arabicdetect.sh"
!
```

3. If you are running a compiled configuration, compile your MTA configuration with the following command: ***MessagingServer_home/bin/imsimta cbuild***
4. Restart with the command: ***MessagingServer_home/bin/imsimta restart***

Character Set Conversion and Message Reformatting

This section describes character set, formatting, and labelling conversions performed internally by the MTA. Note that some of the examples in this section use old or obsolete technology like DEC VMS, or the **d** channels. Although these technologies are old or obsolete, this does not make the examples DEC- or **d channel**-specific. The examples are still valid in describing how the conversion technology works. We may update the examples in a later release.

One very basic mapping table in Messaging Server is the character set conversion table. The name of this table is **CHARSET-CONVERSION**. It is used to specify what sorts of channel-to-channel character set conversions and message reformatting should be done.

On many systems there is no need to do character set conversions or message reformatting and therefore this table is not needed. Situations arise, however, where character conversions must be done. For example, sites running Japanese OpenVMS may need to convert between DEC Kanji and the ISO-2022 Kanji currently used on the Internet. Another possible use of conversions arises when multinational characters are so heavily used that the slight discrepancies between the DEC Multinational Character Set (DEC-MCS) and the ISO-8859-1 character set specified for use in MIME may become an issue, and actual conversion between the two may therefore be needed.

The **CHARSET-CONVERSION** mapping table can also be used to alter the format of messages. Facilities are provided to convert a number of non-MIME formats into MIME. Changes to MIME encodings and structure are also possible. These options are used when messages are being relayed to systems that only support MIME or some subset of MIME. And finally, conversion from MIME into non-MIME formats is provided in a small number of cases.

The MTA will probe the **CHARSET-CONVERSION** mapping table in two different ways. The first probe is used to determine whether or not the MTA should reformat the message and if so, what formatting options should be used. (If no reformatting is specified, the MTA does not bother to check for specific character set conversions.) The input string for this first probe has the general form:

```
IN-CHAN=_in-channel_;OUT-CHAN=_out-channel_;CONVERT
```

Here *in-channel* is the name of the source channel (where the message comes from) and *out-channel* is the name of the destination channel (where the message is going). If a match occurs the resulting string should be a comma-separated list of options.

Table 10–8 lists the options.

Table 10–8 CHARSET-CONVERSION Mapping Table Options

Option	Description
Always	Force conversion even the message is going to be passed through the conversion channel before going to <i>out-channel</i> .
Appledouble	Convert other MacMIME formats to Appledouble format.
Applesingle	Convert other MacMIME formats to Applesingle format.
BASE64	Switch MIME encodings to BASE64. This option only applies to message parts that are already encoded. Messages with Content-transfer-encoding: 7BIT or 8bit do not require any special encoding and therefore this BASE64 option will have no effect on them.
Binhex	Convert other MacMIME formats, or parts including Macintosh type and Mac creator information, to Binhex format.
Block	Extract just the data fork from MacMIME format parts.
Bottom	"Flatten" any message/rfc822 body part (forwarded message) into a message content part and a header part.
Delete	"Flatten" any message/rfc822 body part (forwarded message) into a message content part, deleting the forwarded headers.
Level	Remove redundant multipart levels from message.
Macbinary	Convert other MacMIME formats, or parts including Macintosh type and Macintosh creator information, to Macbinary format.
No	Disable conversion.
QUOTED-PRINTABLE	Switch MIME encodings to QUOTED-PRINTABLE.
Record,Text	Line wrap text/plain parts at 80 characters.
Record,Text= n	Line wrap text/plain parts at n characters.
RFC1154	Convert message to RFC 1154 format.
Top	"Flatten" any message/rfc822 body part (forwarded message) into a header part and a message content part.
UUENCODE	Switch MIME encodings to X-UUENCODE.
Yes	Enable conversion.

Character Set Conversion

If the MTA probes and finds that the message is to be reformatted, it will proceed to check each part of the message. Any text parts are found and their character set parameters are used to generate the second probe. Only when the MTA has checked

and found that conversions may be needed does it ever perform the second probe. The input string in this second case looks like this:

```
IN-CHAN=_in-channel_;OUT-CHAN=_out-channel_;IN-CHARSET=_in-char-set_
```

The *in-channel* and *out-channel* are the same as before, and the *in-char-set* is the name of the character set associated with the particular part in question. If no match occurs for this second probe, no character set conversion is performed (although message reformatting, for example, changes to MIME structure, may be performed in accordance with the option matched on the first probe). If a match does occur it should produce a string of the form:

```
OUT-CHARSET=_out-char-set_
```

Here *out-char-set* specifies the name of the character set to which the **in-char-set** should be converted. Note that both of these character sets must be defined in the character set definition table, **charsets.txt**, located in the MTA table directory. No conversion will be done if the character sets are not properly defined in this file. This is not usually a problem since this file defines several hundred character sets; most of the character sets in use today are defined in this file. See the description of the **imsimta chbuild** (UNIX and NT) utility for further information on the **charsets.txt** file.

If all the conditions are met, the MTA will proceed to build the character set mapping and do the conversion. The converted message part will be relabelled with the name of the character set to which it was converted.

The charset-conversion mapping has been extended to provide several additional capabilities:

- A **IN-CHARSET** option can be specified in the output template of a mapping entry. If present this overrides the charset specified in the encoded-word.
- A **RELABEL-ONLY** option that accepts an integer 0 or 1 can be specified. If this option has a value of 1 the **OUT-CHARSET** simply replaces the **IN-CHARSET**; no relabelling is done.
- If the **IN-CHARSET** option is used to set the input charset to * the charset will be "sniffed" to determine an appropriate label.

Converting ISO-8859-1 to UTF-8 and back

Stopped here

Suppose that ISO-8859-1 is used locally, but this needs to be converted to UTF-8 for use on the Internet. In particular, suppose the connection to the Internet is via the **tcp_local** and **tcp_internal** and **ims-ms** are where internal messages originate and are delivered. The **CHARSET-CONVERSION** table shown below brings such conversions about. Note that each **IN-CHAN** entries must be on a single line. The backslash (\) is used to signify this.

```
CHARSET-CONVERSION
```

```
IN-CHAN=tcp_internal;OUT-CHAN=tcp_local;CONVERT Yes
IN-CHAN=tcp_local;OUT-CHAN=tcp_internal;CONVERT Yes
IN-CHAN=tcp_local;OUT-CHAN=ims-ms;CONVERT Yes
IN-CHAN=*;OUT-CHAN=*;CONVERT No
IN-CHAN=tcp_internal;OUT-CHAN=tcp_local;IN-CHARSET=ISO-8859-1 OUT-CHARSET=UTF-8
IN-CHAN=tcp_local;OUT-CHAN=tcp_internal;IN-CHARSET=UTF-8 OUT-CHARSET=ISO-8859-1
IN-CHAN=tcp_local;OUT-CHAN=ims-ms;IN-CHARSET=UTF-8 OUT-CHARSET=ISO-8859-1
```

Converting EUC-JP to ISO-2022-JP and Back

The **CHARSET-CONVERSION** table shown below specifies a conversion between local usage of EUC-JP and the ISO 2022 based JP code.

```
CHARSET-CONVERSION

IN-CHAN=ims-ms;OUT-CHAN=ims-ms;CONVERT No
IN-CHAN=tcp_internal;OUT-CHAN=ims-ms;CONVERT No
IN-CHAN=tcp_internal;OUT-CHAN=tcp_internal;CONVERT No
IN-CHAN=tcp_internal;OUT-CHAN=*;CONVERT Yes
IN-CHAN=*;OUT-CHAN=ims-ms;CONVERT Yes
IN-CHAN=*;OUT-CHAN=tcp_internal;CONVERT Yes
IN-CHAN=tcp_internal;OUT-CHAN=*;IN-CHARSET=EUC-JP OUT-CHARSET=ISO-2022-JP
IN-CHAN=*;OUT-CHAN=ims-ms;IN-CHARSET=ISO-2022-JP OUT-CHARSET=EUC-JP
IN-CHAN=*;OUT-CHAN=tcp_internal;IN-CHARSET=ISO-2022-JP OUT-CHARSET=EUC-JP
```

Message Reformatting

As described above, the **CHARSET-CONVERSION** mapping table is also used to effect the conversion of attachments between MIME and several proprietary mail formats.

The following sections give examples of some of the other sorts of message reformatting which can be affected with the **CHARSET-CONVERSION** mapping table.

Non-MIME Binary Attachment Conversion

Mail in certain non-standard (non-MIME) formats; for example, mail in certain proprietary formats or mail from the Microsoft Mail (MSMAIL) SMTP gateway is automatically converted into MIME format if **CHARSET-CONVERSION** is enabled for any of the channels involved in handling the message. If you have a **tcp_local** channel then it is normally the incoming channel for messages from a Microsoft Mail SMTP gateway, and the following will enable the conversion of messages delivered to your local users:

```
CHARSET-CONVERSION

IN-CHAN=tcp_local;OUT-CHAN=ims-ms;CONVERT Yes
```

You may also wish to add entries for channels to other local mail systems. For instance, an entry for the **tcp_internal** channel:

```
CHARSET-CONVERSION

IN-CHAN=tcp_local;OUT-CHAN=l;CONVERT Yes
IN-CHAN=tcp_local;OUT-CHAN=tcp_internal;CONVERT Yes
```

Alternatively, to cover every channel you can simply specify **OUT-CHAN=*** instead of **OUT-CHAN=ims-ms**. However, this may bring about an increase in message processing overhead as all messages coming in the **tcp_local** channel will now be scrutinized instead of just those bound to specific channels.

More importantly, such indiscriminate conversions might place your system in the dubious and frowned upon position of converting messages--not necessarily your own site's--which are merely passing through your system, a situation in which you should merely be acting as a transport and not necessarily altering anything beyond the message envelope and related transport information.

To convert MIME into the format Microsoft Mail SMTP gateway understands, use a separate channel in your MTA configuration for the Microsoft Mail SMTP gateway; for example, `tcp_msmail`, and put the following in the mappings. file:

```
CHARSET-CONVERSION
```

```
IN-CHAN=*;OUT-CHAN=tcp_msmail;CONVERT RFC1154
```

Relabelling MIME Headers

Some user agents or gateways may emit messages with MIME headers that are less informative than they might be, but that nevertheless contain enough information to construct more precise MIME headers. Although the best solution is to properly configure such user agents or gateways, if they are not under your control, you can instead ask the MTA to try to reconstruct more useful MIME headers.

If the first probe of the **CHARSET-CONVERSION** mapping table yields a **Yes** or **Always** option, then the MTA will check for the presence of conversion entries. If a conversions entries exist, then the MTA will look in them for an entry with **RELABEL=1** and if it finds such an entry, the MTA will then perform any MIME relabelling specified in the entry. See "[To Control Conversion Processing](#)" for information on conversions entries.

For example, the combination of a **CHARSET-CONVERSION** table such as:

```
CHARSET-CONVERSION
```

```
IN-CHAN=tcp_local;OUT-CHAN=tcp_internal;CONVERT Yes
```

and MTA conversions entries of

```
out-chan=ims-ms; in-type=application; in-subtype=octet-stream;
in-parameter-name-0=name; in-parameter-value-0=*.ps;
out-type=application; out-subtype=postscript;
parameter-copy-0=*; relabel=1
```

```
out-chan=ims-ms; in-type=application; in-subtype=octet-stream;
in-parameter-name-0=name; in-parameter-value-0=*.msw;
out-type=application; out-subtype=msword;
parameter-copy-0=* relabel=1
```

will result in messages that arrive on the `tcp_local` channel and are routed to the `ims-ms` channel, and that arrive originally with MIME labelling of application/octet-stream but have a filename parameter with the extension `ps` or `msw`, being relabelled as application/postscript or application/msword, respectively. (Note that this more precise labelling is what the original user agent or gateway should have performed itself.) Such a relabelling can be particularly useful in conjunction with a **MIME-CONTENT-TYPES-TO-MR** mapping table, used to convert such resulting MIME types back into appropriate **MRTYPE** tags, which needs precise MIME labelling in order to function optimally; if all content types were left labelled only as application/octet-stream, the **MIME-CONTENT-TYPES-TO-MR** mapping table could only, at best, unconditionally convert all such to one sort of **MRTYPE**.

With the above example and **MIME-CONTENT-TYPES-TO-MR** mapping table entries including

```
APPLICATION/POSTSCRIPT PS
```

```
APPLICATION/MSWORD MW
```

a labelling coming in as, for example,

Content-type: application/octet-stream; name=stuff.ps

would be relabelled as

Content-type: application/postscript

and then converted into an **MRTYPE** tag **PS** to let Message Router know to expect PostScript.

Sometimes it is useful to do relabelling in the opposite sort of direction, "downgrading" specific MIME attachment labelling to application/octet-stream, the label for generic binary data. In particular, "downgrading" specific MIME labelling is often used in conjunction with the **convert_octet_stream** channel option on the **mime_to_x400** channel (PMDf-X400) or **xapi_local** channel (PMDf-MB400) to force all binary MIME attachments to be converted to X.400 bodypart 14 format.

For instance, the combination of a **CHARSET-CONVERSION** mapping table such as

```
CHARSET-CONVERSION
```

```
IN-CHAN=*;OUT-CHAN=mime_to_x400*;CONVERT Yes
```

and conversions entries of

```
out-chan=mime_to_x400*; in-type=application; in-subtype=*;  
out-type=application; out-subtype=octet-stream; relabel=1
```

```
out-chan=mime_to_x400*; in-type=audio; in-subtype=*;  
out-type=application; out-subtype=octet-stream; relabel=1
```

```
out-chan=mime_to_x400*; in-type=image; in-subtype=*;  
out-type=application; out-subtype=octet-stream; relabel=1
```

```
out-chan=mime_to_x400*; in-type=video; in-subtype=*;  
out-type=application; out-subtype=octet-stream; relabel=1
```

will result in downgrading various specific MIME attachment labelling to the generic application/octet-stream labelling (so that **convert_octet_stream** will apply) for all messages going to **mime_to_x400*** channels.

MacMIME Format Conversions

Macintosh files have two parts, a resource fork that contains Macintosh specific information, and a data fork that contains data usable on other platforms. This introduces an additional complexity when transporting Macintosh files, as there are four different formats in common use for transporting the Macintosh file parts. Three of the formats, Applesingle, Binhex, and Macbinary, consist of the Macintosh resource fork and Macintosh data fork encoded together in one piece. The fourth format, Appledouble, is a multipart format with the resource fork and data fork in separate parts. Appledouble is hence the format most likely to be useful on non-Macintosh platforms, as in this case the resource fork part may be ignored and the data fork part is available for use by non-Macintosh applications. But the other formats may be useful when sending specifically to Macintoshes.

The MTA can convert between these various Macintosh formats. The **CHARSET-CONVERSION** options **Appledouble**, **Applesingle**, **Binhex**, or **Macbinary** tell the MTA to convert other MacMIME structured parts to a MIME structure of multipart/appledouble, application/applefile, application/mac-binhex40, or application/macbinary, respectively. Further, the **Binhex** or **Macbinary** options also request conversion to the specified format of non-MacMIME format parts that do nevertheless contain **X-MAC-TYPE** and **X-MAC-CREATOR** parameters on the MIME Content-type: header. The **CHARSET-CONVERSION** option **Block** tells the MTA to

extract just the data fork from MacMIME format parts, discarding the resource fork; (since this loses information, use of **Appledouble** instead is generally preferable).

For example, the following **CHARSET-CONVERSION** table would tell the MTA to convert to Appledouble format when delivering to the VMS MAIL mailbox or a GroupWise postoffice, and to convert to Macbinary format when delivering to the Message Router channel:

```
CHARSET-CONVERSION

IN-CHAN=*;OUT-CHAN=l;CONVERT Appledouble
IN-CHAN=*;OUT-CHAN=wpo_local;CONVERT Appledouble
IN-CHAN=*;OUT-CHAN=tcp_internal;CONVERT Macbinary
```

```
{}
```

The conversion to Appledouble format would only be applied to parts already in one of the MacMIME formats. The conversion to Macbinary format would only be applied to parts already in one of the MacMIME formats, or non-MacMIME parts which included X-MAC-TYPE and X-MAC-CREATOR parameters on the MIME Content-type: header.

When doing conversion to Appledouble or Block format, the **MAC-TO-MIME-CONTENT-TYPES** mapping table may be used to indicate what specific MIME label to put on the data fork of the Appledouble part, or the Block part, depending on what the Macintosh creator and Macintosh type information in the original Macintosh file were. Probes for this table have the form *format | type | creator | filename* where format is one of SINGLE, BINHEX or MACBINARY, where type and creator are the Macintosh type and Macintosh creator information in hex, respectively, and where filename is the filename.

For example, to convert to Appledouble when sending to the **ims-ms** channel and when doing so to use specific MIME labels for any MS Word or PostScript documents converted from MACBINARY or BINHEX parts, appropriate tables might be:

```
CHARSET-CONVERSION

IN-CHAN=*;OUT-CHAN=ims-ms;CONVERT Appledouble

MAC-TO-MIME-CONTENT-TYPES

! PostScript
MACBINARY|45505346|76677264|* APPLICATION/POSTSCRIPT$Y
BINHEX|45505346|76677264|* APPLICATION/POSTSCRIPT$Y
! Microsoft Word
MACBINARY|5744424E|4D535744|* APPLICATION/MSWORD$Y
BINHEX|5744424E|4D535744|* APPLICATION/MSWORD$Y
```

Note that the template (right hand side) of the mapping entry must have the \$Y flag set in order for the specified labelling to be performed. Sample entries for additional types of attachments may be found in the file **mac_mappings.sample** in the MTA table directory.

If you wish to convert non-MacMIME format parts to Binhex or Macbinary format, such parts need to have X-MAC-TYPE and X-MAC-CREATOR MIME Content-type: parameter values provided. Note that MIME relabelling can be used to force such parameters onto parts that would not otherwise have them.

Service Conversions

The MTA's conversion service facility may be used to process with site-supplied procedures a message so as to produce a new form of the message. Unlike either the sorts of **CHARSET-CONVERSION** operations discussed above or the **conversion** channel, which operate on the content of individual MIME message parts, conversion services operate on entire MIME message parts (MIME headers and content) as well as entire MIME messages. Also, unlike other **CHARSET-CONVERSION** operations or conversion channel operations, conversion services are expected to do their own MIME disassembly, decoding, re-encoding, and reassembly.

Like other **CHARSET-CONVERSION** operations, conversion services are enabled through the **CHARSET-CONVERSION** mapping table. If the first probe of the **CHARSET-CONVERSION** mapping table yields a **Yes** or **Always** option, then the MTA will check for the presence of conversions entries. If conversions entries exist, then the MTA will look in them for an entry specifying a **SERVICE-COMMAND**, and if it finds such an entry, execute it. The conversions entries should be created using **msconfig edit conversions** and have the form:

```
in-chan=channel-pattern;  
in-type=type-pattern; in-subtype=subtype-pattern;  
service-command=command
```

Of key interest is the command string. This is the command that should be executed to perform a service conversion (for example, invoke a document converter). The command must process an input file containing the message text to be serviced and produce as output a file containing the new message text. On UNIX, the command must exit with a 0 if successful and a non-zero value otherwise.

For instance, the combination of a **CHARSET-CONVERSION** table such as

```
CHARSET-CONVERSION
```

```
IN-CHAN=bsout_;OUT-CHAN=;CONVERT Yes
```

and an MTA conversions entry on UNIX of

```
in-chan=bsout_*; in-type=*; in-subtype=*;  
service-command="/pmdf/bin/compress.sh compress $INPUT_FILE $OUTPUT_FILE"
```

will result in all messages coming from a **BSOUT** channel being compressed.

Environment variables are used to pass the names of the input and output files as well as the name of a file containing the list of the message's envelope recipient addresses. The names of these environment variables are:

- **INPUT_FILE** - Name of the input file to process
- **OUTPUT_FILE** - Name of the output file to produce
- **INFO_FILE** - Name of the file containing envelope recipient addresses

The values of these three environment variables may be substituted into the command line by using standard command line substitution: that is, preceding the variable's name with a dollar character on UNIX. For example, when **INPUT_FILE** and **OUTPUT_FILE** have the values **a.in** and **a.out**, then the following declaration on UNIX:

```
in-chan=bsout_*; in-type=*; in-subtype=*;  
service-command="/pmdf/bin/convert.sh $INPUT_FILE $OUTPUT_FILE"
```

executes the command

`/pmdf/bin/convert.sh a.in a.out`

LMTP Delivery

This chapter provides an overview of the Local Mail Transfer Protocol (LMTP), and describes how to configure both the LMTP client and server.

Overview of LMTP

The Oracle Communications Messaging Server MTA can use Local Mail Transfer Protocol (LMTP), as defined in RFC 2033, to deliver messages to the message store in a multi-tiered Messaging Server deployment. In this scenario, the front-end relays become responsible for address expansion and delivery methods such as autoreply and forwarding, and also for mailing list expansion. Delivery to the back-end stores historically has been over SMTP, which requires the back-end system to look up the recipient addresses in the LDAP directory again, thereby engaging the full machinery of the MTA. For speed and efficiency, the MTA can use LMTP rather than SMTP to deliver messages to the back-end store. The Messaging Server's LMTP server is not intended as a general purpose LMTP server, but rather as a private protocol between the relays and the back-end message stores. For simplicity of discussion, examples involving two-tiered deployments are used.

Note: LMTP is recommended for use in multi-tiered deployments. Also, the Messaging Server's LMTP service as implemented is not designed to work with third-party LMTP servers or third-party LMTP clients.

LMTP Delivery Features

The MTA's LMTP server is more efficient for delivering to the back-end message store because it:

- Reduces the load on the back-end stores. Because relays are horizontally scalable and back-end stores are not, it is good practice to push as much processing to the relays as possible.
- Reduces the load on the LDAP servers. The LDAP infrastructure is often a limiting factor in large messaging deployments.
- Reduces the number of message queues. Having queues on both the relay and the back-end store makes finding a lost message that much harder for administering a messaging deployment.

Support for LMTP Client and Server to Detect and Respond to Certain Conditions

Support has been added to the LMTP client and server to detect and respond to the condition where a given host's LMTP server is responding but isn't associated with the master store replica. When this happens the LMTP server produces a banner or **MAIL FROM** response of the form:

```
423 4.3.2 Host not master for store; correct master host is <HOSTNAME>
```

or, if the correct master host is not known:

```
421 4.3.2 Host not master for store; cannot determine correct master host
```

When the LMTP client sees the 423 banner it will immediately disconnect and reconnect to the correct host. The affinity subsystem is also informed of the new master host.

The 421 banner is treated like any other 4YZ banner response; the client disconnects and tries the next host in the affinity group.

A 423 **MAIL FROM** response now engages the LMTP client's fast retry logic; the client will abort the delivery attempt and disconnect but schedule the message for fast retry.

The handling of a 421 **MAIL FROM** response is unchanged.


Two new LMTP channel-specific parameters have been added to support this new capability:

- **WRONG_MASTER_FAST_RETRY** - Controls the actual value used to override the normal backoff value. The default is 1, meaning to ask the Job Controller for a fairly quick retry. Setting the option to 0 disables backoff override in this case. Positive values greater than 1 result in a retry, but not quite as quick of one. The formula used is $n + \text{rand}() \% (15 * n)$ seconds, where n is the **WRONG_MASTER_FAST_RETRY** value.
- **WRONG_MASTER_FALLBACK_ATTEMPTS** - Controls the number of times the LMTP client will honor a 423 banner redirect. The default is 1, meaning a single redirect will be honored for a given message.

Messaging Processing in a Two-Tiered Deployment Without LMTP

Figure 11-1 shows message processing in a two-tiered deployment without LMTP.

Figure 11-1 Two-Tiered Deployment Without LMTP



Without LMTP, in a two-tiered deployment with relays "in front" of the store systems, inbound message processing begins with a connection on the SMTP port picked up by the dispatcher on the relay machine and handed off to a **tcp_smtp_server** process. This process does a number of things with the inbound message including:

- Looking up the user in the directory
- Determining if the user is within a domain hosted by this email deployment
- Determining if the user is a valid user in the domain
- Rewriting the envelope address as **@mailhost:user@domain**
- Enqueuing the message for delivery to the mailhost

The **smtp_client** process then picks up the mail message from the queue and sends it to the mailhost. On the mailhost, some very similar processing takes place. A connection on the SMTP port is picked up by the dispatcher and handed off to a **tcp_smtp_server** process. This process does a number of things to the message, including:


- Looking up the user in the directory
- Determining if the user is within a domain hosted by this email deployment
- Determining if the user is a valid user in the domain
- Rewriting the envelope address to direct the message to the **ims_ms** channel
- Enqueuing the message for delivery to the store

Then the **ims_ms** process picks up the mail message and attempts to deliver it to the store. In this scenario, the enqueuing processing is performed twice, and the MTAs each perform an LDAP lookup.

Messaging Processing in a Two-Tiered Deployment With LMTP

Figure 11–2 shows message processing in a two-tiered deployment scenario with LMTP.

Figure 11–2 Two-Tiered Deployment With LMTP



With LMTP in place, a connection on the SMTP port of the relay machine is picked up by the dispatcher and handed off to a **tcp_smtp_server** process. This process does a number of things with the inbound message including:

- Looking up the user in the directory

- Determining if the user is within a domain hosted by this email deployment
- Determining if the user is a valid user in the domain
- Determining which back end message store machine hosts the mailbox for the user
- Enqueuing the message for delivery to the mailhost

On the store machine, a connection to the LMTP port is received by the dispatcher and handed off to the **lmtp_server** process. The LMTP server then inserts the message into the user's mailbox or into the UNIX native mailbox. If message delivery is successful, the message is dequeued on the relay machine. If unsuccessful, the message remains on the relay machine. Note that the LMTP process on the message store does not engage any MTA machinery for processing addresses or messages.

LMTP Architecture

An LMTP configuration consists of setting up LMTP channels, one for the LMTP client and one for the LMTP server. LMTP channels are special cases of TCP/IP channels. The general SMTP-over-TCP/IP channel is often configured for bidirectional use. An LMTP channel, on the other hand, is configured to be dedicated for either client or server use. You configure the LMTP client channel on the MTA front-end relay host and the LMTP server channel on the back-end message store. Running **commpkg install** on both the front-end relay and back-end store system will configure LMTP properly.

For the most part, the MTA itself can be basically absent from the back-end server. The following items are the only necessary MTA components on the back end:

- The dispatcher
- The LMTP server
- A simple MTA configuration including **mappings**

The dispatcher must run on the back-end server so that it can start the LMTP servers that run under it. Because the dispatcher and the LMTP server use various functions of **libimta**, this needs to be present on the back-end server as well.

The LMTP server does not perform any of the usual MTA enqueueing or dequeuing functions, header processing, or address translations. The front-end relay system performs all the manipulation of the content of the messages and addresses, which then presents to the LMTP server the message in exactly the form to be delivered to the message store and with the delivery address already in the form required by the store. Additional recipient information that is usually available as a message that is delivered to the store, such as the user's quota, is presented along with the recipient address as LMTP options. Should a delivery attempt fail, the message is left enqueued in the LMTP queue on the relay system.

Configuring LMTP

Setting up LMTP requires that you configure both the front-end relay hosts and back-end message store hosts. On the relays, you must change the **mta.delivery_options** option so that messages being delivered to the stores are passed to the LMTP channel. The back-end store must be configured with the dispatcher, but does not need the job controller. On the LMTP client, you must configure the job controller to run the LMTP client.

This section describes how to configure the LMTP front ends and back ends by using Unified Configuration recipes. You could also perform the same configuration

manually by using the **msconfig** command, however, using recipes is faster and makes the task repeatable. You run one recipe on the LMTP front ends and one on the LMTP back ends. For more information on Unified Configuration recipes, see the discussion on recipe language in the *Messaging Server Reference* and "Using Recipes."

Topics in this section:

- [Before You Begin](#)
- [To Configure the Front-end MTA Relay with LMTP](#)
- [To Configure Back-End Stores with LMTP and a Minimal MTA](#)

Before You Begin

- To see if any LMTP options are already enabled on your Messaging Server hosts, run the following command:

```
msconfig show mta.delivery_options -default
```

To Configure the Front-end MTA Relay with LMTP

Use the following "Example Recipe to Configure LMTP Front-end Relay with LMTP" to configure the inbound MTA relay for LMTP.

1. Make a copy of the example recipe file and save it as *recipe.rcp* in the **config/recipes** directory.
2. In this example, replace the following items with your site-specific values:
 - **myIP**
 - **myNetmaskbits**
3. To run the recipe, type the following command:

```
cd MessagingServer_home/bin
msconfig run MessagingServer_home/config/recipes/<recipe_name>
```

4. Recompile if running a compiled configuration

```
imsimta cnbuild
```

5. Restart Messaging Server.

```
cd MessagingServer_home/bin
stop-msg
start-msg
```

Example Recipe to Configure LMTP Front-end Relay with LMTP

```
# -*- mode: sieve; -*-
description("frontendMTA to backend LMTP store");
keywords(["frontend", "MTA", "LMTP"]);
#
# Sample recipe for a frontend MMP/MTA to a backend LMTP store
# the corresponding recipe for the backend store/LMTP is backendLMTP.rcp
#
#####
# !!!!!!!!!!!!!!!!!!!!!!!
# CHANGE THESE
# !!!!!!!!!!!!!!!!!!!!!!!
```

```

#
# constants - supplied input
#
# my IP address
myIP = "10.133.158.10";
# network portion of IP address in number of bits
myNetmaskbits = "8";

#####
#
# configure LMTP frontend
#

#
# disable store
#
set_option("store.enable", "0");
set_option("imap.enable", "0");
set_option("pop.enable", "0");

#
# add to rewrite rules
# .lmtpl $E$F$U%H.lmtpl@lmtplcs-daemon
# .lmtpl $B$F$U%H@$H@lmtplcs-daemon
#
# should really check to see if the rewrite rule exists instead
# of unilaterally appending it
#
append_rewrites([".lmtpl", "$E$F$U%H.lmtpl@lmtplcs-daemon"]);
append_rewrites([".lmtpl", "$B$F$U%H@$H@lmtplcs-daemon"]);

#
# add channel tcp_lmtplcs
# ,----
# | tcp_lmtplcs defragment lmtpl multigate connectcanonical fileinto @$40:$U+$S@$D
f\
# | lagtransfer multigate connectcanonical port 225 nomx sin\
# | gle_sys pool SMTP_POOL dequeue_removeoute
# | lmtplcs-daemon
# `----
#
set_option("channel:tcp_lmtplcs.connectcanonical");
set_option("channel:tcp_lmtplcs.defragment");
set_option("channel:tcp_lmtplcs.fileinto", "@$40:$U+$S@$D");
set_option("channel:tcp_lmtplcs.flagtransfer");
set_option("channel:tcp_lmtplcs.lmtpl");
set_option("channel:tcp_lmtplcs.multigate");
set_option("channel:tcp_lmtplcs.nomx");
set_option("channel:tcp_lmtplcs.pool", "SMTP_POOL");
set_option("channel:tcp_lmtplcs.port", "225");
set_option("channel:tcp_lmtplcs.single_sys");
set_option("channel:tcp_lmtplcs.dequeueremoveoute");
set_option("channel:tcp_lmtplcs.official_host_name", "lmtplcs-daemon");

#
# ! The modified DELIVERY_OPTIONS which activate LMTP
# ! delivery from a frontend relay to the backend
#
# DELIVERY_OPTIONS=\
#      #*mailbox=@$X.LMTP:$M%$\'$2I$\'_+$2S@lmtplcs-daemon,\

```

```

#      #&members=*, \
#      #*native=@$X.LMTPN:$M+$2S@native-daemon, \
#      #*unix=@$X.LMTPN:$M, \
#      #*file=@$X.LMTPN:+$F, \
#      #&@members_offline=*, \
#      #/hold=@hold-daemon:$A, \
#      #program=$M$P@pipe-daemon, \
#      #forward=**, \
#      #*^!autoreply=$M+$D@bitbucket
#
# NOTE NOTE NOTE - have to escape the backslash in the "mailbox=..."
set_option("mta.delivery_options", "#*mailbox=@$X.LMTP:$M%\$2I$_
+$2S@lmtpcs-daemon, #&members=*, #*native=@$X.LMTPN:$M+$2S@native-daemon, #*unix=@$X.
LMTPN:$M, #*file=@$X.LMTPN:+$F, #&@members_
offline=*, #/hold=@hold-daemon:$A, #program=$M$P@pipe-daemon, #forward=**, #*^!autore
ply=$M+$D@bitbucket");

#####
write("use \"write -remark frontendLMTP.rcp\" to write out the changes\n");

```

This recipe performs the following configuration on the LMTP front end:

1. Disables the message store, as it is not needed.
2. Enables the necessary rewrite rules. The recipe language **append_rewrites** function adds entries to existing rewrite rules, or, if none exists, adds them as new rules.
3. Adds the **tcp_lmtpcs** channel with the appropriate options:
 - **connectcanonical**: Tells the MTA to compare the recipient envelope address domain with the channel host proper names, and if the domain name matches one of the channel's host proper names, then connect to the host name corresponding to that host proper name.
 - **defragment**: Any message/partial messages queued to the channel are placed in the **defragmentation** channel queue instead. Once all the parts have arrived, the message is rebuilt and sent on its way.
 - **fileinto**, value of **@\$4O:\$SU+\$S@\$D**: Specifies how to alter an address when a Sieve filter "fileinto" action is applied. In \$4O, the O is the capital or majuscule letter "o", not the numeral zero 0. The effect is that the explicit source route to the mailhost should be preserved if present, and the foldername should be inserted as a subaddress into the original address, replacing any originally present subaddress.
 - **flagtransfer**: Enables SMTP client support of the XDFLG private SMTP extension command.
 - **lmtp**: Specifies that channel supports LMTP protocol.
 - **multigate**: Instructs the MTA to route the message to the daemon mailbox specified by the daemon channel option on the system specified in the message's To: address.
 - **nomx**: Disables MX lookups.
 - **pool**, value of **SMTP_POOL**: Specifies the **SMTP_POOL** pool where the jobs are created for this channel.
 - **port**, value of **225**: Specifies dispatcher port number.

- **single_sys**: Creates a single copy of the message for each destination system (more precisely, each destination domain name) associated with a recipient address.
 - **dequeueremoveoute**: Causes source routes to be stripped from envelope recipient addresses when the channel dequeues messages (but after the channel has determined how to route the message).
 - **official_host_name**, value of **lmtpcs-daemon**: Specifies the "name" of the system with which this channel communicates. (In legacy configuration, the official host name is specified as the first name on the second line of a channel definition.)
4. Finally, the delivery options are set, which activates LMTP delivery from the front-end relay to the back-end message store. In the ***mailbox=@\$X.LMTP:\$M%\$\\\$2I\$_+\$2S@lmtpcs-daemon**, portion of the delivery options, the script "escapes" the forward slash, so that the actual entry in the recipe is as follows:

```
... "#*mailbox=@$X.LMTP:$M%$\\$2I$_+$2S@lmtpcs-daemon, #&members=*, ...
```

Note: You need to add the full **delivery_options** as shown in the example. You cannot override just the default for one option. If you specify **delivery_options** at all, you must define them all.

To Configure Back-End Stores with LMTP and a Minimal MTA

Use the following ["Example Recipe to Configure LMTP Back-end Store with LMTP"](#) to configure the back-end store for LMTP. This recipe:

- Prompts for **myIP** and **feIPs**, if you do not specify them in the recipe (but does not validate the IP addresses)
- Checks if the **tcp_lmtpps** channel exists before creating it
- Checks if the dispatcher group exists before creating it
- Checks if **PORT_ACCESS** entries exist before creating them
- Treats **feIPs** as a list to allow multiple front ends
- Can be run multiple times correctly

To create and run the recipe, follow these steps:

1. Make a copy of the example recipe file and save it as *recipe.rcp* in the **config/recipes** directory.
2. In this example, replace the following items with your site-specific values:
 - **myIP**
 - **feIPs**
3. To run the recipe, type the following command:

```
cd MessagingServer_home/bin
msconfig run MessagingServer_home/config/recipes/<recipe_name>
```

4. Recompile if running a compiled configuration.

```
imsimta cnbuild
```

5. Restart Messaging Server.

```
cd MessagingServer_home/bin
stop-msg
start-msg
```

Example Recipe to Configure LMTP Back-end Store with LMTP

```
# -*- mode: sieve; -*-
description("backend store via LMTP");
keywords(["backend", "store", "LMTP"]);
#
# Sample recipe for a backend store via LMTP
# the corresponding recipe for the frontend MMP/MTA is frontendLMTP.rcp
#####
# !!!!!!!!!!!!!!!!!!!!!!!
# CHANGE THESE
# !!!!!!!!!!!!!!!!!!!!!!!
#
# constants - supplied input
# if you leave it blank, the script will prompt for it
# sample entry
#myIP = "10.133.152.193";
myIP = "";

#
# list of frontend machines that access this store via LMTP
# if you leave it blank, the script will prompt for it
# sample entry:
#feIPs = ["10.133.152.192", "10.133.152.193"];
feIPs = [];

#####
# prompt for myIP and feIP if needed

if (length(myIP) <= 0) {
    myIP = read("Enter the IP address of this host: ");
}

if (length(feIPs) <= 0) {
    loop {
        ip = read("Enter the IP address of a frontend machine (<RET> if no more): ");
        exitif (ip == "");
        push(feIPs, ip);
    }
}

#####
#
# configure LMTP backend

#
# create tcp_lmtpss channel. In legacy config, this would show up as:
# tcp_lmtpss lmtpl flagtransfer identnonnumeric
# tcp_lmtpss-daemon
if exists_channel("tcp_lmtpss") {
    warn("-- WARNING: tcp_lmtpss channel already exists.");
} else {
    # This creates the channel by using individual options
```

```

#set_option("channel:tcp_lmtpss.flagtransfer");
#set_option("channel:tcp_lmtpss.identnonnumeric");
#set_option("channel:tcp_lmtpss.lmtp");
#set_option("channel:tcp_lmtpss.official_host_name", "tcp_lmtpss-daemon");
# an alternative way of doing it as a one-liner
print("-- INFO: Adding tcp_lmtpss channel\n");
add_channel("tcp_lmtpss",
            ["flagtransfer", "",
             "identnonnumeric", "",
             "lmtp", "",
             "official_host_name", "tcp_lmtpss-daemon"]);
}

#
# dispatcher.cnf
# uncomment [SERVICE=LMPSS] block
#
if exists_group("dispatcher.service:LMPSS") {
    warn("-- WARNING: dispatcher.service:LMPSS group already exists");
} else {
    print("-- INFO: Creating dispatcher.service:LMPSS group\n");
    # This creates the group by using individual options
    #set_option("dispatcher.service:LMPSS.image", "IMTA_BIN:tcp_lmtp_server");
    #set_option("dispatcher.service:LMPSS.logfilename", "IMTA_LOG:tcp_lmtpss_
server.log");
    #set_option("dispatcher.service:LMPSS.option", "CHANNEL=tcp_lmtpss");
    #set_option("dispatcher.service:LMPSS.tcp_ports", "225");
    #set_option("dispatcher.service:LMPSS.stacksize", "2048000");
    #set_option("dispatcher.service:LMPSS.enable", "1");

    # alternate way of doing this in one line
    add_group("dispatcher.service:LMPSS",
             ["image", "IMTA_BIN:tcp_lmtp_server",
              "logfilename", "IMTA_LOG:tcp_lmtpss_server.log",
              "option", "CHANNEL=tcp_lmtpss",
              "tcp_ports", "225",
              "stacksize", "2048000",
              "enable", "1"]);
}

#
# add PORT_ACCESS mapping entries
#
# allow frontends (feIPs) to access LMTP port
# TCP|*|225|10.133.152.192|* $Y
# TCP|*|226|10.133.152.192|* $Y
#
# ! Allow 'msprobe' on this host (myIP) to connect to the LMTP ports
# !
# TCP|*|225|10.133.152.193|* $Y
# TCP|*|226|10.133.152.193|* $Y

portAccess_optlist = get_mapping("PORT_ACCESS");
#print ("\n -- DEBUG optlist for PORT_ACCESS" . portAccess_optlist . "\n");

# list of IP addresses
ipaddrs = [feIPs];
push(ipaddrs, myIP);
numips = length(ipaddrs);
ports = ["225", "226"];

```

```

numports = length(ports);

i = 1;
loop {                                #loop over all ipaddrs
    p = numports;
    loop {
        tmp = "TCP|*|" + ports[p] + "|" + ipaddrs[i] + "|*";
        if exists_optlist(get_mapping("PORT_ACCESS"), tmp) {
            warn ( "-- WARNING: optlist for " . tmp . " exists in PORT_ACCESS");
        } else {
            print("-- INFO: Adding IP " + ipaddrs[i] + " port " + ports[p] + " to PORT_
ACCESS mapping\n");
            prepend_mapping("PORT_ACCESS", [tmp, "$Y"]);
        }
        exitif(p==1);
        p--;
    }
    exitif(i==numips);
    i++;
}

#####
write("use \"write -remark backendLMTP.rcp\" to write out the changes\n");

```

LMTP Protocol as Implemented

This section provides a sample LMTP dialogue with an explanation of what is seen in that dialogue. The LMTP client on the relay uses standard LMTP protocol to talk to the LMTP server on the back end store. However the protocol is used in specific ways. For example:

```

---> LHLO
<--- 250 OK

```

No action is taken on the **LHLO** message. The reply is always **250 OK**.

```

---> MAIL FROM: address size=messageSizeInBytes
<--- 250 OK

```

No checks or conversions are made on the originator address. The **size=** option gives a size in bytes for the message that is to be delivered. This is the size of the message exactly as it appears in the protocol. It is not necessarily the exact size of the message, but the actual message size will not exceed this size. The LMTP server allocates a memory buffer of this size to receive the message.

```

---> RCPT TO: uid+folder@domain xquota=size,number xdfld=xxx
<--- 250 OK

```

No checks are made on the recipient addresses at the time they are received, but a list of recipients is built for later use. Note that the **@domain** part of the address is omitted for **uids** in the primary domain, and that the **+folder** part is optional. This is the same address format used by the message store channel in the MTA.

The **xquota=** option gives the user's message quotas which consist of the maximum total size and the maximum number of messages. The MTA provides this information which it retrieves while performing an LDAP lookup on the user to do the address translation. This information is used to keep the quota information in the message store synchronized with the directory. Getting the quota information does not result in an additional performance hit.

The `xdfld=` option specifies a number which is interpreted as a bit field. These bits control how the message is delivered. For example, the bit whose value is 2, if set, guarantees delivery of the message even if the user is over quota. (Note that `xdfld` is an internal option and the bits in it are subject to change or addition without notice. Oracle does not support other clients using this extension with Messaging Server, nor does Oracle support using the Messaging Server LMTP client with some other server and this option.)

This interaction may be repeated many times, once for each recipient.

```
--->DATA
---> <the message text>
--->.
```

The LMTP client then sends the entire message, dot-stuffed, just as SMTP does. The message finishes with a dot (.) alone on a line. If the message size is exceeded the LMTP server sends:

```
<--- 500 message too big
```

and ends the connection.

Assuming that the message is received correctly, the LMTP server then sends back to the LMTP client the status for each recipient given in the **RCPT TO:** lines. For instance, if the message is delivered successfully, the response is:

```
<--- 250 2.5.0 address OK
```

where **address** is exactly as it appeared on the **RCPT TO:** line.

The conversation can either repeat with another **MAIL FROM:** line or end with the following interaction:

```
---> quit
<--- 221 OK
```

Table 11-1 shows the possible status codes for each recipient. This three-column table shows the short code in the first column, its long-code equivalent in the second column and the status text in the third column. 2.x.x status codes are success codes, 4.x.x codes are retryable errors, and 5.x.x codes are non-retryable errors.

Table 11-1 LMTP Status Codes for Recipients

Short Code	Long Code	Status Text
250	2.5.0	OK
420	4.2.0	Mailbox Locked
422	4.2.2	Quota Exceeded
420	4.2.0	Mailbox Bad Formats
420	4.2.0	Mailbox not supported
430	4.3.0	IMAP IOERROR
522	5.2.2	Persistent Quota Exceeded
523	5.2.3	Message too large
511	5.1.1	mailbox nonexistent
560	5.6.0	message contains null
560	5.6.0	message contains nl
560	5.6.0	message has bad header

Table 11-1 (Cont.) LMTP Status Codes for Recipients

Short Code	Long Code	Status Text
560	5.6.0	message has no blank line

Otherwise, there are changes to the delivery options for mailbox, native (and, therefore, UNIX), and file. The object of these rules is to generate addresses that will cause the messages to be sent through the appropriate LMTP channel to the back end servers. The addresses generated are source routed addresses of the form:

```
@sourceroute:_localpart_@_domain_
```

Vacation Automatic Message Reply

For automatically generated responses to email (autoreply), specifically vacation messages, the MTA uses Message Disposition Notifications (*MDNs*) and the Sieve scripting language. MDNs are email messages sent by the MTA to a sender and/or postmaster reporting on a message's delivery disposition. MDNs are also known as read receipts, acknowledgments, receipt notifications, or delivery receipts. The Sieve is a simple scripting language used to create mail filters.

This information describes the vacation autoreply mechanism. In most cases, you do not need to modify the default configuration. However, you might want to configure your system such that an MTA relay performs vacation processing rather than the back-end message stores.

Topics:

- [Vacation Autoreply Overview](#)
- [Configuring Autoreply](#)
- [Vacation Autoreply Theory of Operation](#)
- [Vacation Autoreply Attributes](#)
- [Other Auto Reply Tasks and Issues](#)

Vacation Autoreply Overview

Vacation Sieve scripts are generated automatically from the various LDAP vacation attributes (see "[Vacation Autoreply Attributes](#)"). They can also be specified explicitly for additional flexibility. The underlying mechanism for tracking vacations is a set of files, one per intended recipient, that keep track of when replies were sent to the various senders.

By default, the MTA evaluates vacation on the back-end store systems. However, because MTA relays do not do as much work as back-end stores, for performance reasons, you can have the MTA evaluate vacation on the mail relay machines instead of on the back-end store. Use of this feature, however, can result in vacation responses being sent out more often than intended because different relays handle different messages. If you do not want vacation messages to be sent out more often than you intend, you may share the tracking of files between the relays. If this is also unacceptable to you, you can always have vacation evaluated on the back-end store systems.

Configuring Autoreply

Delivery addresses are generated through a set of patterns. The patterns used depend upon the values defined for the **mailDeliveryOption** attribute. One delivery address is generated for each valid **mailDeliveryOption** that is set in the recipient's LDAP entry. The patterns are defined by the setting of the **delivery_options** MTA option, which consists of a comma-separated list of *option=pattern* pairs. The default autoreply *option=pattern* pair is:

```
^*!autoreply=$M+$D@bitbucket
```

Table 12–1 shows the prefix characters used for the autoreply rule in the first column and their definitions in the second column.

Table 12–1 Prefix Characters for the Autoreply Rule in *delivery_options*

Prefix Character	Definition
!	Enables the generation of the autoreply Sieve script.
#	Allows the processing to take place on relays.
^	Option is only evaluated if the vacation dates indicate that it should be evaluated.
*	Option is only evaluated for users, not groups.
@	Extracts preferred language information from various message header fields as well as from LDAP entries associated with envelope From: addresses. For this information to be available at the correct time, the message must pass through the reprocess channel when autoreply is engaged. This is done by adding the @ flag to the autoreply delivery option. The addition of a channel hop increases message processing overhead.

The autoreply rule itself specifies an address destined for the bitbucket channel. The mail is considered delivered by this method once the autoreply is generated, but the MTA machinery requires a delivery address. Anything delivered to the bitbucket channel is discarded.

To Configure Autoreply on the Back-end Store System

The default autoreply rule in **delivery_options** causes the autoreply to take place on the mail server that serves the user. If you want vacation messages to be evaluated on the back-end store system, you do not have to configure anything. This is the default behavior.

To Configure Autoreply on a Relay

If you want to evaluate vacation on the relay rather than on the back-end store system to enhance performance, add the # prefix to the autoreply pattern. This can be done by performing the following steps:

1. Use the **msconfig** command to determine the current setting of **delivery_options**:

```
msconfig
msconfig> show delivery_options
```

2. If the **delivery_options** MTA option is not set, determine the current default by using the **show -default** command and then cutting and pasting to set the option to that default:

To Share Autoreply Information Between Relays

For information about sharing autoreply information between relays, see the discussion of the **vacation_template** autoresponse periodicity MTA option in the *Messaging Server Reference*.

Vacation Autoreply Theory of Operation

The vacation action, when invoked, works as follows:

1. Oracle Communications Messaging Server checks to make sure that the vacation action was performed by a user-level rather than a system-level Sieve script. An error results if vacation is used in a system-level script.
2. The "no vacation notice" internal MTA flag is checked. If it is set, processing terminates and no vacation notice is sent.
3. The return address for the message is checked next. If it is blank, processing terminates and no vacation notice is sent.
4. The MTA checks to see if the user's address or any of the additional addresses specified in the **:addresses** tagged argument appear in a **To:**, **Cc:**, **Resent-to:**, or **Resent-cc:** header field for the current message. Processing terminates and no vacation notice is sent if none of the addresses is found in any of the header fields.
5. Messaging Server constructs a hash of the **:subject** argument and the reason string. This string, along with the return address of the current message, is checked against a per-user record of previous vacation responses. Processing terminates and no response is sent if a response has already been sent within the time allowed by the **:days** argument.
6. Messaging Server constructs a vacation notice from the **:subject** argument, reason string, and **:mime** argument. Two basic forms for this response message are possible:
 - A message disposition notification of the form specified in RFC 2298, with the first part containing the reason text.
 - A single part text reply. (This form is only used to support the "reply" autoreply mode attribute setting.)

The **mailautoreplymode** is automatically set to **reply** when vacation messages are configured through Messenger Express.

The "no vacation notice" MTA flag is clear by default. It can be set by a system-level Sieve script through the use of the nonstandard **novacation** action. The **novacation** Sieve action is only allowed in a system-level Sieve script. It generates an error if it is used in a user-level script. You can use this action to implement site-wide restrictions on vacation replies such as blocking replies to addresses containing the substring **"MAILER-DAEMON."**

Per-user per-response information is stored in a set of flat text files, one per local user. The location and naming scheme for these files is specified by the setting of the **mta.vacation_template** option. This option should be set to a **file: URL**.

Maintenance of these files is automatic and controlled by the **mta.vacation_cleanup** integer MTA option setting. Each time one of these files is opened, the value of the current time in seconds modulo this value is computed. If the result is zero the file is scanned and all expired entries are removed. The default value for the option is 200, which means that there is 1-in-200 chance that a cleanup pass is performed.

The machinery used to read and write these flat text files is designed in such a way that it should be able to operate correctly over NFS. This enables multiple MTAs to share a single set of files on a common file system.

Vacation Autoreply Attributes

The set of user LDAP directory attributes that the vacation action uses are:

- Attribute defined by the MTA option **mta.ldap_autoreply_addresses** This attribute provides the ability to generate **:addresses** arguments to sieve vacation. This option has no value by default. The attribute can be multivalued, with each value specifying a separate address to pass to the **:addresses** vacation option.
- Attribute defined by **mta.ldap_personal_name** Alias processing keeps track of personal name information specified in this attribute and uses this information to construct From: fields for any MDNs or vacation replies that are generated. Use with caution to avoid exposing personal information.
- **vacationStartDate** Vacation start date and time. The value is in the format **YYYYMMDDHHMMSSZ**. This value is normalized to GMT. An autoreply should only be generated if the current time is after the time specified by this attribute. No start date is enforced if this attribute is missing. The MTA can be instructed to look at a different attribute for this information by setting the **mta.lda_start_date** MTA option to a different attribute name. This attribute is read and checked by the code that generated the Sieve script. Vacation processing is aborted if the current date is before the vacation start date. This attribute cannot be handled by the script itself because at present Sieve lacks date/time testing and comparison facilities.
- **vacationEndDate** Vacation end date and time. The value is in the format **YYYYMMDDHHMMSSZ**. This value is normalized to GMT. An autoreply should only be generated if the current time is before the time specified by this attribute. No end date is enforced if this attribute is missing. The MTA can be instructed to look at a different attribute for this information by setting the **mta.ldap_end_date** MTA option to a different attribute name. This attribute is and checked by the code that generated the Sieve script. Vacation processing is aborted if the current date is after the vacation end date. This attribute cannot be handled in the script itself because at present Sieve lacks date/time testing and comparison facilities.
- **mailAutoReplyMode** Specifies autoreply mode for the user mail account. Valid values of this attribute are:
 - **echo** - Create a multipart that echoes the original message text in addition to the added **mailAutoReplyText** or **mailAutoReplyTextInternal** text.
 - **reply** - Send a single part reply as specified by either **mailAutoReplyText** or **mailAutoReplyTextInternal** to the original sender. These modes appear in the Sieve script as nonstandard **:echo** and **:reply** arguments to the vacation action. **echo** produces a "processed" message disposition notification (MDN) that contains the original message as returned content. **reply** produces a pure reply containing only the reply text. An illegal value does not manifest as any argument to the vacation action and this produces an MDN containing only the headers of the original message. Selecting an autoreply mode of echo causes an automatic reply to be sent to every message regardless of how recently a previous reply was sent. The MTA can be instructed to use a different attribute for this information by setting the **mta.ldap_autoreply_mode** MTA option to a different attribute name.
- **mailAutoReplySubject** Specifies the contents of the subject field to use in the autoreply response. This must be a UTF-8 string. This value gets passed as the

:subject argument to the vacation action. The MTA can be instructed to use a different attribute for this information by setting the **mta.ldap_autoreply_subject** MTA option to a different attribute name.

- **mailAutoReplyText** Autoreply text sent to all senders except users in the recipient's domain. If not specified, external users receive no vacation message. The MTA can be instructed to use a different attribute for this information by setting the **mta.ldap_autoreply_text** MTA option to a different attribute name.
- **mailAutoReplyTextInternal** Auto-reply text sent to senders from the recipients domain. If not specified, then internal users get the mail autoreply text message. The MTA can be instructed to use a different attribute for this information by setting the **mta.ldap_autoreply_text_int** MTA option to a different attribute name. The MTA passes either the **mailAutoReplyText** or **mailAutoReplyTextInternal** attribute value as the reason string to the vacation action.
- **mailAutoReplyTimeOut** Duration, in hours, for successive autoreply responses to any given mail sender. Used only when **mailAutoReplyMode=reply**. If value is *0* then a response is sent back every time a message is received. This value is converted to the nonstandard **:hours** argument to the vacation action. (Normally the Sieve vacation action only supports the **:days** argument for this purpose and does not allow a value of *0*.) If this attribute does not appear on a user entry, a default time-out is obtained from the **mta.autoreply_timeout_default** MTA option. The MTA can be instructed to use a different attribute for this information by setting the **mta.ldap_autoreply_timeout** MTA option.

The MTA can choose between multiple LDAP attributes and attribute values with different language tags and determine the correct value to use. The language tags in effect are compared against the preferred language information associated with the envelope from address. Currently the only attributes receiving this treatment are **mta.ldap_autoreply_subject** (normally **mailAutoReplySubject**), **mta.ldap_autoreply_text** (normally **mailAutoReplyText**), **mta.ldap_autoreply_text_int** (normally **mailAutoReplyTextInternal**), **mta.ldap_spare_4**, **mta.ldap_spare_5**, **mta.ldap_prefix_text** and **mta.ldap_suffix_text**.

It is expected that each attribute value has a different language tag value. If different values have the same tag value the choice between them is essentially random.

Other Auto Reply Tasks and Issues

This section describes auto reply tasks and issues not described in the configuration section.

Topic in this section:

- [To Send Autoreply Messages for Email That Have Been Automatically Forwarded from Another Mail Server](#)

To Send Autoreply Messages for Email That Have Been Automatically Forwarded from Another Mail Server

An autoreply problem can occur when the MTA receives a message that has been automatically forwarded from another system in some other administrative domain. For example, if a customer has a home account with **example.com** and the customer sets that account to automatically forward messages to their work account at **example.org** and if **example.org** uses Messaging Server and that user has set his account to autoreply a vacation message, then Messaging Server has a problem sending out a vacation message.

The problem occurs because the **example.com** mail server changes the envelope **To:** address from **user@example.com** to **user@example**, but it does not change the **To:** header, which remains **user@example.com**. When the MTA receives the message, it looks at the header address only. It attempts to match this address with an address in the LDAP user directory. If it finds a match with someone who has set autoreply, then a vacation message is sent. Because there is no LDAP address match to **user@example.com**, no vacation message is sent. The problem is that the actual address is in the envelope and not in the header.

Because the recipient's address known to the remote system doing automatic forwarding is not known to correspond to the user by the local system, there needs to be a way for the recipient to make such addresses known to the local system so vacation replies are sent when necessary.

The **:addresses** argument to the Sieve **vacation** action provides this capability. It accepts a list of addresses that correspond to the recipient for purposes of making this check. The attribute defined by the MTA option **ldap_autoreply_addresses** allows specification of such addresses in the user's LDAP entry.

To provide autoreply capability for messages that have been automatically forwarded from mail servers in some other administrative domain, the user or administrator would set the email addresses from where those messages may be forwarded to the attribute defined by **mta.ldap_autoreply_addresses**.

Mail Filtering and Access Control

This information discusses how to filter mail based on its source (sender, IP address and so on) or header strings. Two mail filtering mechanisms are used, controlling access to the MTA by using mapping tables and Sieve server-side rules (SSR).

Limiting access to the MTA by using the mapping tables enables messages to be filtered based on From: and To: addresses, IP address, port numbers, and source or destination channels. Mapping tables permit SMTP relaying to be enabled or disabled. Sieve is a mail filtering script that enables messages to be filtered based on strings found in headers.

If envelope-level controls are desired, use mapping tables to filter mail. If header-based controls are desired, use Sieve server-side rules.

Topics:

"PART 1. MAPPING TABLES." Enables the administrator to control access to MTA services by configuring certain mapping tables. The administrator can control who can or cannot send mail, receive mail through Oracle Communications Messaging Server.

"PART 2. MAILBOX FILTERS." Enables users and administrators to filter messages and specify actions on those filtered messages based on a string found in the message header. Uses the Sieve filter language and can filter at the channel, MTA or user level.

PART 1. MAPPING TABLES

The following topics contain the specifics of each mapping table. See the *Messaging Server Reference* for the details about the format and operation, patterns, templates, and metacharacters which are common to all mapping tables.

- [Controlling Access with Mapping Tables](#)
- [Access Control Mapping Table Flags](#)
- [Send Access and Mail Access Mapping Tables](#)
- [FROM_ACCESS Mapping Table](#)
- [PORT_ACCESS Mapping Table](#)
- [IP_ACCESS Mapping Table](#)
- [When Access Controls Are Applied](#)
- [To Test Access Control Mappings](#)
- [To Limit Specified IP Address Connections to the MTA](#)
- [To Add SMTP Relaying](#)

- [Configuring SMTP Relay Blocking](#)
- [Handling Large Numbers of Access Entries](#)
- [Controlling the Envelope From: Address in Mappings Strings and Mailing List Named Options and LDAP Attributes](#)

Controlling Access with Mapping Tables

You can control access to your mail services by configuring certain mapping tables. These mapping tables allow you to control who can or cannot send mail, receive mail, or both. Table 13–1, "Access Control Mapping Tables" lists the mapping tables described in this section. The application information string supplied to the **FROM_ACCESS**, **MAIL_ACCESS**, and **ORIG_MAIL_ACCESS** mappings includes the system name claimed in the **HELO/EHLO** SMTP command. This name appears at the end of the string and is separated from the rest of the string (normally "SMTP") by a slash. The claimed system name can be useful in blocking some worms and viruses.

Access Control Mapping Tables - Operation

Access control mapping tables follow the same general format and operations as all mappings tables. (See the *Messaging Server Reference* for more details about the format and operation, and patterns, templates, and metacharacters which apply to all mapping tables.) They consist of a mapping table name, followed by a blank line, followed by one or more mapping entries. Mapping entries consist of a *search pattern* on the left side and a *template* on the right side. The search pattern filters specific messages and the template specifies actions to take on the message. For example:

```
SEND_ACCESS

*|Elvis1@example.org|*|*      $Y
*|Nelson7@example.org|*|*    $Y
*|AkiraK@example.org|*|*     $Y
*|*example.org|*|*           $NMail$ Blocked
```

In this example all email from the domain **example.org** except those of **Elvis1**, **Nelson**, **AkiraK** are blocked.

The search pattern for access control mapping entries consist of a number of search criteria separated by vertical bars (|). The order of the search criteria depends on the access mapping table and is described in subsequent sections. But as an example, the **SEND_ACCESS** mapping table has the following search form:

```
src-channel|from-address|dst-channel|to-address
```

where *src-channel* is the channel queuing the message; *from-address* is the address of the message's originator; *dst-channel* is the channel to which the message will be queued; and *to-address* is the address to which the message is addressed. Use of an asterisk in any of these four fields causes that field to match any channel or address, as appropriate.

The *from-address* is a positional content string placeholder. In some circumstances, you can modify the default type of information, for example, by using the **use_auth_return** MTA option. An option might affect only substrings in certain mappings and might be subject to precedence rules of options affecting the same mapping table substring. For more information, see "Controlling the Envelope From: Address in Mappings Strings and Mailing List Named Options and LDAP Attributes."

Note: Whenever a mappings table is modified, you must recompile the configuration. See "Compiling the MTA Configuration."

Table 13–1 describes the different mapping tables.

Table 13–1 Access Control Mapping Tables

Mapping Table	Description
SEND_ACCESS (See SEND_ACCESS and ORIG_SEND_ACCESS Mapping Tables)	Used to block incoming connections based on envelope From address, envelope To address, source and destination channels. The To address is checked after rewriting, alias expansion, and so on, have been performed.
ORIG_SEND_ACCESS (See SEND_ACCESS and ORIG_SEND_ACCESS Mapping Tables)	Used to block incoming connections based on envelope From address, envelope To address, source and destination channels. The <i>original To</i> address after rewriting but before alias expansion is used, but the function of alias expansion has already been performed.
MAIL_ACCESS (See MAIL_ACCESS and ORIG_MAIL_ACCESS Mapping Tables)	Used to block incoming connections based on combined information found in SEND_ACCESS and PORT_ACCESS tables: that is, the channel and address information found in SEND_ACCESS combined with the IP address and port number information found in PORT_ACCESS .
ORIG_MAIL_ACCESS (See MAIL_ACCESS and ORIG_MAIL_ACCESS Mapping Tables)	Used to block incoming connections based on combined information found in ORIG_SEND_ACCESS and PORT_ACCESS tables: that is, the channel and address information found in ORIG_SEND_ACCESS combined with the IP address and port number information found in PORT_ACCESS .
FROM_ACCESS Mapping Table	Used to filter mail based on envelope From addresses. Use this table if the To address is irrelevant.
PORT_ACCESS Mapping Table	Used to block incoming connections based on IP address and TCP port.
IP_ACCESS Mapping Table.	Used to control outgoing connections based on IP address information.

The **MAIL_ACCESS** and **ORIG_MAIL_ACCESS** mappings are the most general, having available not only the address and channel information available to **SEND_ACCESS** and **ORIG_SEND_ACCESS**, but also any information that would be available via the **PORT_ACCESS** mapping table, including IP address and port number information.

Access Control Mapping Table Flags

This section consists of the following subsections:

- [Input vs Output Flags](#)
- [Output Flag Argument Order](#)

See the *Messaging Server Reference* for more details about the format and operation, and patterns, templates, and metacharacters which apply to all mapping tables.

Input vs Output Flags

Mapping table templates can test input flags and can set output flags. Input and output flags should be thought of as being stored in separate namespaces. For

example, the **\$A** input flag may have nothing to do with a **\$A** output flag. "**\$A**" is a short hand way of saying "the **A** flag" (as opposed to the string "**A**"). There are input flags set by the MTA when it does a probe - so there is the probe string, plus any probe input flags. Then there is the output string plus output flags that were set on the right hand side of the mapping entry.

So **\$A** could be set by the MTA as an input flag, then testable (on the right hand side of an entry) via **\$.A** or **\$.A**. Then the right hand side of an entry may set its own output flags (that is, **\$A**) which as output means something entirely different - and which do not count as input flags for testing purposes (regardless of whether the mapping "iterates" - say with **\$C**). That is, the flag tests such as **\$.flag-char** are always on the *original* (set by the MTA, or via **-flags** input setting in **imsimta test -mapping**) input flags, not on any flags that may have been set in the right-hand-side of the mapping.

For example, try using **imsimta test -mapping**, especially with its **-flags** switch to set some input flags. Try using some arbitrary, private mapping table. Use a few **\$.flag** or **\$.flag** tests on your right hand side. Also notice that the output lists (separately):

```
Output string: ...
Output flags: [...]
```

where the output flags (the stuff inside the **[]**'s) correspond to the output flags set. That is, with a mapping:

```
X-FLAGS
```

```
*           $C$.AA$ set$.Y
*           $.AA$ not$. set$.Y
```

try something like:

```
# imsimta test -mapping -flags=A -table=x-flags
```

and note how it yields

```
Output string: A set
Output flags: [0, 'Y' (89)]
```

meaning that it output the string "A set", matched on the 0th iteration (as opposed to iterative or recursive mappings where you'll see the number go up) and output the **Y** flag (or **\$.Y** if you want to write it that way) but the **A** flag that was an input flag is *not* an output flag. **'Y'** is character position 89 in ASCII - the character position is shown for extra clarity.

In contrast,

```
# imsimta test -mapping -flags=B -table=x-flags
```

yields

```
Output string: A not set
Output flags: [0, 'Y' (89)]
```

Output Flag Argument Order

Flags with arguments must have those arguments arranged in the reading order shown in the table. For example:

```
ORIG_SEND_ACCESS
```

```
tcp_local|*|tcp_local|*   $.N$.D30|Relaying$. not$. allowed
```

In this case, the proper order is the delay period followed by the rejection string. Note that the flags themselves can be in any order. So the following entries have identical results:

```
30|Relaying$ not$ allowed$D$N
$N30|Relaying$ not$ allowed$D
30|$N$DRelaying$ not$ allowed
```

For the details of the flags and metacharacters supported in each table, see:

- [Send Access and Mail Access Mapping Tables](#)
- [FROM_ACCESS Mapping Table](#)
- [PORT_ACCESS Mapping Table](#)
- [IP_ACCESS Mapping Table](#)

Send Access and Mail Access Mapping Tables

The following flags and metacharacters apply to the **SEND_ACCESS**, **ORIG_SEND_ACCESS**, **MAIL_ACCESS**, and **ORIG_MAIL_ACCESS** mapping tables. The **FROM_ACCESS** and **PORT_ACCESS** mapping tables have different sets of flags - see the "FROM_ACCESS Mapping Table" and "PORT_ACCESS Mapping Table" sections.

- [Table 13–2, "Send and Mail Access Input Flags"](#)
- [Table 13–3, "Send and Mail Access Output Flags"](#)
- [SEND_ACCESS and ORIG_SEND_ACCESS Mapping Tables](#)
- [MAIL_ACCESS and ORIG_MAIL_ACCESS Mapping Tables](#)

Input flags are set by the process that calls the mapping table and tested for by using the **\$:** and **\$;** metacharacters in the template. The **\$:** will succeed if the flag is set or fail if it is not set. The **\$;** will succeed if the flag is not set or fail if it is set. [Table 13–2](#) shows the **\$:x** form to test if the flag is set. The **\$;x** form can also be used to test if condition is not true. [Table 13–2](#) describes the send and mail access input flags.

Table 13–2 *Send and Mail Access Input Flags*

Flag	Description
\$:l	Match only if external material (that is, an envelope address) in the probe contained a vertical bar.
\$:A	Match only if SASL (SMTP AUTH) has been used.
\$:D	Match only if DELAY delivery receipts have been requested (that is, NOTIFY=DELAY).
\$:E	Match only if ESMTP has been used.
\$:F	Match only if FAILURE delivery receipts have been requested (that is, NOTIFY=FAILURE).
\$:L	Match only if LMTP has been used.
\$:P	Match only if POP-before-SMTP was used.
\$:S	Match only if SUCCESS delivery receipts have been requested (that is, NOTIFY=SUCCESS).
\$:T	Match only if TLS has been used.
\$:V	Match only if recipient address expanded via an alias.

Output flags are set by the template in the mapping table entry and consumed upon return to the calling program. [Table 13–3](#) describes the send and mail access output flags.

Table 13–3 Send and Mail Access Output Flags

Flag	Description
\$.	Establishes a string which will be processed as the mapping entry result in the event of a temporary LDAP lookup failure. By default a temporary failure string remains set only for the duration of the current rule. "\$.." can be used to return to the default state where no temporary failure string is set and temporary LDAP failures cause mapping entry or rewrite rule failure. Note that all errors other than failure to match an entry in the directory are considered to be temporary errors; in general it isn't possible to distinguish between errors caused by incorrect LDAP URLs and errors caused by directory server configuration problems.
\$*	If used with \$N, force disconnect of the SMTP session.
\$B	Redirect the message to the bitbucket; the effect is per-recipient in the Send/Mail Access mapping tables, vs entire message in the " FROM_ACCESS Mapping Table ."
\$H	Hold the message as a .HELD file; the effect is per-recipient in the Send/Mail Access mapping tables, vs entire message in the " FROM_ACCESS Mapping Table ."
\$J	Used with \$M (see below) to cause generation of envelope "journal" format rather than the default DSN encapsulated format for the "capture" message copy (added in 7u3p16)
\$O	Forces single-copy-per-recipient message copy "split up." For more information see the discussion on addresses per message copy in the <i>Messaging Server Reference</i> .
\$P	Force to enqueue to the reprocess channel. (For instance, this might be useful as part of \$.text. handling when attempting an LDAP lookup that encountered an LDAP directory temporary problem; that is, in case of an LDAP lookup problem, defer to the reprocess channel.)
\$V	Force Sieve filter discard behavior for all recipients of this message copy.
\$v	Force Sieve filter discard behavior for solely this recipient.
\$Y	Allow access.
\$Z	Force Sieve filter jettison behavior (non-overridable discard) for all recipients of this message copy.
\$z	Force Sieve filter jettison behavior (non-overridable discard) for solely this recipient.

[Table 13–4](#) describes send and mail access flags with arguments. For more information see "[Output Flag Argument Order](#)." Note: Do **NOT** alphabetize this list.

Table 13–4 Send and Mail Access Flags with Arguments

Flag	Description
\$U <i>integer</i>	Enable channel (slave) debugging; if the optional <i>n</i> argument is specified, it also sets the specified value for enqueue debugging (see <code>MM_DEBUG</code> option). See the discussion on debugging channel master and slave programs in the <i>Messaging Server Reference</i> .
\$I <i>user identifier</i>	Check specified user for specified groupid (UNIX) and if not in the group, reject access (effectively sets the <code>\$N</code> output flag).
\$< <i>string</i>	+++ Send <i>string</i> to UNIX syslog (user.notice facility and severity) if probe matches. (see <code>SNDOPR_PRIORITY</code> option)
\$> <i>string</i>	+++ Send <i>string</i> to UNIX syslog (user.notice facility and severity) if access is rejected. (see <code>SNDOPR_PRIORITY</code> option)
\$D <i>delay</i>	Delay response for an interval of <i>delay</i> hundredths of seconds; a positive value causes the delay to be imposed on each command in the transaction; a negative value causes the delay to be imposed only on the address handover (SMTP MAIL FROM: command for the <code>FROM_ACCESS</code> table; SMTP RCPT TO: command for the other tables).
\$T <i>tag</i>	Prefix subsequent <i>*ACCESS mapping table probes with tag _tag</i>
\$A <i>header</i>	Add the header line <i>header</i> to the message.
\$G <i>conversion_tag</i>	If used in <code>ORIG_SEND_ACCESS</code> , <code>SEND_ACCESS</code> , <code>ORIG_MAIL_ACCESS</code> , and <code>MAIL_ACCESS</code> , it reads a value from the mapping result and treats it as a set of conversion tags to be applied to the current recipient. If used with <code>FROM_ACCESS</code> , conversion tags are applied to all recipients. \$G is positioned after \$A (header address) in the sequence of arguments read from the mappings. See "Mail Conversion Tags."
\$M <i>address</i>	Capture a copy of the message, sending the captured copy to <i>address</i> . By default, this captured copy is DSN encapsulated---but see the no-argument, non- <code>FROM_ACCESS</code> <code>\$J</code> output flag above.
\$, <i>spamadjust_arg</i>	Set a spam level value <i>x</i> (between -10000.0 and 10000.0). If the message already had a spam level, this is a spamadjust effect (adds or subtracts the specified amount <i>x</i> from the prior spam level). Note that such a spam level/spamadjust effect applies to all recipients (even for the recipient-specific mapping tables such as <code>SEND_ACCESS</code>) in order for tests to see if one of the recipients is a "honeypot" address. Allows you to perform a sieve spamadjust operation from the access mapping tables. Argument takes the same form as a spamadjust argument. Note also that some of these mappings are applied on a per-recipient basis. Any spamadjust operation that is done applies to all recipients.
\$+ <i>ename value</i>	Set the environment variable <i>name</i> to <i>value</i> .
\$X <i>error-code</i>	Lets you specify the extended SMTP <i>error-code</i> (the digit.digit.digit part), and if the first digit is a 4 rather than a 5, then you'll get a 452 SMTP temporary error, rather than the usual 550 SMTP permanent error. For example: <code>ORIG_SEND_ACCESS</code> <...probe...> \$N\$X4.5.9 Temporary\$ problem\$ with\$ address;\$ try\$ later\$

Table 13–4 (Cont.) Send and Mail Access Flags with Arguments

Flag	Description
\$N <i>string</i> \$F <i>string</i>	Reject access with the optional error text <i>string</i> . \$F and \$N are synonyms.
\$S <i>url</i>	Apply the Sieve filter obtained from resolving <i>url</i> .
\$+R <i>n</i> <i>string</i>	Opt in to spam filtering. Only applied if neither \$N nor \$F is set. <i>n</i> is which spam/virus filter package; <i>string</i> is the option string to pass to that spam filter.

{+} To use multiple flags with arguments, separate the arguments with the vertical bar character, |, placing the arguments in the order listed in this table. For more information see "[Output Flag Argument Order](#)."

+++ It is a good idea to use the **\$D** flag when dealing with problem senders, to prevent a denial of service attack. In particular, it is a good idea to use **\$D** in any **\$>** entry or **<** entry rejecting access.

SEND_ACCESS and ORIG_SEND_ACCESS Mapping Tables

You can use the **SEND_ACCESS** and **ORIG_SEND_ACCESS** mapping tables to control who can or cannot send mail, receive mail, or both. The access checks have available a message's envelope **From:** address and envelope **To:** addresses, and knowledge of what channel the message came in, and what channel it would attempt to go out.

If a **SEND_ACCESS** or **ORIG_SEND_ACCESS** mapping table exists, then for each recipient of every message passing through the MTA, the MTA will scan the table with a string of the following form (note the use of the vertical bar character, |):

```
src-channel | from-address | dst-channel | to-address
```

The *src-channel* is the channel queuing the message; *from-address* is the address of the message's originator; *dst-channel* is the channel to which the message will be queued; and *to-address* is the address to which the message is addressed. Use of an asterisk in any of these four fields causes that field to match any channel or address, as appropriate.

The addresses here are envelope addresses; that is, envelope **From:** address and envelope **To:** address. In the case of **SEND_ACCESS**, the envelope **To:** address is checked after rewriting, alias expansion, etc., have been performed; in the case of **ORIG_SEND_ACCESS** the originally specified envelope **To:** address is checked after rewriting, but before alias expansion. Note however that the function of alias expansion has already been performed at this point, the **ORIG_*** variant simply uses a copy of the address from before expansion.

If the search string matches a pattern (that is, the left-hand side of an entry in the table), then the resulting output of the mapping is checked. If the output contains the flags **\$Y** or **\$y**, then the enqueue for that particular **To:** address is permitted. If the output contains any of the flags **\$N**, **\$n**, **\$F**, or **\$f**, then the enqueue to that particular address is rejected. In the case of a rejection, optional rejection text may be supplied in the mapping output. This string will be included in the rejection error the MTA issues. If no string is output (other than the **\$N**, **\$n**, **\$F**, or **\$f** flag), then default rejection text will be used. For descriptions of additional flags, see "[Access Control Mapping Table Flags](#)."

See "Send Access and Mail Access Mapping Tables" for the complete list of flags which apply to the **SEND_ACCESS** and **ORIG_SEND_ACCESS** mapping tables.

Setting the MTA option **ACCESS_ORCPT** to 1 adds an additional vertical bar delimited field to the probe value passed to the **SEND_ACCESS**, **ORIG_SEND_ACCESS**, **MAIL_ACCESS**, and **ORIG_MAIL_ACCESS** mapping tables that contains the original recipient (ORCPT) address. If the message doesn't have an ORCPT address, the original unmodified RCPT TO: address is used instead. The default is 0, and the probe value is at the end:

```
src-channel|from-address|dst-channel|to-address|ORCPT_address
```

In the following example, mail sent from UNIX user agents such as mail, Pine, and so on, originates from the local, **l**, channel and messages to the Internet go out a TCP/IP channel of some sort. Suppose that local users, with the exception of the postmaster, are not allowed to send mail to the Internet but can receive mail from there. Then the **SEND_ACCESS** mapping table shown in the example below is one possible way to enforce this restriction. In the mapping table, the local host name is assumed to be **example.org**. In the channel name "**tcp_***," a wild card is used so as to match any possible TCP/IP channel name (for example, **tcp_local**).

Example - SEND_ACCESS Mapping Table

```
SEND_ACCESS
```

```
*|postmaster@example.org|*|*      $Y
*|*|*|postmaster@example.org      $Y
l|*@example.org|tcp_*|*           $NInternet$ postings$ are$ not$ permitted
```

In the rejection message, dollar signs are used to quote spaces in the message. Without those dollar signs, the rejection would be ended prematurely and only read "Internet" instead of "Internet postings are not permitted." Note that this example ignores other possible sources of "local" postings such as from PC-based mail systems or from POP or IMAP clients.

Note: The client attempting to send the message determines whether the MTA rejection error text is actually presented to the user who attempted to send the message. If **SEND_ACCESS** is used to reject an incoming SMTP message, the MTA merely issues an SMTP rejection code including the optional rejection text; it is up to the sending SMTP client to use that information to construct a bounce message to send back to the original sender.

MAIL_ACCESS and ORIG_MAIL_ACCESS Mapping Tables

The **MAIL_ACCESS** mapping table is a superset of the **SEND_ACCESS** and **PORT_ACCESS** mapping tables. It combines both the channel and address information of **SEND_ACCESS** with the IP address and port number information of **PORT_ACCESS**. Similarly, the **ORIG_MAIL_ACCESS** mapping table is a superset of the **ORIG_SEND_ACCESS** and **PORT_ACCESS** mapping tables. The format for the probe string for **MAIL_ACCESS** is:

```
port-access-probe-info|app-info|submit-type|send_access-probe-info
```

Similarly, the format for the probe string for **ORIG_MAIL_ACCESS** is:

```
port-access-probe-info|app-info|submit-type|orig_send_access-probe-info
```

Here *port-access-probe-info* consists of all the information usually included in a **PORT_ACCESS** mapping table probe in the case of incoming SMTP messages; otherwise, it is blank. *app-info* includes the system name claimed in the **HELO/EHLO** SMTP command. This name appears at the end of the string and is separated from the rest of the string (normally "SMTP*") by a slash. The claimed system name can be useful in blocking some worms and viruses. *submit-type* may be one of MAIL, SEND, SAML, or SOML, corresponding to how the message was submitted into Messaging Server. Normally the value is MAIL, meaning it was submitted as a message; SEND, SAML, or SOML can occur in the case of broadcast requests (or combined broadcast/message requests) submitted to the SMTP server. And for the **MAIL_ACCESS** mapping, *send-access-probe-info* consists of all the information usually included in a **SEND_ACCESS** mapping table probe. Similarly for the **ORIG_MAIL_ACCESS** mapping, *orig-send-access-probe-info* consists of all the information usually included in an **ORIG_SEND_ACCESS** mapping table probe.

See "Send Access and Mail Access Mapping Tables" for the list of flags which apply to the **MAIL_ACCESS** and **ORIG_MAIL_ACCESS** mapping tables.

See the *Messaging Server Reference* for more details about the format and operation, and patterns, templates, and metacharacters which apply to all mapping tables.

Setting the MTA option **ACCESS_ORCPT** to 1 adds an additional vertical bar delimited field to the probe value passed to the **SEND_ACCESS**, **ORIG_SEND_ACCESS**, **MAIL_ACCESS**, and **ORIG_MAIL_ACCESS** mapping tables that contains the original recipient (**ORCPT**) address. If the message doesn't have an **ORCPT** address, the original unmodified **RCPT TO:** address is used instead. The default is 0, and the probe value is at the end. Example:

```
port-access-probe-info|app-info|submit-type|send_access-probe-info|ORCPT_address
```

Having the incoming TCP/IP connection information available in the same mapping table as the channel and address information makes it more convenient to impose certain sorts of controls, such as enforcing what envelope **From:** addresses are allowed to appear in messages from particular IP addresses. This can be desirable to limit cases of email forgery, or to encourage users to configure their POP and IMAP clients' **From:** address appropriately. For example, a site that wishes to allow the envelope **From:** address **vip@example.com** to appear only on messages coming from the IP address 1.2.3.1 and 1.2.3.2, and to ensure that the envelope **From:** addresses on messages from any systems in the 1.2.0.0 subnet are from **example.com**, might use a **MAIL_ACCESS** mapping table as shown in the following example.

Example - MAIL_ACCESS Mapping Table

```
MAIL_ACCESS
```

```
! Entries for vip's two systems
!
TCP|*|25|1.2.3.1|*|SMTP*|MAIL|tcp_*|vip@example.com|*|* $Y
TCP|*|25|1.2.3.2|*|SMTP*|MAIL|tcp_*|vip@example.com|*|* $Y
!
! Disallow attempts to use vip's From: address from other
! systems
!
TCP|*|25|*|*|SMTP*|MAIL|tcp_*|vip@example.com|*|* \
    $N500$ Not$ authorized$ to$ use$ this$ From:$ address
!
! Allow sending from within our subnet with example.com From:
! addresses
```

```

!
TCP|*|25|1.2.*.*|*|SMTP*|MAIL|tcp_*|*|example.com|*|* $Y
!
! Allow notifications through
!
TCP|*|25|1.2.*.*|*|SMTP*|MAIL|tcp_*|*|* $Y
!
! Block sending from within our subnet with non-example.com
! addresses
!
TCP|*|25|1.2.*.*|*|SMTP*|MAIL|tcp_*|*|*|* \
$NOnly$ example.com$ From:$ addresses$ authorized

```

FROM_ACCESS Mapping Table

The following flags and metacharacters apply to the **FROM_ACCESS** mapping table. The **SEND_ACCESS**, **ORIG_SEND_ACCESS**, **MAIL_ACCESS**, and **ORIG_MAIL_ACCESS** mapping tables have a similar set of flags but there are significant differences. See "Send Access and Mail Access Mapping Tables" for more information. The "PORT_ACCESS Mapping Table" has a different set of flags.

- Table 13–5, "FROM_ACCESS Input Flags"
- Table 13–6, "FROM_ACCESS Output Flags"
- FROM_ACCESS Description

See the *Messaging Server Reference* for more details about the format and operation, and patterns, templates, and metacharacters which apply to all mapping tables.

Input flags are set by the process that calls the mapping table and tested for by using the **\$:** and **\$;** metacharacters in the template. The **\$:** will succeed if the flag is set or fail if it is not set. The **\$;** will succeed if the flag is not set or fail if it is set. Table 13–5 shows the **\$:x** form to test if the flag is set. The **\$;x** form can also be used to test if condition is not true.

Table 13–5 FROM_ACCESS Input Flags

Flag	Description
\$: 	Match only if external material (that is, an envelope address) in the probe contained a vertical bar.
\$:A	Match only if SASL (SMTP AUTH) has been used.
\$:E	Match only if ESMTP has been used.
\$:L	Match only if LMTP has been used.
\$:P	Match only if POP-before-SMTP was used.
\$:T	Match only if TLS has been used.

Output flags are set by the template in the mapping table entry and consumed upon return to the calling program. Table 13–6 shows the FROM_ACCESS output flags.

Table 13–6 FROM_ACCESS Output Flags

Flag	Description
\$.	Establishes a string which will be processed as the mapping entry result in the event of a temporary LDAP lookup failure. By default a temporary failure string remains set only for the duration of the current rule. "\$.." can be used to return to the default state where no temporary failure string is set and temporary LDAP failures cause mapping entry or rewrite rule failure. Note that all errors other than failure to match an entry in the directory are considered to be temporary errors; in general it isn't possible to distinguish between errors caused by incorrect LDAP URLs and errors caused by directory server configuration problems.
\$/	Set the "fast disconnect" flag for sessions that have not yet succeeded in starting a transaction; for such sessions, any subsequent disconnect is done with SO_LINGER enabled and a timeout of 0, which may clear slots quicker on intermediate firewalls and proxies.
#!	Disable (Sieve requested) vacation messages regarding this message.
\$*	If used with \$N, force disconnect of the SMTP session.
\$B	Redirect the message to the bitbucket; effects the entire message (all recipients) when used in FROM_ACCESS , vs. per-recipient in Table 13–3, "Send and Mail Access Output Flags".
\$H	Hold the message as a .HELD file; effects the entire message (all recipients) when used in FROM_ACCESS , vs. per-recipient in Table 13–3, "Send and Mail Access Output Flags".
\$O	Forces single-copy-per-recipient message copy "split up," as if the single channel option were set on the relevant destination channel(s). For more information see the discussion on addresses per message copy in the <i>Messaging Server Reference</i> .
\$P	Force to enqueue to the reprocess channel. (For instance, this might be useful as part of \$.text. handling when attempting an LDAP lookup that encountered an LDAP directory temporary problem; that is, in case of an LDAP lookup problem, defer to the reprocess channel.)
\$V	Force Sieve filter discard behavior for all recipients of this message copy.
\$Y	Allow access.
\$Z	Force Sieve filter jettison behavior (non-overridable discard) for all recipients of this message copy.
\$z	Force Sieve filter jettison - synonymous with \$Z. Also see per-recipient behavior in Table 13–3, "Send and Mail Access Output Flags".

Table 13–7 describes FROM_ACCESS flags with arguments. For more information see "Output Flag Argument Order." Note: Do **NOT** alphabetize this list.

Table 13-7 FROM_ACCESS Flags with Arguments

Flag	Description
\$Uinteger	Enable channel (slave) debugging; if the optional <i>n</i> argument is specified, it also sets the specified value for enqueue debugging (see MM_DEBUG option). See the discussion on debugging channel master and slave programs in the <i>Messaging Server Reference</i> for more information.
\$~channel-name	Change source channel to <i>_channel-name_</i> . This may be of especial interest for the case of incoming notification messages (incoming messages with empty envelope From); note that when this feature is used and it actually changes the source channel, then the FROM_ACCESS check process is restarted; also, \$~ can only be applied once.
\$Jaddress	Replace original envelope From: address with specified <i>address</i> .
\$Kaddress	Replace the MTA's internal sender address with specified <i>address</i> for subsequent checking. Note: to also add a Sender: header, the source channel must include the authrewrite channel option.
\$User \ identifier	Check specified user for specified groupid (UNIX) and if not in the group, reject access (effectively sets the \$N output flag).
\$<string	+++ Send <i>string</i> to UNIX syslog (user.notice facility and severity) if probe matches. (see SNDOPR_PRIORITY option).
\$>string	+++ Send <i>string</i> to UNIX syslog (user.notice facility and severity) if access is rejected. (see SNDOPR_PRIORITY option).
\$Ddelay	Delay response for an interval of <i>delay</i> hundredths of seconds; a positive value causes the delay to be imposed on each command in the transaction; a negative value causes the delay to be imposed only on the address handover (SMTP MAIL FROM: command for the FROM_ACCESS table; SMTP RCPT TO: command for the other tables).
\$Ttag	Prefix subsequent *_ACCESS mapping table probes with tag <i>tag</i>
\$Aheader	Add the header line <i>header</i> to the message.
\$Gconversion_tag	If used in ORIG_SEND_ACCESS , SEND_ACCESS , ORIG_MAIL_ACCESS , and MAIL_ACCESS , it reads a value from the mapping result and treats it as a set of conversion tags to be applied to the current recipient. If used with FROM_ACCESS , conversion tags are applied to all recipients. \$G is positioned after \$A (header address) in the sequence of arguments read from the mappings. See "Mail Conversion Tags."
\$Maddress	Capture a copy of the message, sending the captured copy to <i>address</i> . By default, this captured copy is DSN encapsulated---but see the no-argument, non-FROM_ACCESS \$J flag above.
\$Sx,y,z	Set a blocklimit , and optionally recipientlimit , and optionally recipientcutoff for the transaction. These values override any global MTA option or source-based such limits that may already be effect (but destination-based limits will still be applied, later). See the discussion on message size limits in the <i>Messaging Server Reference</i> .

Table 13-7 (Cont.) FROM_ACCESS Flags with Arguments

Flag	Description
<code>\$,spamadjust_arg</code>	Set a spam level value <i>x</i> (between -10000.0 and 10000.0). If the message already had a spam level, this is a spamadjust effect (adds or subtracts the specified amount <i>x</i> from the prior spam level). Note that such a spam level/spamadjust effect applies to all recipients (even for the recipient-specific mapping tables such as <code>SEND_ACCESS</code>) in order for tests to see if one of the recipients is a "honeypot" address. Allows you to perform a sieve <code>spamadjust</code> operation from the access mapping tables. Argument takes the same form as a <code>spamadjust</code> argument. Note also that some of these mappings are applied on a per-recipient basis. Any <code>spamadjust</code> operation that is done applies to all recipients.
<code>\$(postmaster-address</code>	Set an override postmaster address.
<code>)postmaster-address</code>	Set an override postmaster address if none was previously set.
<code>\$+Ename value</code>	Set the environment variable <i>name</i> to <i>value</i> .
<code>\$+Rn string</code>	Opt in to spam filtering. <i>n</i> specifies which spam/virus filter package; <i>string</i> is the optin string to pass to that spam filter.
<code>\$Xerror-code</code>	Lets you specify the extended SMTP <i>error-code</i> (the digit.digit.digit part), and if the first digit is a 4 rather than a 5, then you'll get a 452 SMTP temporary error, rather than the usual 550 SMTP permanent error. For example: ORIG_SEND_ACCESS <...probe...> \$N\$X4.5.9 Temporary\$ problem\$ with\$ address;\$ try\$ later\$
<code>\$Nstring</code>	Reject access with the optional error text <i>string</i> .
<code>\$Fstring</code>	<code>\$F</code> and <code>\$N</code> are synonyms.

{+} To use multiple flags with arguments, separate the arguments with the vertical bar character, |, placing the arguments in the order listed in this table. For more information see "Output Flag Argument Order."

+++ It is a good idea to use the `$D` flag when dealing with problem senders, to prevent a denial of service attack. In particular, it is a good idea to use `$D` in any `$>` entry or `<` entry rejecting access.

FROM_ACCESS Description

The `FROM_ACCESS` mapping table may be used to control who can send mail, or to override purported **From:** addresses with authenticated addresses, or both.

The input probe string to the `FROM_ACCESS` mapping table is similar to that for a `MAIL_ACCESS` mapping table, minus the destination channel and address, and with the addition of authenticated sender information, if available. Thus, if a `FROM_ACCESS` mapping table exists, then for each attempted message submission, Messaging Server will search the table with a string of the form (note the use of the vertical bar character, |):

```
port-access-probe-info|app-info|submit-type|src-channel|from-address|auth-from
```

Here *port-access-probe-info* consists of all the information usually included in a `PORT_ACCESS` mapping table probe in the case of incoming SMTP messages; otherwise, it is blank. *app-info* includes the system name claimed in the `HELO/EHLO` SMTP command. This name appears at the end of the string and is separated from the rest of the string (normally "SMTP*") by a slash. The claimed system name can be useful in

blocking some worms and viruses. *submit-type* may be one of MAIL, SEND, SAML, or SOML, corresponding to how the message was submitted into the MTA. Normally the value is MAIL, meaning it was submitted as a message; SEND, SAML, or SOML can occur in the case of broadcast requests (or combined broadcast/message requests) submitted to the SMTP server. *src-channel* is the channel originating the message (that is, queuing the message); *from-address* is the address of the message's purported originator; and *auth-from* is the authenticated originator address, if such information is available, or blank if no authenticated information is available.

If the probe string matches a pattern (that is, the left-hand side of an entry in the table), the resulting output of the mapping is checked. If the output contains the flags \$Y or \$y, then the enqueue for that particular **To:** address is permitted. If the output contains any of the flags \$N, \$n, \$F, or \$f, then the enqueue to that particular address is rejected. In the case of a rejection, optional rejection text may be supplied in the mapping output. This string will be included in the rejection error Messaging Server issues. If no string is output (other than the \$N, \$n, \$F, or \$f flag), then default rejection text will be used. For descriptions of additional flags, see "[Access Control Mapping Table Flags](#)."

Besides determining whether to allow a message to be submitted based on the originator, **FROM_ACCESS** can also be used to alter the envelope **From:** address via the **\$J** flag, or to modify the effect of the **authrewrite** channel option (adding a **Sender:** header address on an accepted message) via the **\$K** flag. For instance, this mapping table can be used to cause the original envelope **From:** address to simply be replaced by the authenticated address.

Example - FROM_ACCESS Mapping Table

FROM_ACCESS

```
* |SMTP*|*|tcp_auth|*|      $Y
* |SMTP*|*|tcp_auth|*|*    $Y$J$4
```

When using the **FROM_ACCESS** mapping table to modify the effect on having **authrewrite** set to a nonzero value on some source channel, it is not necessary to use **FROM_ACCESS** if the authenticated address is going to be used verbatim.

For example, with **authrewrite 2** set on the **tcp_local** channel, the following **FROM_ACCESS** mapping table would not be necessary because **authrewrite** alone is sufficient to get this effect (adding the authenticated address verbatim):

FROM_ACCESS

```
* |SMTP*|*|tcp_auth|*|      $Y
* |SMTP*|*|tcp_auth|*|*    $Y$K$4
```

However, the real purpose of **FROM_ACCESS** is to permit more complex and subtle alterations, as shown in the example below. The **authrewrite** channel option alone is appropriate if you want to add a **Sender:** header line (showing the SMTP AUTH authenticated submitter address) to incoming messages. However, suppose you want to force the addition of such a **Sender:** header line to incoming messages only if the SMTP AUTH authenticated submitter address differs from the envelope **From:** address (that is, not bother to add a **Sender:** header line if the addresses match), and suppose further that you wish the SMTP AUTH and envelope **From:** addresses will not be considered to differ merely because the envelope **From:** includes optional subaddress information.

FROM_ACCESS

```

! If no authenticated address is available, do nothing
*|SMTP*|*|tcp_auth|*| $Y
! If authenticated address matches envelope From:, do nothing
*|SMTP*|*|tcp_auth|*|$3* $Y
! If authenticated address matches envelope From: sans
! subaddress, do nothing
*|SMTP*|*|tcp_auth|*+*@$3*$5* $Y
! Fall though to...
! ...authenticated address present, but didn't match, so force
! Sender: header
*|SMTP*|*|tcp_auth|*|* $Y$K$4

```

The **\$**(metacharacter in a **FROM_ACCESS** specifies that an address should be read from the result string and used to replace the current overriding postmaster address. **\$)** has the same effect with the added constraint that the overriding postmaster address must not be set prior to invoking the mapping. This allows for specific postmaster addresses to be used with addresses in nonlocal domains - domain postmaster addresses by definition only work with locally defined domains. The override address is (currently) the last string read from the **FROM_ACCESS** result prior to reading any **\$N/\$F** failure result.

PORT_ACCESS Mapping Table

The Dispatcher is able to selectively accept or reject incoming connections based on IP address and port number. As the Dispatcher accepts incoming connections, it searches the **PORT_ACCESS** table using probes of the form:

```
TCP|server-address|server-port|client-address|client-port
```

The Dispatcher tries to match against all **PORT_ACCESS** mapping entries. If the result of the mapping contains **\$N** or **\$F**, the connection will be immediately closed. Any other result of the mapping indicates that the connection is to be accepted. **\$N** or **\$F** may optionally be followed by a rejection message. If present, the message will be sent back down the connection just prior to closure. Note that a CRLF terminator will be appended to the string before it is sent back down the connection.

The **tcp_smtp_server** and **tcp_lmtp_server** processes also probe the **PORT_ACCESS** table. Checks that should not be duplicated in both the **dispatcher** and SMTP or LMTP server processes should be qualified using **\$:** or **\$;** and the **A** or **S** flags. See "[PORT_ACCESS Input Flags](#)" for more information.

Note: The MMP does not make use of mapping tables. If you wish to reject SMTP connections from certain IP addresses and you are using the MMP, you must use the **TCPAccess** option. See "[To Configure Mail Access with MMP](#)."

To control SMTP connections using mapping tables, use the **INTERNAL_IP** mapping table (see "[Allowing SMTP Relaying for External Sites](#).")

The flag **\$<** followed by an optional string causes Messaging Server to send the string to syslog (UNIX) or to the event log (NT) if the mapping probe matches. The flag **\$>** followed by an optional string causes Messaging Server to send the string as to syslog (UNIX) or to the event log (NT) if access is rejected. If bit 1 of the **LOG_CONNECTION** MTA option is set and the **\$N** flag is set so that the connection is rejected, then also specifying the **\$T** flag will cause a "T" entry to be written to the connection log. If bit 4 of the **LOG_CONNECTION** MTA option is set, then

site-supplied text may be provided in the **PORT_ACCESS** entry to include in the "C" connection log entries. To specify such text, include two vertical bar characters in the right-hand side of the entry, followed by the desired text.

PORT_ACCESS is evaluated unconditionally as soon as the SMTP server thread starts, before the banner is sent.

The following flags and metacharacters apply to the **PORT_ACCESS** mapping table. The "Send Access and Mail Access Mapping Tables" and "FROM_ACCESS Mapping Table" have different sets of flags.

- [Table 13–8, "PORT_ACCESS Input Flags"](#)
- [Table 13–9, "PORT_ACCESS Output Flags"](#)
- [PORT_ACCESS Table Examples](#)

See the *Messaging Server Reference* for more details about the format and operation, and patterns, templates, and metacharacters which apply to all mapping tables.

Input flags are set by the process that calls the mapping table and tested for by using the **\$:** and **\$;** metacharacters in the template. The **\$:** will succeed if the flag is set or fail if it is not set. The **\$;** will succeed if the flag is not set or fail if it is set. [Table 13–8](#) shows the **\$.x** form to test if the flag is set. The **\$.x** form can also be used to test if condition is not true.

The **PORT_ACCESS** table is consulted by both the **dispatcher** and the **tcp_smtp_server** or **tcp_lmtp_server** processes. Mapping table entries that perform callouts to external functions that may incur some overhead should be done in one or the other, but not both.

Table 13–8 PORT_ACCESS Input Flags

Flag	Description
\$.A	Match only when the probe is performed by the Dispatcher.
\$.S	Match only when the probe is performed by an SMTP server or LMTP server.

See "To Use DNS Lookups Including RBL Checking for SMTP Relay Blocking" for examples of using **\$.A**, **\$.A;**, **\$.S**, and **\$.S;**.

Output flags are set by the template in the mapping table entry and consumed upon return to the calling program. The *Valid* column in the tables below indicates whether the flag is used by **dispatcher** (D), **tcp_smtp_server** (S), or **tcp_lmtp_server** (L). The [Table 13–8, "PORT_ACCESS Input Flags"](#) should be used to ensure that rules are only applied in the appropriate context. [Table 13–9](#) shows the **PORT_ACCESS** output flags.

Table 13–9 PORT_ACCESS Output Flags

Flag	Valid	Description
\$U	S	Enable channel (slave) debugging. If MM_DEBUG or OS_DEBUG are set, enable those as well. See the discussion on debugging channel master and slave programs in the <i>Messaging Server Reference</i> for more information.
\$V	S	Enable the MTA's private SMTP extensions XADR, XCIR, XGEN, and XSTA, overriding any SMTP server DISABLE_* channel options.

Table 13–9 (Cont.) PORT_ACCESS Output Flags

Flag	Valid	Description
\$/	DS	Set the "fast disconnect" flag for sessions that have not yet succeeded in starting a transaction; for such sessions, any subsequent disconnect is done with SO_LINGER enabled and a timeout of 0, which may clear slots quicker on intermediate firewalls and proxies.
\$Y	DSL	Allow access.
\$T	D	If bit 1 (value 2) of the LOG_CONNECTION MTA option is set and the \$N flag is set so that the connection is rejected, then \$T outputs the entire right hand side text in a "T" (or "X" ++) record. The T log entry will include the entire mapping result string (\$N and its string). In contrast, bit 4 of LOG_CONNECTION is a different effect: it will cause material after two vertical bars to be included in normal "C" (connection close) records.

Table 13–10 describes PORT_ACCESS flags with arguments. For more information see "Output Flag Argument Order." Note: Do **NOT** alphabetize this list.

Table 13–10 PORT_ACCESS Flags with Arguments

Flag	Valid	Description
\$<string	DSL	Send <i>string</i> to syslog if probe matches.
\$>string	DSL	Send <i>string</i> to syslog if access is rejected.
\$N string \$F string	DSL	Reject access with the optional error text <i>string</i> . \$F is a synonym for \$N
SASL Ruleset	S	Not used, but you must enter an empty value (double bar with no space, " ") if you want to use any of the following flags.
SASL Realm	S	Not used, but you must enter an empty value (double bar with no space, " ") if you want to use any of the following flags.
Application Info	SL	If the LOG_CONNECTION MTA option is set to bit 4 (value 16), PORT_ACCESS is allowed to add text to application information string. This is where the string can be specified. If it is not used, you must enter an empty value (double bar with no space, " ") if you want to use any of the flags below.
\$D delay	S	Specifies the number of centiseconds to delay before purging and sending the banner. A value of 0 disabled both the delay and purge. This value has the same semantics as the BANNER_PURGE_DELAY value. Note that any PORT_ACCESS mapping setting overrides the BANNER_PURGE_DELAY SMTP channel option. For details on using this anti-spam feature see the <i>Messaging Server Reference</i> .
\$S channel	S	Set <i>channel</i> as the source channel for this SMTP session.
Certificate Nickname	S	Comma-separated list of Server Certificates to send to client for TLS.

{+} To use multiple flags with arguments, separate the arguments with the vertical bar character, |, placing the arguments in the order listed in this table. For more information see "Output Flag Argument Order."

++ The T record discussed in the **\$T** flag above occurs when the PORT_ACCESS table entry is being evaluated by the **dispatcher** process. If it is evaluated by the **tcp_smtp_server** it will produce an X record instead.

PORT_ACCESS Table Examples

For example, the following mapping will only accept SMTP connections (to port 25, the normal SMTP port) from a single network, except for a particular host singled out for rejection without explanatory text:

```
PORT_ACCESS
```

```
TCP|*|25|192.123.10.70|* $N500
TCP|*|25|192.123.10.*|* $Y
TCP|*|25|*|* $N500$ Bzzzt$ thank$ you$ for$ playing.
```

You need to restart the Dispatcher after making any changes to the **PORT_ACCESS** mapping table so that the Dispatcher will see the changes. (If you are using a compiled MTA configuration, you will first need to recompile your configuration using the **imsimta cnbuild** command to get the change incorporated into the compiled configuration.)

The **PORT_ACCESS** mapping table is specifically intended for performing IP-based rejections. For more general control at the email address level, the **SEND_ACCESS** or **MAIL_ACCESS** mapping table, might be more appropriate.

Table 13–9, "PORT_ACCESS Output Flags" lists the order in which arguments are consumed from the template for certain flags and other arguments that can be returned from the **PORT_ACCESS** table.

```
PORT_ACCESS
```

```
*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
* $YEXTERNAL
```

In this default configuration, it appears as if the **\$Y** is being given a mysterious "EXTERNAL" argument. But **\$Y** does not take any arguments. That "EXTERNAL" is actually the value for the "SASL Ruleset" argument. If other argument-consuming flags were specified, you would use "|" to separate the arguments.

SASL Ruleset

If you set a ruleset, then **msconfig** options of the form **authoption-name** apply.

Note that the **PORT_ACCESS** mapping supplied by the initial configuration does set the ruleset to "EXTERNAL" for incoming connections that don't match in **INTERNAL_IP** - that way you can, in theory, configure different SASL options for the "internal" connections (in the "DEFAULT" ruleset) vs. the "external" connections (in the "EXTERNAL" ruleset).

SASL Realm

Realm is used in place of **base.defaultdomain** if the userid supplied for authentication is unqualified.

Application Info

If the **LOG_CONNECTION** MTA option is set to bit 4 (value 16), **PORT_ACCESS** is allowed to add text to application information string. This is where the string can be specified. If the option is enabled but you do not want to specify any additional app info in a rule, you must enter an empty value (double bar with no space, "||") if you want to use any of the later flags/arguments.

Certificate Nickname

```
PORT_ACCESS
```

```
*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
TCP|1.2.3.4|*|* $;A$Y||xyzzy
* $YEXTERNAL
```

In this example, if the server has multiple IP addresses, for connections from clients to the server's IP address 1.2.3.4, the `tcp_smtp_server` process will use the Certificate with the nickname `xyzzzy` for TLS instead of the default certificate. The " | | " supplies null values for "SASL Ruleset" and "SASL Realm." If the `LOG_CONNECTION` MTA option is set to bit 4 (value 16), an additional " | " would be needed, for example:

```
PORT_ACCESS

*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
TCP|1.2.3.4|*|*|* $;A$Y|||xyzzzy
* $YEXTERNAL
```

IP_ACCESS Mapping Table

The `IP_ACCESS` Mapping Table can be used to do a last moment check on the IP address to which the MTA is about to connect. The connection attempt can then be aborted or redirected. This can be useful under certain special circumstances, for example, security concerns about a destination IP address to which should never be connected, or where it is wished to avoid connecting to known-to-be-bogus destination IP addresses (for example, 127.0.0.1), or where you wish to attempt to fail over to another destination IP address similar to a `lastresort` channel option effect (see the discussion on specifying a last resort host for delivery in the *Messaging Server Reference*).

This access mapping is consulted during SMTP client operations just prior to attempting to open connections to a remote server. The mapping probe has the following format:

```
source-channel | address-current | address-count | ip-current | hostname
```

source-channel is the channel from which the message is being dequeued. **address-count** is the total number of IP addresses for the remote server. **address-current** is the index of the current IP address being tried. **ip-current** is the current IP address. **hostname** is the symbolic name of the remote server. [Table 13–11](#) shows the `IP_ACCESS` flags.

See the *Messaging Server Reference* for more details about the format and operation, and patterns, templates, and metacharacters which apply to all mapping tables.

Table 13–11 *IP_ACCESS Mapping Table Flags*

Flag	Description
\$N	Immediately reject the message with an "invalid host/domain error." Any supplied text will be logged as the reason for rejection but will not be included in the DSN.
\$I	Skip the current IP without attempting to connect.
\$A	Replace the current IP address with the mapping result.

When Access Controls Are Applied

Messaging Server checks access control mappings as early as possible. Exactly when this happens depends upon the email protocol in use - when the information that must be checked becomes available.

For the SMTP protocol, a `FROM_ACCESS` rejection occurs in response to the `MAIL FROM:` command, before the sending side can send the recipient information or the

message data. A **SEND_ACCESS** or **MAIL_ACCESS** rejection occurs in response to the **RCPT TO:** command, before the sending side gets to send the message data. If an SMTP message is rejected, Messaging Server never accepts or sees the message data, thus minimizing the overhead of performing such rejections.

If multiple access control mapping tables exist, Messaging Server checks them all. That is, a **FROM_ACCESS**, a **SEND_ACCESS**, an **ORIG_SEND_ACCESS**, a **MAIL_ACCESS**, and **ORIG_MAIL_ACCESS** mapping tables may all be in effect.

PORT_ACCESS is called from dispatcher as soon as it accepts the incoming TCP connection. It is also called from **tcp_smtp_server** when any of the **maysaslserver** or **mustsaslserver** channel options are present on the source channel. (See the discussion on SMTP authentication and SASL in the *Messaging Server Reference*.)

FROM_ACCESS is used by the **tcp_smtp_server** when processing the **MAIL FROM SMTP** command.

SEND_ACCESS and **ORIG_SEND_ACCESS** tables are used by the **tcp_smtp_server** when processing the **RCPT TO SMTP** command.

MAIL_ACCESS and **ORIG_MAIL_ACCESS** tables are used by the **tcp_smtp_server** when processing the **RCPT TO SMTP** command.

To Test Access Control Mappings

The **imsimta test -rewrite** utility - particularly with the **-from**, **-source_channel**, **-sender** and **-destination_channel** options - can be useful in testing access control mappings. See "imsimta test" for details. The example below shows a sample **SEND_ACCESS** mapping table and the resulting probe.

MAPPING TABLE:

ORIG_SEND_ACCESS

```
tcp_local|friendly@example.com|*|User@example.org    $Y
tcp_local|unwelcome@example.edu|*|User@example.org  $NGo$ away!
```

PROBE:

```
$ imsimta test -rewrite -from="friendly@example.com" -source=tcp_local
User@example.org
...
Submitted address list:
  ims-ms
    user@ims-ms-daemon (orig User@example.org, inter User@example.org, host
ims-ms-daemon)
*NOTIFY-FAILURES* *NOTIFY-DELAYS*
```

Submitted notifications list:

```
$ imsimta test -rewrite -from="unwelcome@example.edu" -source=tcp_local
User@example.org
...
Submitted address list:
Address list error -- 5.7.1 Go away!: User@example.org

Submitted notifications list:
```

To Limit Specified IP Address Connections to the MTA

To limit how often a particular IP address can connect to the MTA, see "[Using and Configuring MeterMaid for Access Control](#)." Limiting connections by particular IP addresses can be useful for preventing excessive connections used in denial-of-service attacks. In the past, this function was performed using the shared library, **conn_throttle.so** in the Port Access mapping table. No new enhancements are planned for **conn_throttle.so** and MeterMaid is its more effective replacement.

conn_throttle.so is a shared library used in a **PORT_ACCESS** mapping table to limit MTA connections made too frequently from particular IP addresses. All configuration options are specified as options to the connection throttle shared library as follows:

```
$[MessagingServer_home/lib/conn_throttle.so,throttle,IP-address,max-rate]
```

IP-address is the dotted-decimal address of the remote system. *max-rate* is the connections per minute that shall be the enforced maximum rate for this IP-address.

The routine name **throttle_p** may be used instead of **throttle** for a penalizing version of the routine. **throttle_p** will deny connections in the future if they've connected too many times in the past. If the maximum rate is 100, and 250 connections have been attempted in the past minute, not only will the remote site be blocked after the first 100 connections in that minute, but they'll also be blocked during the second minute. In other words, after each minute, *max-rate* is deducted from the total number of connections attempted and the remote system is blocked as long as the total number of connections is greater than the maximum rate.

If the IP-address specified has not exceeded the maximum connections per minute rate, the shared library callout will fail.

If the rate has been exceeded, the callout will succeed, but will return nothing. This is done in a **\$C/\$E** combination as in the example:

```
PORT_ACCESS
    TCP|*|25|*|* \
    $C$[_MessagingServer_home_/lib/conn_throttle.so,throttle,$1,10] \
    $N421$ Connection$ not$ accepted$ at$ this$ time$E
```

Where,

\$C continues the mapping process starting with the next table entry; uses the output string of this entry as the new input string for the mapping process.

\$[MessagingServer_home/lib/conn_throttle.so,throttle,\$1,10] is the library call with **throttle** as the library routine, **\$1** as the server IP Address, and 10 the connections per minute threshold.

\$N421\$ Connection\$ not\$ accepted\$ at\$ this\$ time rejects access and returns the 421 SMTP code (transient negative completion) along with the message "Connection not accepted at this time".

\$E ends the mapping process now. It uses the output string from this entry as the final result of the mapping process.

To Add SMTP Relaying

Messaging Server is, by default, configured to block attempted SMTP relays. That is, it rejects attempted message submissions to external addresses from unauthenticated external sources (external systems are any other system than the host on which the

server itself resides). This default configuration is quite aggressive in blocking SMTP relaying in that it considers all other systems to be external systems.

IMAP and POP clients that attempt to submit messages via the Messaging Server system's SMTP server destined for external addresses, and who do not authenticate using SMTP AUTH (SASL), will find their submission attempts rejected. Thus, you will likely want to modify your configuration so that it recognizes your own internal systems and subnets from which relaying should always be accepted.

Which systems and subnets are recognized as internal is normally controlled by the **INTERNAL_IP** mapping table, which is located in the *MessagingServer_homedir/config/mappings* directory.

For instance, on a Messaging Server whose IP address is 123.45.67.89, the default **INTERNAL_IP** mapping table would appear as follows:

```
INTERNAL_IP

$(123.45.67.89/32)  $Y
127.0.0.1          $Y
*                  $N
```

Here the initial entry, using the **\$(IP-pattern/significant-prefix-bits)** syntax, is specifying that any IP address that matches all 32 bits of 123.45.67.89 should match and be considered internal. The second entry recognizes the loopback IP address 127.0.0.1 as internal. The final entry specifies that all other IP addresses should not be considered internal. All entries must be preceded by at least one space.

You add additional entries by specifying additional IP addresses or subnets before the final \$N entry. These entries must specify an IP address or subnet (using the **\$(.../...)** syntax to specify a subnet) on the left side and \$Y on the right side. Or you may modify the existing **\$(.../...)** entry to accept a more general subnet.

For instance, if this same sample site has a class-C network, that is, it owns all of the 123.45.67.0 subnet, then the site would want to modify the initial entry by changing the number of bits used in matching the address. In the mapping table below, we change from 32 bits to 24 bits. This allows all clients on the class-C network to relay mail through this SMTP relay server.

```
INTERNAL_IP

$(123.45.67.89/24)  $Y
127.0.0.1          $Y
*                  $N
```

Or if the site owns only those IP addresses in the range 123.45.67.80-123.45.67.99, then the site would want to use:

```
INTERNAL_IP

! Match IP addresses in the range 123.45.67.80-123.45.67.95
$(123.45.67.80/28) $Y
! Match IP addresses in the range 123.45.67.96-123.45.67.99
$(123.45.67.96/30) $Y
127.0.0.1 $Y
* $N
```

Note that the **imsimta test -match** utility can be useful for checking whether an IP address matches a particular **\$(.../...)** test condition. The **imsimta test -mapping** utility can be more generally useful in checking that your **INTERNAL_IP** mapping table returns the desired results for various IP address inputs.

After modifying your **INTERNAL_IP** mapping table, be sure to issue the **imsimta restart** command (if you are not running with a compiled configuration) or an **imsimta cnbuild** followed by an **imsimta restart smtp** (if you are running with a compiled configuration) so that the changes take effect.

Further information on the mapping file and general mapping table format, as well as information on **imsimta** command-line utilities, can be found in the *Messaging Server Reference*.

Allowing SMTP Relaying for External Sites

All internal IP addresses should be added to the **INTERNAL_IP** mapping table as discussed above. If you have friendly or companion systems/sites from which you wish to allow SMTP relaying, the simplest approach is to include them along with your true internal IP addresses in your **INTERNAL_IP** mapping table.

If you don't wish to consider these as true internal systems/sites, (for instance, if for logging or other control purposes you wish to distinguish between *true internal systems* versus the *friendly non-internal systems with relay privileges*), there are other ways to configure the system.

One approach is to set up a special channel for receiving messages from such friendly systems. Do this by creating a *tcp_friendly* channel akin to your existing *tcp_internal* channel with official host name *tcp_friendly-daemon*, and a **FRIENDLY_IP** mapping table akin to your **INTERNAL_IP** mapping table that lists the friendly system IP addresses. Then right after the current rewrite rule:

```
! Do mapping lookup for internal IP addresses
[]   $E$R${INTERNAL_IP,$L}$U%[$L]@tcp_intranet-daemon
```

add a new rewrite rule:

```
! Do mapping lookup for "friendly", non-internal IP addresses
[]   $E$R${FRIENDLY_IP,$L}$U%[$L]@tcp_friendly-daemon
```

An alternate approach is to add to your **ORIG_SEND_ACCESS** mapping table above the final **\$N** entry, new entries of the form:

```
tcp_local|*@example.com|tcp_local|*   $Y
```

where **example.com** is the name of a friendly domain, and to add an **ORIG_MAIL_ACCESS** mapping table of the form:

```
ORIG_MAIL_ACCESS

TCP|*|25|$(match-example.com-IP-addresses)|*|SMTP*|MAIL| \
tcp_local|*@example.com|tcp_local|*   $Y
TCP|*|*|*|SMTP*|MAIL|tcp_local|*|tcp_local|* $N
```

where the **\$(...)** IP address syntax is the same syntax described in the previous section. The **ORIG_SEND_ACCESS** check will succeed as long as the address is ok, so we can go ahead and also do the **ORIG_MAIL_ACCESS** check which is more stringent and will only succeed if the IP address also corresponds to an example.com IP address.

Configuring SMTP Relay Blocking

You can use access control mappings to prevent people from relaying SMTP mail through your Messaging Server system. For example, you can prevent people from using your mail system to relay junk mail to hundreds or thousands of Internet mailboxes.

By default, Messaging Server prevents all SMTP relaying activity, including relaying by local POP and IMAP users.

Blocking unauthorized relaying while allowing it for legitimate local users requires configuring Messaging Server to know how to distinguish between the two classes of users. For example, local users using POP or IMAP depend upon Messaging Server to act as an SMTP relay.

To prevent SMTP relay, you must be able to:

- [Differentiate Between Internal and External Mail](#)
- [Differentiate Authenticated Users' Mail](#)
- [Prevent Mail Relay](#)

To enable SMTP relay by internal hosts and clients, you must add your "internal" IP addresses or subnets to the `INTERNAL_IP` mapping table.

Differentiate Between Internal and External Mail

To block mail relaying activities, the MTA must first be able to differentiate between internal mail originated at your site and external mail originated out on the Internet and passing through your system back out to the Internet. The former class of mail you want to permit; the latter class you want to block. This differentiation is achieved using the `switchchannel` option on your inbound SMTP channel, usually the `tcp_local` channel, and is set by default.

The `switchchannel` channel option works by causing the SMTP server to look at the actual IP address associated with the incoming SMTP connection. Messaging Server uses that IP address, in conjunction with your rewrite rules, to differentiate between an SMTP connection originated within your domain and a connection from outside of your domain. This information can then be used to segregate the message traffic between internal and external traffic.

The MTA configuration described below is setup by default so that the server can differentiate between your internal and external message traffic.

- In the configuration file, immediately before the local channel, is a `defaults` channel with the `noswitchchannel` option:

```
defaults notices 1 2 4 7 noswitchchannel immnonurgent maxjobs 7 defaulthost
example.com example.com
```

- The incoming TCP/IP channel specifies the `switchchannel` and `remotehost` options; for example:

```
tcp_local smtp single_sys mx switchchannel remotehost
tcp-daemon
```

- After the incoming TCP/IP channel definition, is a similar channel with a different name which specifies the `allowswitchchannel` option; for example:

```
tcp_intranet smtp single_sys mx allowswitchchannel
tcp_intranet-daemon
```

With the above configuration settings, SMTP mail generated within your domain will come in via the `tcp_intranet` channel. All other SMTP mail will come in via the `tcp_local` channel. Mail is distinguished between internal and external based upon which channel it comes in on.

How does this work? The key is the **switchchannel** option. The option is applied to the **tcp_local** channel. When a message comes in your SMTP server, that option causes the server to look at the source IP address associated with the incoming connection. The server attempts a reverse-pointing envelope rewrite of the literal IP address of the incoming connection, looking for an associated channel. If the source IP address matches an IP address or subnet in your **INTERNAL_IP** mapping table, the rewrite rule which calls out to that mapping table causes the address to rewrite to the **tcp_intranet** channel.

Since the **tcp_intranet** channel is marked with the **allowswitchchannel** option, the message is switched to the **tcp_intranet** channel and comes in on that channel. If the message comes in from a system whose IP address is not in the **INTERNAL_IP** mapping table, the reverse-pointing envelope rewrite will either rewrite to the **tcp_local** or, perhaps to some other channel. However, it will not rewrite to the **tcp_intranet** channel and since all other channels are marked **noswitchchannel** by default, the message will not switch to another channel and will remain with the **tcp_local** channel.

Note: Any mapping table or conversion file entries which use the string "**tcp_local**" may need to be changed to either "**tcp_***" or "**tcp_intranet**" depending upon the usage.

Differentiate Authenticated Users' Mail

Your site might have "local" client users who are not part of your physical network. When these users submit mail, the message submissions come in from an external IP address, for instance, arbitrary Internet Service Providers. If your users use mail clients that can perform SASL authentication, then their authenticated connections can be distinguished from arbitrary other external connections. The authenticated submissions you can then permit, while denying non-authenticated relay submission attempts. Differentiating between authenticated and non-authenticated connections is achieved using the **saslswitchchannel** option on your inbound SMTP channel, usually the **tcp_local** channel.

The **saslswitchchannel** option takes an argument specifying the channel to switch to; if an SMTP sender succeeds in authenticating, then their submitted messages are considered to come in the specified switched to channel.

To Add Distinguishing Authenticated Submissions

1. In your configuration file, add a new TCP/IP channel definition with a distinct name; for example:

```
tcp_auth smtp mustsaslsrvr noswitchchannel
tcp_auth-daemon
```

This channel should not allow regular channel switching (that is, it should have **noswitchchannel** on it either explicitly or implied by a prior defaults line). This channel should have **mustsaslsrvr** on it.

2. Modify your **tcp_local** channel by adding **maysaslsrvr** and **saslswitchchannel tcp_auth**, as shown in the following example:

```
tcp_local smtp mx single_sys maysaslsrvr saslswitchchannel tcp_auth
switchchannel
tcp-daemon
```

With this configuration, SMTP mail sent by users who can authenticate with a local password will now come in the `tcp_auth` channel. Unauthenticated SMTP mail sent from internal hosts will still come in `tcp_internal`. All other SMTP mail will come in `tcp_local`.

Prevent Mail Relay

Now to the point of this example: preventing unauthorized people from relaying SMTP mail through your system. First, keep in mind that you want to allow local users to relay SMTP mail. For instance, POP and IMAP users rely upon using Messaging Server to send their mail. Note that local users may either be physically local, in which case their messages come in from an internal IP address, or may be physically remote but able to authenticate themselves as local users.

You want to prevent random people out on the Internet from using your server as a relay. With the configuration described in the following sections, you can differentiate between these classes of users and block the correct class. Specifically, you want to block mail from coming in your `tcp_local` channel and going back out that same channel. To that end, an `ORIG_SEND_ACCESS` mapping table is used.

An `ORIG_SEND_ACCESS` mapping table may be used to block traffic based upon the source and destination channel. In this case, traffic from and back to the `tcp_local` channel is to be blocked. This is realized with the following `ORIG_SEND_ACCESS` mapping table:

```
ORIG_SEND_ACCESS
```

```
tcp_local|*|tcp_local|* $N$D30|Relaying$ not$ allowed
```

In this example, the entry states that messages cannot come in the `tcp_local` channel and go right back out it. That is, this entry disallows external mail from coming in your SMTP server and being relayed right back out to the Internet.

An `ORIG_SEND_ACCESS` mapping table is used rather than a `SEND_ACCESS` mapping table so that the blocking will not apply to addresses that originally match the `ims-ms` channel (but which may expand via an alias or mailing list definition back to an external address). With a `SEND_ACCESS` mapping table one would have to go to extra lengths to allow outsiders to send to mailing lists that expand back out to external users, or to send to users who forward their messages back out to external addresses.

To Use DNS Lookups Including RBL Checking for SMTP Relay Blocking

In the Messaging Server, there are a number of different ways to ensure that all mail accepted for delivery or forwarding comes from an address with a valid DNS name. The simplest way is to put the `mailfromdnsverify` channel option on the `tcp_local` channel.

Messaging Server also provides the `dns_verify` program which allows you to ensure that all mail accepted for delivery or forwarding comes from an address with a valid DNS name using the following rule in `ORIG_MAIL_ACCESS`:

```
ORIG_MAIL_ACCESS
```

```
TCP|*|*|*|SMTP*|MAIL|*|*@*|*|* \
$_MessagingServer_home_/lib/dns_verify.so,\
dns_verify,$7|$$y|$$NInvalid$ host:$ $$7$ -$ %e]
```

The line breaks in the above example are syntactically significant in such mapping entries. The backslash character is a way of legally continuing on to the next line.

The **dns_verify** image can also be used to check incoming connections against things like the RBL (Realtime Blackhole List), MAPS (Mail Abuse Prevention System, DUL (Dial-up User List), or ORBS (Open Relay Behavior-modification System) lists as another attempt to protect against UBE. As with the new **mailfromdnsverify** option, there's also a separate "simpler to configure" approach one can use for such checks rather than doing the **dns_verify** callout. The simpler approach is to use the **DNS_VERIFY_DOMAIN** option in the dispatcher configuration. For example, for the smtp service, set instances of the option to the various lists you want to check against:

```
msconfig
msconfig# set dispatcher.service:smtp.dns_verify_domain sbl-xbl.spamhaus.org
msconfig# set dispatcher.service:smtp.dns_verify_domain list.dsbl.org
msconfig# write
```

In this case, messages are rejected at the SMTP level, that is, the messages are rejected during the SMTP dialogue and thus never sent to the MTA. The disadvantage of this simpler approach is that it does the checks for all normal incoming SMTP messages including those from internal users. This is less efficient and potentially problematic if your Internet connectivity goes down. An alternative is to call out to **dns_verify** from a **PORT_ACCESS** mapping table or **ORIG_MAIL_ACCESS** mapping table. In the **PORT_ACCESS** mapping table, you can have an initial entry or entries that don't check for local internal IP addresses or message submitters and a later entry that does the desired check for everyone else. Or, in an **ORIG_MAIL_ACCESS** mapping table, if you only apply the check on messages coming in the tcp_local channel then you're skipping it for messages coming from your internal systems/clients. Examples using the entry points to **dns_verify** are shown below.

PORT_ACCESS

```
! Allow internal connections in unconditionally
*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
! Check other connections against RBL list
TCP|*|25|*|* \
$C$[_MessagingServer_home_/lib/dns_verify.so,dns_verify_domain_port,$1,\
dnsblock.example.com.,Your$ host$ ($1)$ found$ on$ dnsblock$ list]$E
* $YEXTERNAL
```

ORIG_MAIL_ACCESS

```
TCP|*|25|*|*|SMTP*|*|tcp_local|*@*|*|* \
$C$[_MessagingServer_home_/lib/dns_verify.so,\
dns_verify_domain,$1,sbl-xbl.spamhaus.org.]$E
```

For more information see the discussion on performance tuning DNS realtime blocklists (RBL) lookups in the *Messaging Server Installation and Configuration Guide*.

The **PORT_ACCESS** table is probed both by the dispatcher, when accepting connections, and by the **tcp_smtp_server** process under certain circumstances.

The **tcp_smtp_server** processes always check the **PORT_ACCESS** table for channels marked **maysaslsrver** or **mustsaslsrver**, and they will do it for all channels if bit 4 (value 16) of the MTA **LOG_CONNECTION** option is set.

So if either of those are true, the customer needs to be aware that his **PORT_ACCESS** table is being processed twice for every connection. This may be a trivial overhead except that callouts to things like DNS RBLs or LDAP lookups can be relatively expensive in terms of their impact on SMTP server response time. For this reason, you want to avoid doing them any more than necessary. That is one reason the callout in the example above is coded after the **INTERNAL_IP** lookup. If the SMTP client is in

your **INTERNAL_IP** table, then you allow the connection without doing the RBL lookup. Reversing that order would mean your local clients would be delayed by RBL lookups. To prevent doing this twice, add a check for one of the flags set by dispatcher or **tcp_smtp_server** so that you only process that table entry when one of those is set or not. For example, add **\$.A** to the entry:

```
TCP|*|25|*|* \
$C$:$A$_MessagingServer_home_/lib/dns_verify.so,dns_verify_domain_port,$1,\
dnsblock.example.com.,Your$ host$ ($1)$ found$ on$ dnsblock$ list]$E
* $YEXTERNAL
```

We added the **\$.A** after the **\$.C** so the table processing will continue down the table if this does not match. The **\$.A** specifies that this entry be processed only if it is being done by the dispatcher, that is, not when done by **tcp_smtp_server**. Alternatively, if you wanted to cause it to be done only in **tcp_smtp_server**, use **\$.S** instead:

```
TCP|*|25|*|* \
$C$:$S$_MessagingServer_home_/lib/dns_verify.so,dns_verify_domain_port,$1,\
dnsblock.example.com.,Your$ host$ ($1)$ found$ on$ dnsblock$ list]$E
* $YEXTERNAL
```

The negative check would be **\$.A** to check that **A** is not set, or **\$.S** to check that **S** is not set.

Support for DNS-based Databases

The **dns_verify** program supports DNS-based databases used to determine incoming SMTP connections that might send unsolicited bulk mail. Some of the publicly available DNS databases do not contain TXT records that are typically used for this purpose. Instead, they only contain **A** records.

In a typical setup, the **TXT** record found in the DNS for a particular IP address contains an error message suitable to return to the SMTP client when refusing a message.

dns_verify supports an option that specifies a default text that is used in the event that no **TXT** record is available. For example, the following **PORT_ACCESS** mapping table shows how to enable this option:

```
PORT_ACCESS
*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
TCP|*|25|*|* \
$C$[_MessagingServer_home_/lib/dns_verify.so \
,dns_verify_domain_port,$1,dnsblock.example.com,Your$ host$ ($1)$ \
found$ on$ dnsblock$ list]$E
* $YEXTERNAL
```

In this example, if the remote system is found in a query in the domain **dnsblock.example.com**, but no **TXT** record is available, then the following message is returned, "Your host a.b.c.d found on dnsblock list."

Handling Large Numbers of Access Entries

Sites that use very large numbers of entries in mapping tables should consider organizing their mapping tables to have a few general wildcarded entries that call out to the general text database for the specific lookups. It is much more efficient to have a few mapping table entries calling out to the general text database for specific lookups than to have huge numbers of entries directly in the mapping table.

One case in particular is that some sites like to have per user controls on who can send and receive Internet email. Such controls are conveniently implemented using an access mapping table such as **ORIG_SEND_ACCESS**. For such uses, efficiency and performance can be greatly improved by storing the bulk of the specific information (that is, specific addresses) in the general text database with mapping table entries structured to call out appropriately to the general text database.

For example, consider the **ORIG_SEND_ACCESS** mapping table shown below.

```
ORIG_SEND_ACCESS

! Users allowed to send to Internet
!
*|adam@example.com|tcp_local|*    $Y
*|betty@example.com|tcp_local|*    $Y
! ...etc...
!
! Users not allowed to send to Internet
!
*|norman@example.com|tcp_local|*    $NInternet$ access$ not$ permitted
*|opal@example.com|tcp_local|*    $NInternet$ access$ not$ permitted
! ...etc...
!
! Users allowed to receive from the Internet
!
tcp_*|*|*|adam@example.com        $Y
tcp_*|*|*|betty@example.com        $Y
! ...etc...
!
! Users not allowed to receive from the Internet
!
tcp_*|*|*|norman@example.com        $NInternet$ e-mail$ not$ accepted
tcp_*|*|*|opal@example.com          $NInternet$ e-mail$ not$ accepted
! ...etc...
```

Rather than using such a mapping table with each user individually entered into the table, a more efficient setup (much more efficient if hundreds or thousands of user entries are involved) is shown in the example below, which shows sample source text file for a general database and a sample **ORIG_SEND_ACCESS** mapping table. For more information, see the discussion on MTA text databases in the *Messaging Server Reference*.

DATABASE ENTRIES

```
SEND|adam@domain.com    $Y
SEND|betty@domain.com   $Y
! ...etc...
SEND|norman@domain.com  $NInternet$ access$ not$ permitted
SEND|opal@domain.com    $NInternet$ access$ not$ permitted
! ...etc...
RECV|adam@domain.com    $Y
RECV|betty@domain.com   $Y
! ...etc...
RECV|norman@domain.com  $NInternet$ e-mail$ not$ accepted
RECV|opal@domain.com    $NInternet$ e-mail$ not$ accepted
```

MAPPING TABLE

```
ORIG_SEND_ACCESS

! Check if may send to Internet
```

```

!
*|*|*|tcp_local      $$${SEND|$1}$$E
!
! Check if may receive from Internet
!
tcp_*|*|*|*          $$${RECV|$3}$$E

```

In this example, the use of the arbitrary strings **SEND|** and **RECV|** in the general database left-hand sides (and hence in the general database probes generated by the mapping table) provides a way to distinguish between the two sorts of probes being made. The wrapping of the general text database probes with the **\$C** and **\$E** flags, as shown, is typical of mapping table callouts to the general text database.

The above example showed a case of simple mapping table probes getting checked against general text database entries. Mapping tables with much more complex probes can also benefit from use of the general text database.

Controlling the Envelope From: Address in Mappings Strings and Mailing List Named Options and LDAP Attributes

The `use_auth_return`, `use_canonical_return`, and `use_orig_return` MTA options control which envelope address to use in certain mapping tables or mailing list named options (if using the aliases file for mailing lists) or attributes (if using LDAP mailing lists).

For more information on these options, see the discussion on return address type used in checks MTA options in the *Messaging Server Reference*.

PART 2. MAILBOX FILTERS

Mailbox filters, also called Sieve filters, filter messages containing specified strings found in the message headers and apply specified actions to these mail message. Administrators can filter mail streams going to a user, through a channel, or through an MTA. Messaging Server filters are stored on the server and evaluated by the server, hence, they are sometimes called server-side rules (SSR).

This section contains the following topics:

- [Sieve Filter Support](#)
- [Sieve Filtering Overview](#)
- [To Create User-level Filters](#)
- [To Create Channel-level Filters](#)
- [Managing Sieve Scripts](#)
- [To Create MTA-Wide Filters](#)
- [To Debug User-level Filters](#)

Sieve Filter Support

Messaging Server filters are based on the Sieve filtering language, Draft 9 of the Sieve Internet Draft. See RFC3028 for more information about Sieve syntax and semantics. In addition, Messaging Server also supports the following Sieve extensions:

- **jettison**. Similar to **discard** in that it causes messages to be silently discarded, but unlike **discard**, which does nothing but cancel the implicit keep, **jettison** forces a

discard to be performed. The behavioral difference is only relevant when multiple Sieve filters are involved. For example, a system level discard can be overridden by a user Sieve filter explicitly specifying **keep**, whereas a system level **jettison** will override anything done by a user Sieve.

- **Head-of-household Sieve filters.** Provides a means by which one user can specify a Sieve filter for another user. Uses two LDAP attributes in a user entry controlled by these MTA options:

- **LDAP_PARENTAL_CONTROLS** - Specifies an attribute containing a string value of either **Yes** or **No**. **Yes** means a head of household Sieve is to be applied to this entry, **No** means no such Sieve is to be applied. No default.
- **LDAP_FILTER_REFERENCE** - Specifies an attribute containing a DN pointing to a directory entry where the head of household Sieve can be found. No default.

The entry containing the head of household Sieve must contain two attributes specified by the following MTA options:

- **LDAP_HOH_FILTER** - Specifies an attribute containing the head of household Sieve. The value of this option defaults to **mailSieveRuleSource**.
- **LDAP_HOH_OWNER** - Specifies an attribute containing the email address of the owner of the head of household. The value of this option defaults to **mail**. Both of these attributes must be present for the head of household Sieve to work.
- Sieve redirect can now add three header fields:

```
resent-date: _date-of-resend-operation_
resent-to: _address-specified-in-redirect_
resent-from: _address-of-sieve-owner_
```

The new **:resent** and **:noresent** arguments to redirect can be used to control whether or not these fields are added. If neither argument is specific the system default is used. The system default is controlled by the new **SIEVE_REDIRECT_ADD_RESENT** MTA option. Setting the option to 1 causes these fields to be generated unless **:noresent** used. A setting of 0 causes the fields to be generated only if **:resent** is used. The option defaults to 1, which means the fields are generated by default for regular redirects.

- Sieve redirect has been enhanced with three new arguments:
 - :resetmailfrom** - Reset the envelope **FROM:** address to that of the current Sieve owner.
 - :keepmailfrom** - Preserve the envelope **FROM:** address from the original message.
 - :notify** - Specify a new set of notification flags for the redirected message. A single option is required giving a list of notification flags. The same set of flags accepted by the **NOTIFY** option of the DSN SMTP extension are accepted here: **SUCCESS**, **FAILURE**, **DELAY** and **NEVER**. Note that these flags are specified as a Sieve list, for example:

```
redirect :notify ["SUCCESS", "FAILURE"] "foo@example.com";
```

The default if **:notify** isn't specified as the normal SMTP default of **FAILURE**, **DELAY**. **:keepmailfrom** is the default unless **:notify** is specified, in which case the default switches to **:resetmailfrom**. The one additional exception is that specification of the **SUCCESS** flag forces the use of **:resetmailfrom** unconditionally.

Sieve Filtering Overview

A Sieve filter consists of one or more conditional actions to apply to a mail message depending upon a string found in the message header. As an administrator, you can create channel-level filters and MTA-wide filters to prevent delivery of unwanted mail. Users can create per-user filters for their own mailboxes by using the mail filter interface in either Communications Express or Convergence.

The server applies filters in the following priority:

1. **User-level filters** If a personal mailbox filter explicitly accepts or rejects a message, then filter processing for that message finishes. But if the recipient user had no mailbox filter, or if the user's mailbox filter did not explicitly apply to the message in question, Messaging Server next applies the channel-level filter.
2. **Channel-level filter** If the channel-level filter explicitly accepts or rejects a message, then filter processing for that message finishes. Otherwise, Messaging Server next applies the MTA-wide filter, if there is one.
3. **MTA-wide filter**

By default, each user has no mailbox filter. When a user uses either the Communications Express or Convergence interfaces to create one or more filters, then their filters are stored in the Directory and retrieved by the MTA during the directory synchronization process.

To Create User-level Filters

Per-user mail filters apply to messages destined for a particular user's mailbox. Per-user mail filters can only be created by using either the Communications Express or Convergence interfaces.

To Create Channel-level Filters

Channel-level filters apply to each message enqueued to a channel. A typical use for this type of filter is to block messages going through a specific channel. [Table 13–12](#) shows the **filter** channel option URL-pattern substitution tags.

Table 13–12 *filter Channel Option URL-pattern Substitution Tags (Case-insensitive)*

Tag	Meaning
*	Perform group expansion.
**	Expand the attribute mailForwardingAddress . This can be a multivalued attribute resulting in several delivery addresses being produced.
\$\$	Substitute in the \$ character.
\$	Force subsequent text to lower case.
\$\$	Force subsequent text to upper case.
\$_	Perform no case conversion on subsequent text.
\$~	Substitute in the file path for the home directory associated with the local part of the address.
\$1S	As \$\$, but if no subaddress is available just insert nothing.
\$2S	As \$\$, but if no subaddress is available insert nothing and delete the preceding character.

Table 13–12 (Cont.) filter Channel Option URL-pattern Substitution Tags

Tag	Meaning
\$3S	As \$S, but if no subaddress is available insert nothing and ignore the following character.
\$A	Substitute in the address, local-part@host.domain.
\$D	Substitute in host.domain.
\$E	Insert the value of the second spare attribute, LDAP_SPARE_1.
\$F	Insert the name of the delivery file (mailDeliveryFileURL attribute).
\$G	Insert the value of the second spare attribute, LDAP_SPARE_2.
\$H	Substitute in host.
\$I	Insert the hosted domain (part of UID to the right of the separator specified by domainUidSeparator). Fail if no hosted domain is available.
\$II	As \$I, but if no hosted domain is available just insert nothing.
\$2I	As \$I, but if no hosted domain is available insert nothing and delete the preceding character.
\$3I	As \$I, but if no hosted domain is available insert nothing and ignore the following character.
\$L	Substitute in local-part.
\$M	Insert the UID, stripped of any hosted domain.
\$P	Insert the method name (mailProgramDeliveryInfo attribute).
\$S	Insert the subaddress associated with the current address. The subaddress is that part of the user part of the original address after the subaddress separator, usually {+}, but can be specified by the MTA option SUBADDRESS_CHAR . Fail if no subaddress is given
\$U	Insert the mailbox part of the current address. This is either the whole of the address to the left of the @ sign, or that part of the left hand side of the address before the subaddress separator, {+}.

To Create a Channel-level Filter

1. Write the filter using Sieve.
2. Store the filter in a file in the following directory:*MessagingServer_home/config/file.filter* The file must be world readable and owned by the MTA's uid.
3. Include the following in the channel configuration:

```
destinationfilter file:IMTA_TABLE:<file>.filter
```
4. Recompile the configuration and restart the Dispatcher. Note that changes to the filter file do not require a recompile or restart of the Dispatcher. The **destinationfilter** channel option enables message filtering on messages enqueued to the channel to which it is applied. The **sourcefilter** channel option enables message filtering on messages enqueued *by* (from) the channel to which it is

applied. These options each have one required option which specifies the path to the corresponding channel filter file associated with the channel.

The syntax for the **destinationfilter** channel option is:

destinationfilter*URL-pattern*

The syntax for the **sourcefilter** channel option is: **sourcefilter***URL-pattern*

where *URL-pattern* is a URL specifying the path to the filter file for the channel in question. In the following example, *channel-name* is the name of the channel.

```
destinationfilter file:///usr/tmp/filters/<channel-name>.filter
```

The **filter** channel option enables message filtering on the channels to which it is applied. The option has one required option which specifies the path to the filter files associated with each envelope recipient who receives mail via the channel.

The syntax for the **filter** channel option is:

filter*URL-pattern*

URL-pattern is a URL that, after processing special substitution sequences, yields the path to the filter file for a given recipient address. *URL-pattern* can contain special substitution sequences that, when encountered, are replaced with strings derived from the recipient address, **local-part@host.domain** in question. These substitution sequences are shown in Table 13–12, "filter Channel Option [URL-pattern Substitution Tags \(Case-insensitive\)](#)". The **fileinto** option specifies how to alter an address when a mailbox filter **fileinto** operator is applied. The following example specifies that the folder name should be inserted as a subaddress into the original address, replacing any originally present subaddress:

fileinto \$U+\$S@\$D

Managing Sieve Scripts

This document provides information about Oracle Communications Messaging Server's implementation and configuration of ManageSieve, a protocol for user management of Sieve mail filters (RFC 5228) specified in RFC 5804.

This section contains the following topics:

- [Location of Managed Sieves](#)
 - [Script Name Storage](#)
 - [Stored Script Semantics](#)
- [ManageSieve Service Configuration](#)
 - [Configuring the ManageSieve Service](#)
 - [Supported Channel Options](#)
 - [Supported ManageSieve-Channel-Specific Options](#)
 - [ManageSieve Server Use of PORT_ACCESS Mapping](#)
- [Basic Command/Response Test Script](#)

Location of Managed Sieves

The ManageSieve server manages a collection of Sieve scripts belonging to each user. You may select one of these scripts as the active script to filter the user's mail. You can store the scripts either in LDAP or in disk files. In either case, the active script is stored

separately from the Sieve collection in either a different attribute or a separate file. You can store the script collection and the active script in different places. For example, you can store the collection on disk and store the active script in the user's LDAP entry.

Active scripts stored in LDAP can consist of multiple values of the same attribute. When the ManageSieve server reads multiple values, it appends them to form a single script. (Note, however, that you lose this storage pattern if the ManageSieve server writes the active script. The server always writes scripts to LDAP as a single attribute value.) The server stores the script collection differently. It uses a single attribute and takes each value to be a separate script.

The server stores active scripts on disk in a single file with a specific name. The server stores each script in the script collection in a separate file. Each file has the same name prefix, but the server appends a unique identifier to create a different file name for each script.

Script Name Storage

Each script in the collection has a name. The server stores this name as a comment at the beginning of the script. The ManageSieve server maintains this comment which does not appear as part of the Sieve data provided to the client.

The server does not store the active script's name in the script itself and therefore the script does not contain an initial comment.

Stored Script Semantics

The ManageSieve server is designed to accommodate the behavior of other agents that both access and modify stored Sieve scripts. It does not assume that stored scripts will remain unaltered and does not cache scripts internally.

Specifically, each ManageSieve operation causes the following initial steps to be performed:

1. The server reads the active script and script collections from their respective locations. It silently ignores scripts in the collection that do not begin with a valid name comment.
2. The server computes a hash for the content of each script that is found. Note that this hash does not cover the initial comment containing the script name.
3. The server performs a sanity check on all scripts in the collection. If two scripts have the same hash value, the server changes one of them by adding sufficient trailing white space to obtain a unique hash. Similarly, the server checks the names of every script in the collection. If two scripts have the same name, the server adds $-n$, where n is an integer, to make the name unique.
4. The server checks the hash of each script in the collection against the hash of the active script. The server considers the script with the matching hash to be the active script. If no script in the collection has a matching hash, the server creates a new script with a name of the form **ACTIVE-*datetime***, where *datetime* is the time the server performs the operation.
5. The server writes immediately to LDAP, disk, or both any changes the preceding two steps made to the collection.

The server performs the requested operation only after it performs these steps. This operation may cause additional changes that the server writes when the operation finishes.

This rather convoluted approach ensures that the ManageSieve server's view of the user's Sieve state remains consistent even if other agents manipulate the underlying data in unexpected ways.

ManageSieve Service Configuration

ManageSieve runs as a service under the dispatcher. To configure the ManageSieve service, you add a ManageSieve service to the dispatcher, create a channel definition, and set optional channel options. Note that the ManageSieve channel does not act as either a source or destination channel in the conventional sense in the MTA. Instead, it is used to store various ManageSieve configuration settings. In particular, the channel is used to specify where Sieves are stored. The **filter** channel option specifies where the active Sieve is stored while the **destinationfilter** channel option specifies where the Sieve collection is stored.

Configuring the ManageSieve Service

To configure the ManageSieve Service:

1. Edit the **dispatcher.cnf** file and add the following lines if you are running in a legacy configuration.

```
[SERVICE=MANAGESIEVE]
PORT=4190
IMAGE=IMTA_BIN:managesieve
LOGFILE=IMTA_LOG:managesieve_server.log
STACKSIZE=204800
PARAMETER=CHANNEL=managesieve
```

If you have a Unified Configuration, run the following **msconfig** commands.

```
msconfig
set dispatcher.service:MANAGESIEVE.tcp_ports 4190
set dispatcher.service:MANAGESIEVE.parameter "CHANNEL=managesieve"
set dispatcher.service:MANAGESIEVE.stacksize 204800
set dispatcher.service:MANAGESIEVE.image "IMTA_BIN:managesieve"
set dispatcher.service:MANAGESIEVE.logfilename "IMTA_LOG:managesieve_
server.log"
write
```

2. Create a channel definition specifying Sieve storage in LDAP by editing the **imta.cnf** configuration file in a legacy configuration and adding the following channel definition.

```
managesieve smtp_crlf filter ldap:/// $A?mailSieveRuleSource \
destinationfilter ldap:/// $A?mailSieveCollection maytlserver
managesieve-daemon
```

If you have a Unified Configuration, run the following commands.

```
msconfig
set channel:managesieve.official_host_name managesieve-daemon
set channel:managesieve.smtp_crlf
set channel:managesieve.filter "ldap:/// $A?mailSieveRuleSource"
set channel:managesieve.destinationfilter "ldap:/// $A?mailSieveCollection"
set channel:managesieve.maytlserver
set channel:managesieve.scriptlimit 32
set channel:managesieve.sievelimit 1048576
set channel:managesieve.sizelimit 4194304
write
```

Note: The values of the last three channel options shown are the defaults. We include these settings to point out the limits imposed by the ManageSieve server.

The specification of LDAP URLs in the channel definitions indicates that scripts are to be stored in LDAP. The **\$A** metacharacter substitutes the DN of the user's LDAP entry that was determined when the user authenticated to the ManageSieve server. Note the specification of the **maytlsserver** channel option. Although TLS is negotiated in the ManageSieve protocol and is therefore optional, most existing ManageSieve clients require it and will not work without TLS being enabled.

3. Set optional ManageSieve-channel-specific options. If you want to configure a custom implementation string in a legacy configuration, for example, add the following lines to a **managesieve_options** file.

```
CUSTOM_BANNER_STRING=Name
CUSTOM_VERSION_STRING=Version
```

To set these option in a Unified Configuration using **msconfig**, run the following commands.

```
msconfig
set channel:managesieve.options.custom_banner_string Name
set channel:managesieve.options.custom_version_string Version
write
```

A recipe, **managesieve.rcp**, is provided to set up a ManageSieve server in a Unified Configuration.

Supported Channel Options

Table 13–13 shows the channel options supported by the ManageSieve server and their descriptions:

Table 13–13 Supported Channel Options

Channel Option	Description
filter url	A URL specifying where the active script is stored. Either a file: or LDAP URL can be used. If an LDAP URL is used it must specify exactly one LDAP attribute. The following substitutions may be useful in constructing the URL: \$A - The DN of the user's LDAP entry determined by the authentication process. \$M - The user's UID determined by the authentication process. \$?whatever? - The store's hashdir algorithm is applied to <i>whatever</i> to produce a directory path.
destinationfilter url	A URL specifying where the user's script collection is stored. Either a file: or LDAP URL can be used. If an LDAP URL is used it must specify exactly one LDAP attribute. The same set of substitutions provided for the filter option are available.
scriptlimit integer	Maximum number of scripts allowed in a single user's collection. The default is 32.
sievelimit integer	The maximum size, in octets, of a single script. Note that this limit is imposed at the protocol, not processing, level. The default is 1048576.
sizelimit integer	The maximum combined size, in octets, of a user's scripts. The default is 4194304.
notlsserver, maytlsserver, musttlsserver	Specifies whether TLS cannot be negotiated, may be negotiated, or must be negotiated, respectively. notlsserver is the default, but note that many ManageSieve clients require the availability and use of TLS.

Table 13–13 (Cont.) Supported Channel Options

Channel Option	Description
disconnectbadauthlimit <i>integer</i>	Limit on the number of failed authentication attempts that can be done before the server disconnects.
disconnectbadcommandlimit <i>integer</i>	Limit on the number of bad commands the server will allow before disconnecting.
disconnectcommandlimit <i>integer</i>	Limit on the number of commands the server will allow before disconnecting.
slave_debug, noslave_debug	Controls whether or not the server produces debug output. noslave_debug is the default. Note that server logging does not normally include authentication information. Authentication information will be included if the AUTH_DEBUG ManageSieve-channel-specific option is set (see below).
smtp_crlf, smtp_cr, smtp_lf, smtp_crorlf	These options do not imply the use of the SMTP protocol. They are simply used to specify the line termination that the server requires for commands. smtp_crlf is the default and, given the use of counted literals in the protocol, should not be changed except under truly extraordinary circumstances.
ident*	Controls the sort of checks done on the client IP address. This information is only used for logging purposes in ManageSieve.
nameservers <i>ip4-addrss-list</i>	Controls what nameservers are used by any check requested by the ident* options. The default is to use the system default nameservers.

Supported ManageSieve-Channel-Specific Options

Table 13–14 shows the ManageSieve-channel-specific options supported by the ManageSieve server and their descriptions:

Table 13–14 Supported ManageSieve-Channel-Specific Options

ManageSieve-Channel-Specific-Option	Description
CUSTOM_BANNER_STRING <i>string</i>	Text to include in the IMPLEMENTATION capability response. The default is the Messaging Server product name.
CUSTOM_VERSION_STRING <i>string</i>	Text to include in the IMPLEMENTATION capability response. The default is the Messaging Server product version and build date.
LOG_CONNECTION <i>integer</i>	Same semantics as the LOG_CONNECTION TCPIP-channel-specific option.
IGNORE_BAD_CERT <i>integer</i>	Same semantics as the IGNORE_BAD_CERT TCPIP-channel-specific option.
COMMAND_TIMEOUT <i>integer</i>	The amount of time in seconds that the ManageSieve server will wait for a command. The default is 300 seconds (5 minutes).
STATUS_TIMEOUT <i>integer</i>	The amount of time in seconds that the ManageSieve server will wait for a status response to be sent. The default is 60 seconds.
TRACE_LEVEL <i>integer</i>	Enables protocol-level tracing of all protocol exchanges just above the TLS layer. Note that this will include authentication exchanges. The default value is 0.
AUTH_DEBUG <i>string</i>	Set of space-separated keys to pass to the authentication subsystem to enable various debug settings. See the HULA documentation for details about what keys are available. Note that setting AUTH_DEBUG and the slave_debug channel option (see above) will cause authentication command and response information to be included in the debug logs. The default is for no keys to be set.

Table 13–14 (Cont.) Supported ManageSieve-Channel-Specific Options

ManageSieve-Channel-Specific-Option	Description
MAX_THREADS <i>integer</i>	Same semantics as the MAX_SERVER_THREADS TCPIP-channel-specific option.

ManageSieve Server Use of PORT_ACCESS Mapping

The dispatcher checks connections to the ManageSieve server with the PORT_ACCESS mapping just as it does with any other connection it handles. Table 13–15 shows the server-specific PORT_ACCESS metacharacters that are relevant to ManageSieve and a description of how they work.

Table 13–15 PORT_ACCESS Metacharacters

PORT_ACCESS Metacharacters	Description
\$U	\$U is the equivalent of setting the <code>slave_debug</code> channel option, except it only affects the current connection.
\$G	Set TRACE_LEVEL ManageSieve-channel-specific option value to 2 for this connection only.

Basic Command/Response Test Script

The following output shows a basic command/response test script.

```
"IMPLEMENTATION" "YoyoDyne file test 0.00000000001"
"VERSION" "1.0"
"MAXREDIRECTS" "32"
"NOTIFY" "mailto"
"SIEVE" "copy date editheader encoded-character envelope-auth envelope
envelope-dsn environment ereject extracttext foreverypart ihave index imap4flags
mime reject relational redirect-dsn subaddress spantest spantestplus variables
virustest body extlists fileinto notify enotify regex vacation vacation-seconds"
"STARTTLS"
"SASL" "PLAIN"
OK
noop
OK "NOOP Completed"
noop "doof"
OK (TAG {4}
doof) "Done"
authenticate "plain" "xxxxxxxxxxxxx"
OK "Authentication successful"
listscripts
OK
putsript "discard" "discard;"
OK "Sieve added"
listscripts
"discard"
OK
getsript "discard"
{8}
discard;
OK
putsript "keep" "keep;"
OK "Sieve added"
listscripts
```

```

"keep"
"discard"
OK
setactive "discard"
OK "Sieve activated"
listscripts
"keep"
"discard" ACTIVE
OK
setactive "keep"
OK "Sieve activated"
listscripts
"keep" ACTIVE
"discard"
OK
setactive "discard"
OK "Sieve activated"
listscripts
"keep"
"discard" ACTIVE
OK
setactive ""
OK "Active Sieve deactivated"
listscripts
"keep"
"discard"
OK
deletescript "keep"
OK "Sieve deleted"
deletescript "discard"
OK "Sieve deleted"
listscripts
OK
putsript "test" "refuse \"this is a test\";"
NO {72}
Refuse not listed in require clause prior to use [ "this is a test" ; ^]
putsript "test" "require \"refuse\"; refuse \"this is a test\";"
OK "Sieve added"
getsript "test"
{42}
require "refuse"; refuse "this is a test";
OK
renamescript "test" "test2"
OK "Sieve renamed"
listscripts
"test2"
OK
getsript "test2"
{42}
require "refuse"; refuse "this is a test";
OK
renamescript "test3" "test"
NO (NONEXISTENT) "Source Sieve does not exist"
putsript "discard" "discard;"
OK "Sieve added"
renamescript "test2" "discard"
NO (ALREADYEXISTS) "Destination Sieve already exists"
putsript "doof" ""
NO "Quoted string too short"
deletescript "discard"

```

```
OK "Sieve deleted"
deletescript "test2"
OK "Sieve deleted"
havespace "doof" 1234
OK
havespace "doof" 12345678
NO (QUOTA/MAXSIZE) "Maximum Sieve size exceeded"
havespace "doof" 123456789
NO (QUOTA/MAXSIZE) "Maximum Sieve size exceeded"
havespace "doof" 1234567890
NO "Number is too large"
logout
BYE "Disconnecting"
```

To Create MTA-Wide Filters

MTA-wide filters apply to all messages enqueued to the MTA. A typical use for this type of filter is to block unsolicited bulk email or other unwanted messages regardless of the messages' destinations. To create an MTA-wide filter:

1. Write the filter using Sieve
2. Store the filter in the following file: *MessagingServer_home/config/imta.filter* This filter file must be world readable. It is used automatically, if it exists.
3. Recompile the configuration and restart the Dispatcher When using a compiled configuration, the MTA-wide filter file is incorporated into the compiled configuration.

Routing Discarded Messages Out the FILTER_DISCARD Channel

By default, messages discarded via a mailbox filter are immediately discarded (deleted) from the system. However, when users are first setting up mailbox filters (and perhaps making mistakes), or for debugging purposes, it can be useful to have the deletion operation delayed for a period.

To have mailbox filter discarded messages temporarily retained on the system for later deletion, first add a **filter_discard** channel to your MTA configuration with the **notices** channel option specifying the length of time (normally number of days) to retain the messages before deleting them, as shown in the following example:

```
filter_discard notices 7
filter-discard
```

Then set the MTA option **FILTER_DISCARD** to 2 by running **msconfig set filter_discard 2**. Messages in the **filter_discard** queue area should be considered to be in an extension of users' personal wastebasket folders. As such, note that warning messages are never sent for messages in the **filter_discard** queue area, nor are such messages returned to their senders when a bounce or return is requested. Rather, the only action taken for such messages is to eventually silently delete them, either when the final notices value expires, or if a manual bounce is requested using a utility such as **imsimta return**.

The use of the **filter_discard** channel by the **jettison** Sieve action is controlled by the **FILTER_JETTISON** option, which takes its default from the **FILTER_DISCARD** setting. **FILTER_DISCARD** in turn defaults to 1 (discards go to the **bitbucket** channel).

To Debug User-level Filters

If a user complains that a Sieve filter is not behaving as expected, there are a number of steps you can take to debug the filters. These are described here.

1. For **fileinto** filtering to work, run **msconfig edit channels** and check that the **ims-ms** channel is marked as follows:

```
fileinto $U+$S@$D
```

2. Get the user level filters from the user's LDAP entry. User level filters are stored in their LDAP entry under the **MailSieveRuleSource** attribute(s). To retrieve this with a **ldapsearch** command, remember they are base64 encoded so you will need to decode the output by using the **-Bo** switch.

```
ldapsearch -D "cn=directory manager" -w password -b
"o=alcatraz.example.org,o=isp" -Bo uid=test
```

The **imsimta test -rewrite** command, described below, will also automatically decode them.

3. Verify that the use's filters are being seen by the MTA. Issue the command:

```
imsimta test -rewrite -filter -debug user@example.org
```

This should output the user's Sieve filters that you retrieve in the previous step. If you do not see the filters, then you need to figure out why the LDAP entry isn't returning them. If the **imsimta test -rewrite** output shows the filters, then you know that the user's filters are being seen by the MTA. The next step is to test the interpretation of the filters by using the **imsimta test -expression** command.

4. Use **imsimta test -exp** to debug the user's filter. The following information is required:

1. The user's Sieve language statements from their **mailSieveRuleSource** attribute. See the steps above.

2. The rfc2822 message that was supposed to trigger the filter.

3. Description of what they expected the filter to do to the message.

5. a text file (example: **temp.filter**) containing the Sieve language statements based on the user's **mailSieveRuleSource: values**. Example:

```
require "fileinto";
if anyof(header :contains
["To", "Cc", "Bcc", "Resent-to", "Resent-cc", "Resent-bcc"] "commsqa"){
    fileinto "QMSG";
}
```

Expected result: if **commsqa** is a recipient for this message, then file the message into a folder called **QMSG**.

6. Create a text file called **test.msg** that contains the contents of the rfc2822 message file supplied by user. You can either use a **.msg** file from the user's message store area, or create a text file called **test_rfc2822.msg** that contains the contents of the rfc2822 message file supplied by user.

7. Use the **imsimta test -exp** command:

```
imsimta test -exp -mm -block -input=temp.filter -message=test_rfc2822.msg
```

8. Examine the output. The last lines of the **imsimta test -exp** command will show the result of the Sieve interpretation. They will look like this:

```
Sieve Result: []
```

or this:

```
Sieve Result: [_action_]
```

where *action* is the action that would be done as a result of applying the Sieve filter on this message. If the criteria of the filter matched, you will have some action displayed as the result. If nothing matched, then the Sieve result will be blank, and there is either a logic error in the Sieve filter or the `.msg` file doesn't contain matching information. If you get any other error, then there is a syntax error in the Sieve script file, and you need to debug it. For more details on the output, see ["imsimta test -exp Output."](#)

9. If the filter is syntactically valid and results are correct, then the next step is to examine a `tcp_local_slave.log` debug log file. It may be that the message file you're testing and the one being sent aren't identical. The only way to see what's being received is to examine a `tcp_local_slave.log` file. This log will show you the exact message being sent to the MTA and how the filter is being applied to that message. For more information on getting a `tcp_local_slave.log` debug file, see the `slave_debug` option in the discussion on debugging channel master and slave programs in the *Messaging Server Reference*.

imsimta test -exp Output

The full command `imsimta test -exp` for is as follows:

```
imsimta test -exp -mm -block -input=temp.filter -message=rfc2822.msg
```

An example of the output is as follows:

Example - imsimta test -exp Output

```
# imsimta test -exp -mm -block -input tmp.filter -message=rfc2822.msg
Expression: if header :contains ["to"] ["pamw"]          (1)
Expression: {
Expression: redirect "usr3@example.org";
Expression: keep;
Expression: }
Expression:
Expression: Dump: header:2000114;0 3 1 :contains 1 "to" 1
"pamw" if 8 ;
Dump: redirect:2000121;0 1 1 "usr3@example.org" ; keep:2000117;0 (2)
Dump: 0
Result: 0
Filter result: [ redirect "usr3@example.org" keep ]      (3)
```

1) The Expression: output lines show the filter being read and parsed from `tmp.filter` text file. These are not particularly useful in debugging the script.

2) The **Dump**: output lines are the result of the computer interpreting the Sieve statements. You should not see any errors and the output should seem to match your input. For example the dump shows the word `redirect`, `usr3@example.org` which is like the line in the filter file `redirect "usr3@example.org";`.

If it didn't show this matching text, then you'd be concerned, otherwise, these also are not particularly useful in debugging the script.

3) At the bottom of the output you will get the **Filter result**: statement. As stated earlier there are two possible results:

Sieve Result: [] or this: **Sieve Result:** [*action*]

where *action* is the action taken by the Sieve script. Note that sometimes the results are expected to be empty. For example, for a discard filter, you should test that it doesn't always discard every **.msg** file you test it against. If there is some action between the brackets, for example:

Filter result: [fileinto "QMSG" keep]

This means the text in the **rfc2822.msg** file matched the filter criteria. In this particular example, the filter will file the mail into folder **QMSG** and keep a copy in the inbox. The resulting actions in this case are **fileinto** and **keep**.

When testing filters you should test various **.msg** files for both results. You should always test that messages that match your filter are filtered, and messages that you do not want to match are not filtered.

Keep in mind that if for wildcard matches, you must use the **:matches** test and not **:contains**. For example, if you wish **from=*@example.org** to match, you must use **:matches** or the test will fail as it will not ever satisfy the test condition.

imsimta test -exp Syntax

imsimta test -exp tests Sieve language statements against a specified RFC2822 message and sends the results of the filter to standard output.

The syntax is as follows:

imsimta test -exp -mm -block -input=*Sieve_language_scriptfile*-message=*rfc2822_message_file*

where,

-block treats the entire input as a single Sieve script. The default is to treat each line as a separate script and to evaluate it separately. The Sieve will only be evaluated once the end of file is reached.

Using and Configuring MeterMaid for Access Control

MeterMaid is a server that can provide centralized metering and management of connections and transactions through monitoring IP addresses and SMTP envelope addresses. Functionally, MeterMaid can be used to limit how often a particular IP address can connect to the MTA. Limiting connections by particular IP addresses is useful for preventing excessive connections used in denial-of-service attacks. MeterMaid supplants `conn_throttle.so` by providing similar functionality, but extending it across the Oracle Communications Messaging Server installation. No new enhancements are planned for `conn_throttle.so` and MeterMaid is its more effective replacement.

Technical Overview

`conn_throttle.so` is a shared library used as a callout from the MTA's mapping table that uses an in-memory table of incoming connections to determine when a particular IP address has recently connected too often and should be turned away for awhile. While having an in-memory table is good for performance, its largest cost is that each individual process on each server maintains its own table.

In most cases, the `conn_throttle.so` callout is done in the `PORT_ACCESS` mapping that is accessed by the Dispatcher, a single process on each system. The only cost is that there is a separate table per server.

The primary improvement by MeterMaid is that it maintains a single repository of the throttling information that can be accessed by all systems and processes within the Messaging Server environment. It continues to maintain an in-memory database to store this data to maximize performance. Restarting MeterMaid will lose all information previously stored, but since the data is typically very short lived, the cost of such a restart (done infrequently) is very low.

Theory of Operations

MeterMaid's configuration is maintained by the `msconfig` command in Unified Configuration or the `configutil` command in legacy configuration.

MeterMaid is accessed from the MTA through a mapping table callout using `check_metermaid.so`. It can be called from any of the `*_ACCESS` tables. When called from the `PORT_ACCESS` table, it can be used to check limits based on the IP address of the connection which will be the most common way to implement MeterMaid as a replacement for the older `conn_throttle.so`. If called from other `*_ACCESS` tables, MeterMaid can also be used to establish limits on other data such as the envelope from or envelope to addresses as well as IP addresses.

This chapter only describes the throttle entry point in `check_metermaid.so`. For a complete list of `check_metermaid.so` entry points, refer to "[MeterMaid Reference](#)." The throttle routine contacts MeterMaid providing two subsequent arguments separated by commas. The first is the name of the table against which the data will be checked, and the second is the data to be checked.

If the result from the probe is that the particular data being checked has exceeded its quota in that table, `check_metermaid.so` returns `success` so that the mapping engine will continue processing this entry. The remainder of the entry would then be used to handle this connection that has exceeded its quota.

```
PORT_ACCESS
```

```
*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
*|*|*|*|* $C$:A$[/opt/sun/comms/messaging64/lib/ check_
metermaid.so,throttle,tablename,$3]\
$N421$ Connection$ declined$ at$ this$ time$E
* $YEXTERNAL
```

Note the `:$A` flag test in the mapping table entry before the call to `check_metermaid.so`. This is to ensure that we only do the MeterMaid probe when `PORT_ACCESS` is being checked by the dispatcher as it will set the `A` flag for its probe.

Options for MeterMaid

MeterMaid's configuration is maintained by setting options using the `msconfig` command in Unified Configuration or using the `configutil` command in legacy configuration. [Table 14–1](#) describes some of the settings currently supported by MeterMaid.

Table 14–1 Options for MeterMaid

Unified Configuration Option	Legacy Configuration Option	Description
<code>metermaid.enable</code>	<code>local.metermaid.enable</code>	This setting must be set to 1 on the system that will run the MeterMaid daemon so that the Watcher will start and control MeterMaid.
<code>metermaid.logfile.*</code>	<code>logfile.metermaid.*</code>	These settings are the same as those used by imap, pop, and other services. By default MeterMaid writes its log file into <code>MessagingServer_home/data/log/metermaid</code> .
<code>metermaid.listenaddr</code>	<code>metermaid.config.listenaddr</code>	The address to which MeterMaid should bind. On most systems, the default would not need to be changed, but for multi-homed or HA systems, specifying the appropriate address here is recommended. Default: <code>(INADDR_ANY)</code>
<code>metermaid.maxthreads</code>	<code>metermaid.config.maxthreads</code>	The MeterMaid server is multithreaded and maintains a pool of threads onto which its tasks are scheduled. This value sets the maximum number of threads that will be used by MeterMaid. On systems with more than 4 CPUs, increasing this value may increase overall throughput. Default: 20

Table 14–1 (Cont.) Options for MeterMaid

Unified Configuration Option	Legacy Configuration Option	Description
<code>metermaid.port</code>	<code>metermaid.config.port</code>	This is the port to which MeterMaid listens for connections and to which MeterMaid clients will connect. Default: 63837
<code>metermaid.secret</code>	<code>metermaid.config.secret</code>	In order to authenticate incoming connections, MeterMaid uses a shared secret that the clients send once they connect to MeterMaid. No default. Value must be supplied.
<code>metermaid.sslusessl</code>	<code>service.metermaid.sslusessl</code>	Requires the use of SSL for all incoming connections to the MeterMaid server. Default: 0

Table 14–2 describes the options used by the `check_metermaid` client:

Table 14–2 `check_metermaid` Options

Unified Configuration Option	Legacy Configuration Option	Description
<code>metermaid_client.connectfrequency</code>	<code>metermaid.mtaclient.connectfrequency</code>	Attempt a connection every <i>connectfrequency</i> seconds. When the client needs to connect to MeterMaid, it uses this as an internal throttle to prevent constant connection attempts when MeterMaid isn't available. During the time that the client is unable to communicate with MeterMaid, it will return a "fail" status to the MTA mapping engine indicating that MeterMaid has not blocked this connection. For example, if <code>check_metermaid.so</code> attempts to connect to MeterMaid, but it fails for some reason, during the next N seconds as specified by <code>metermaid_client.connectfrequency</code> (or <code>metermaid.mtaclient.connectfrequency</code> in legacy configuration), no additional attempts will be attempted. It prevents <code>check_metermaid.so</code> from trying to connect to MeterMaid too frequently if it is not working. Default: 15
<code>metermaid_client.connecttimeout</code>	<code>metermaid.mtaclient.connectwait</code>	When the client is waiting for a connection to MeterMaid (either an initial connection or to reuse another already established connection), it will wait for <i>connecttimeout</i> seconds before returning a fail status and allowing this connection to continue. Default: 5
<code>metermaid_client.debug</code>	<code>metermaid.mtaclient.debug</code>	If this option is enabled, debugging information from the client will be printed into either the server or thread-specific log file for the SMTP server. Default: no

Table 14–2 (Cont.) check_metermaid Options

Unified Configuration Option	Legacy Configuration Option	Description
metermaid_client.max_conns	metermaid.mtaclient.maxconns	In order to support multithreaded servers, the client can maintain a pool of connections to MeterMaid. By doing this, there can be increased concurrency during communications. However, due to internal locking done by MeterMaid, access to a particular table is limited to one request at a time, so multiple connections from a single process may provide limited benefit. Default: 3
metermaid_client.timeout	metermaid.mtaclient.readwait	When communicating with MeterMaid, the client will wait <i>timeout</i> seconds before returning a fail status and allowing this connection to continue. Default: 10
metermaid_client.server_host	metermaid.config.serverhost	This is the host name or IP address to which the clients will connect. It may be the same as metermaid.listenaddr but will most likely have a particular value to direct clients to one system in particular in the Messaging Server environment. No default. Value must be supplied.
metermaid_client.sslusessl	metermaid.mtaclient.sslusessl	Enables SSL communication with an SSL-enabled MeterMaid server. Default: 0

Lastly, the throttling tables are also defined in Unified and legacy configurations. [Table 14–3](#) describes the options defining the throttling tables. The * in each configuration option is the name of the particular table being defined. For example, for a table called `internal`, the first option would be called **metermaid.local_table:internal.data_type** in Unified Configuration or **metermaid.table.internal.data_type** in legacy configuration.

Table 14–3 Options Defining Throttling Tables

Unified Configuration Option	Legacy Configuration Option	Valid for Table Types	Description
metermaid.local_table:*.data_type	metermaid.table:*.data_type	throttle simple greylisting	MeterMaid can support two kinds of data in its tables, string and ipv4. String data is limited to 255 bytes per entry and can be compared using case-sensitive or case-insensitive functions (see metermaid.local_table:*.options below). Default: string

Table 14-3 (Cont.) Options Defining Throttling Tables

Unified Configuration Option	Legacy Configuration Option	Valid for Table Types	Description
<code>metermaid.local_table:*.max_entries</code>	<code>metermaid.table:*.max_entries</code>	throttle simple greylisting	When MeterMaid initializes each table, it pre-allocates this many entries. MeterMaid automatically recycles old entries, even if they haven't yet expired. When a new connection is received, MeterMaid will reuse the least recently accessed entry. A site should specify a value high enough to cache the connections received during <i>quota_time</i> . Default: 1000
<code>metermaid.local_table:*.options</code>	<code>metermaid.table:*.options</code>	throttle	This option is a comma-separated list of keywords that defines behavior or characteristics for the table. Valid keywords are: <ul style="list-style-type: none"> ▪ nocase - When working with the data, all comparisons are done using a case-insensitive comparison function. (This option is valid only for string data). ▪ penalize - After <i>quota_time</i> seconds, throttle will normally reset the connection count to 0, but if the penalize option is enabled, throttle will decrement the connection count by quota (but not less than 0) so that additional connection attempts will penalize future <i>quota_time</i> periods. For example, if <i>quota</i> were 5 with a <i>quota_time</i> of 60, and the system received 12 connection attempts during the first minute, the first 5 connections would be accepted and the remaining 7 would be declined. After 60 seconds has passed, the number of connections counted against the particular address would be reduced to 7, still keeping it above quota and declining connection attempts. Assuming no additional connection attempts were made, after another 60 second period, the number of connections would be further reduced down to 2, and MeterMaid would permit connection attempts again.
<code>metermaid.local_table:*.quota</code>	<code>metermaid.table:*.quota</code>	throttle	When a connection is received, it is counted against quota. If the number of connections received in <i>quota_time</i> seconds exceeds this value, MeterMaid will decline the connection. (The actual effect on the incoming connection is controlled by the mapping table and could result in additional scrutiny, a delay, or denying the connection.) Default: 100
<code>metermaid.local_table:*.quota_time</code>	<code>metermaid.table:*.quota_time</code>	throttle	This specifies the number of seconds during which connections will be counted against quota. After this many seconds, the number of connections counted against the incoming address will be reduced depending on the type of this table. Default: 60

Table 14–3 (Cont.) Options Defining Throttling Tables

Unified Configuration Option	Legacy Configuration Option	Valid for Table Types	Description
<code>metermaid.local_table:*.storage</code>	<code>metermaid.table:*.storage</code>	throttle simple greylisting	MeterMaid can use two different storage methods, hash and splay. The default hash table method is recommended, but under some circumstances a splay tree may provide faster lookups. Default: hash
<code>metermaid.local_table:*.table_type</code>	<code>metermaid.table:*.type</code>	NA	MeterMaid supports three table types: <ul style="list-style-type: none"> throttle (default) - This type of table keeps track of the data, typically IP addresses, and will throttle the incoming connections to quota connections during a period of <code>quota_time</code> seconds. simple - This type of table may be used to store arbitrary data referenced by a key. greylisting - This type of table may be used to provide an anti-spam/anti-virus technique. More information about setting up this type of table can be found in the "Implementing Greylisting by Using MeterMaid" chapter.
<code>metermaid.local_table:*.block_time</code>	<code>metermaid.table:*.block_time</code>	greylisting	Specifies the ISO 8601 duration for how long we temporarily reject each delivery attempt based on sender and recipient information. Default: pt5m (5 minutes)
<code>metermaid.local_table:*.resubmit_time</code>	<code>metermaid.table:*.resubmit_time</code>	greylisting	Specifies the ISO 8601 duration during which, but after <code>block_time</code> , we must receive a subsequent delivery attempt based on the same sender and recipient information previously blocked. This sender and recipient combination is now flagged as permitted. Default: pt4h (4 hours)
<code>metermaid.local_table:*.inactivity_time</code>	<code>metermaid.table:*.inactivity_time</code>	greylisting	Specifies the ISO 8601 duration for how long we will continue to accept messages based on the sender and recipient information previously permitted. This permission expires after <code>inactivity_time</code> from the last allowed delivery. Default: p7d (7 days)

Table 14–4 describes the options used to enable access to multiple MeterMaid Servers from `check_metermaid.so`.

Table 14–4 Options Used to Enable Access to Multiple MeterMaid Servers from check_metermaid.so

Unified Configuration	Legacy Configuration	Description
<code>metermaid_client.remote_table:table.server_nickname</code>	<code>metermaid.mta_client.remote_table.table.server_nickname</code>	Specifies the nickname for a particular MeterMaid server that is responsible for the referenced table. The nickname is a short keyword, consisting only of letters, numbers, and underscores, that will be used in the <code>remote_server</code> option names. No default. Value must be supplied.
<code>metermaid_client.remote_server:nickname.max_conns</code>	<code>metermaid.mta_client.remote_server.nickname.max_conns</code>	Specifies the maximum number of concurrent connections to the MeterMaid server referenced by nickname. Default: 3
<code>metermaid_client.remote_server:nickname.server_host</code>	<code>metermaid.mta_client.remote_server.nickname.server_host</code>	Specifies the host name of the MeterMaid server referenced by nickname. Defaults to the value of <code>metermaid_client.server_host</code> .
<code>metermaid_client.remote_server:nickname.server_port</code>	<code>metermaid.mta_client.remote_server.nickname.server_port</code>	Specifies the port number of the MeterMaid server referenced by nickname. Default: 63837
<code>metermaid_client.remote_server:nickname.sslusessl</code>	<code>metermaid.mta_client.remote_server.nickname.sslusessl</code>	Specifies whether or not to use SSL for the connection to the MeterMaid server referenced by nickname. Defaults to the value of <code>metermaid_client.sslusessl</code> .

Limit Excessive IP Address Connections Using Metermaid – Example

This example uses MeterMaid to throttle IP addresses at 10 connections per minute. For reference, the equivalent `conn_throttle.so` setup in the mappings table would be as follows:

```
PORT_ACCESS
```

```
*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
*|*|*|*|* $C$[/opt/sun/comms/messaging/lib/conn_throttle.so,throttle,$3,10]\
$N421$ Connection$ declined$ at$ this$ time$E
* $YEXTERNAL
```

This `PORT_ACCESS` mapping table implements `conn_throttle.so` to restrict connections to a rate of no more than 10 connections per minute for non-INTERNAL connections.

One fundamental difference between the two technologies is that instead of configuring details such as the rate-limit for throttling directly into the mapping table, MeterMaid uses `msconfig` for these settings, as the following example shows.

On systems running the MeterMaid server:

1. Enable MeterMaid by running the following command:

```
msconfig set metermaid.enable 1
```

or in legacy configuration:

```
configutil -o local.metermaid.enable -v TRUE
```

2. Set an authentication password used to verify communications between the client and MeterMaid server:

```
msconfig set metermaid.secret password
```

or in legacy configuration:

```
configutil -o metermaid.config.secret -v password
```

3. Define a throttling table

MeterMaid's throttling behavior is determined by the use of named throttling tables that define operating characteristics. To define a table that throttles at a rate of 10 connections per minute, set the following options in **msconfig**:

```
set metermaid.local_table:ext_throttle.data_type ipv4
set metermaid.local_table:ext_throttle.quota 10
```

or the following for legacy configuration using **configutil**:

```
configutil -o metermaid.table.ext_throttle.data_type -v ipv4
configutil -o metermaid.table.ext_throttle.quota -v 10
```

ext_throttle is the name of the throttling table. **ipv4** is the data type Internet Protocol version 4 address representation. 10 is the quota (connection limit).

4. On the MeterMaid system, start MeterMaid.

```
start-msg metermaid
```

5. On systems where the MTA will use MeterMaid to do throttling, specify the MeterMaid host and password.

These are required:

```
msconfig set metermaid.secret MeterMaid_Password
msconfig set metermaid_client.server_host name_or_ipaddress_of_MetermaidHost
```

or the following for legacy configuration:

```
configutil -o metermaid.config.secret -v MeterMaid_Password
configutil -o metermaid.config.serverhost -v name_or_ipaddress_of_
MetermaidHost
```

6. Set up the MeterMaid **PORT_ACCESS** table.

This table is similar to the equivalent **conn_throttle.so** setup:

```
PORT_ACCESS

*|*|*|*|* $C$|INTERNAL_IP;$3|$Y$E
*|*|*|*|* $C$:A$[/opt/sun/comms/messaging/lib/check_metermaid.so,throttle,\
ext_throttle,$3]$N421$ Connection$ declined$ at$ this$ time$E
* $YEXTERNAL
```

The first line checks to see if the IP address attempting a connection is internal. If it is, it allows the connection. The second line runs the IP address through MeterMaid and if it has connected too frequently, it declines the connection. The third line allows any other connections through, but flagged as **EXTERNAL**.

This call to **check_metermaid.so** is very similar to the callout to **conn_throttle.so**. The function in **check_metermaid.so** is the same. **throttle** and its arguments are simply the table name as configured using **metermaid.local_table:tablename** and the IP address to check (**\$3**). Like **conn_throttle.so**, this function returns *success* when the limit (as specified in **metermaid.local_table:ext_throttle.quota**) has been reached. This allows the remainder of the mapping entry line is processed, which sends a message (421 SMTP code, transient negative completion, *Connection not accepted at this time*) to the remote SMTP client, and tells the Dispatcher to close the connection.

:\$A ensures that this line will only be processed when being called from the Dispatcher. Without this, the call to **check_metermaid.so** would also happen in the context of the **tcp_smtp_server** processes which also probes the **PORT_ACCESS** mapping table. This would cause MeterMaid to count each incoming connection twice.

This is the basic configuration to set up MeterMaid as a **conn_throttle.so** replacement. For more information, see the discussion on mapping operations in the *Messaging Server Reference*, as well as "[PORT_ACCESS Mapping Table](#)" for information on these topics.

Configuring check_metermaid.so Clients to Access Multiple MeterMaid Servers

The topics in this section include:

- [Considerations for Distributing Load Across Multiple MeterMaid Servers](#)
- [Configuring check_metermaid.so to Access Multiple MeterMaid Servers](#)

Considerations for Distributing Load Across Multiple MeterMaid Servers

If you have multiple MeterMaid tables in your deployment, you may be able to improve overall performance by distributing them across multiple MeterMaid servers. The **check_metermaid.so** client supports associations between MeterMaid tables and the servers that are responsible for their respective tables. The client also supports per server options for concurrency and use of SSL.

Configuring check_metermaid.so to Access Multiple MeterMaid Servers

To configure **check_metermaid.so** to access multiple MeterMaid Servers:

1. Define the list of all tables and associate them with nicknames for each server using **msconfig** (Unified Configuration):

```
set metermaid_client.remote_table:table1.server_nickname "alpha"
set metermaid_client.remote_table:table2.server_nickname "alpha"
set metermaid_client.remote_table:table3.server_nickname "beta"
```

or using **configutil** (legacy configuration):

```
configutil -o metermaid.mtaclient.remote_table.table1.server_nickname -v
"alpha"
configutil -o metermaid.mtaclient.remote_table.table2.server_nickname -v
"alpha"
configutil -o metermaid.mtaclient.remote_table.table3.server_nickname -v
"beta"
```

where *table1*, *table2*, and *table3* are tables defined in your deployment's MeterMaid configuration. The servers' nicknames allow the **check_metermaid.so** client to look for remote servers associated with those nicknames. Since nicknames become part of the configuration option names for the remote server definitions, they must include only letters, numbers, and underscores.

2. Create configuration entries for each server nickname for the host name and port for the server, the maximum number of connections, and whether it uses SSL. Run the following commands if you are using **msconfig** (Unified Configuration):

```
set metermaid_client.remote_server:alpha.max_conns 3
set metermaid_client.remote_server:alpha.server_host "alpha.example.com"
set metermaid_client.remote_server:alpha.server_port 63837
set metermaid_client.remote_server:alpha.sslusessl 0
set metermaid_client.remote_server:beta.max_conns 3
set metermaid_client.remote_server:beta.server_host "beta.example.com"
set metermaid_client.remote_server:beta.server_port 63837
set metermaid_client.remote_server:beta.sslusessl 0
```

Or run the following commands if you are using **configutil** (legacy configuration):

```
configutil -o metermaid.mtaclient.remote_server.alpha.max_conns -v 3
configutil -o metermaid.mtaclient.remote_server.alpha.server_host -v
"alpha.example.com"
configutil -o metermaid.mtaclient.remote_server.alpha.server_port -v 63837
configutil -o metermaid.mtaclient.remote_server.alpha.sslusessl -v 0
configutil -o metermaid.mtaclient.remote_server.beta.max_conns -v 3
configutil -o metermaid.mtaclient.remote_server.beta.server_host -v
"beta.example.com"
configutil -o metermaid.mtaclient.remote_server.beta.server_port -v 63837
configutil -o metermaid.mtaclient.remote_server.beta.sslusessl -v 0
```

For descriptions of these options, please see "[MeterMaid Reference](#)."

3. If you are using Unified Configuration and your configuration is compiled, recompile your configuration. This step is not necessary if you are using legacy configuration:

```
imsimta cnbuild
```

4. Restart components that are using **check_metermaid.so**. You would typically do this by restarting the dispatcher:

```
stop-msg dispatcher
start-msg dispatcher
```

Implementing Greylisting by Using MeterMaid

This chapter describes how to implement greylisting by using MeterMaid as well as enhancing its functionality.

What is Greylisting and How Does It Work?

Greylisting is a technique used by some MTAs as a way to reduce the number of undesirable spam messages they receive. Simply put, greylisting initially gives a temporary rejection to all incoming mail the first time it sees it, but then permits it upon subsequent attempts. It works by making the assumption that most spam is sent by spambots, PCs that have been compromised by a virus or trojan software, that act as mass-mailing clients. In order to send out as much spam as possible, these systems will connect to a mail server and attempt to deliver the spam to the recipient. If it should encounter a failure, it is extremely unlikely to retry delivery. Proper MTA clients will reschedule and reattempt delivery upon receiving a temporary failure, thus allowing the greylisting mail server a second chance to permit the message to be received.

Greylisting matches messages by using a combination of the source IP address, the envelope **FROM:** address, and the envelope **TO:** address. By using this triplet, greylisting works before the message body is sent during the SMTP transaction, making the temporary rejections happen in response to the **RCPT TO:** command. When the same triplet is presented to the mail server during a subsequent delivery attempt, the **RCPT TO:** command returns a successful response and the mail server will then accept the message for delivery.

In some setups, greylisting has been seen to reduce the number of spam messages received by 80-90%.

The downside to greylisting, however, is that it introduces an artificial delay to incoming mail from previously unknown triplets. The length of this delay varies depending on the originating MTA, but could range from thirty minutes to several hours. It is now expected by many that e-mail is nearly instantaneous, and greylisting can have a negative impact on customer's expectations. [Table 15-1](#) provides an example of greylisting in action.

Table 15–1 Greylisting at a Glance: Example

Time	Action	SMTP Result	Explanation
9:45	Incoming SMTP transaction from john@example.com to susan@example.com (local user)	451 4.5.1 Temporary failure - retry later	This is the first time that Oracle Communications Messaging Server has seen mail from john@example.com going to susan@example.com so Messaging Server responds with a temporary rejection.
9:47	Another transaction from john@example.com to susan@example.com	451 4.5.1 Temporary failure - retry later	Because this attempt happened within a short window after the first attempt, it is also temporarily rejected.
10:15	A bit later, the same transaction is retried	250 OK	Now that a subsequent attempt was made within the resubmit window, Messaging Server consider this combination of sender and recipient permitted.
10:20	Mail from stephen@example.com to susan@example.com	451 4.5.1 Temporary failure - retry later	This is a different sender, so it is handled independently from the previously permitted combination. This combination would be permitted after the block period has passed.
The next day	A new message from john@example.com to susan@example.com	250 OK	Since this combination is permitted, it remains valid for an extended period.

MeterMaid supports a greylisting table with additional related tuning options.

Table 15–2 describes the different implementations of greylisting.

Table 15–2 Greylisting Features

Feature	Implementation
Rejects previously unseen triplet	Yes
Continues rejecting for an initial blocking period (to block spambots that automatically retry within a very short period)	Yes
Once accepted, continues to accept messages from that triplet	Yes
Requires subsequent attempt within a specified period of time to register the triplet as permitted	Yes
Allows pre-registration of triplets based on outgoing mail, permitting wildcard source IP address	Yes
Expiration of existing triplets can be extended by recent use	Yes

Greylisting has more functionality than using a throttle table.

Basic Greylisting Implementation

This section contains information on using Greylisting with Messaging Server.

Using Greylisting with Messaging Server

Setting up greylisting with Messaging Server is easy due to MeterMaid's built-in support for greylisting tables. Instead of calling the **throttle** routine in **check_**

metermaid.so, you can use the **greylisting** routine which handles the appropriate return values for allowing Messaging Server to block transactions when the routine returns success.

First, the MeterMaid table definition:

```
metermaid.table.greylisr.type = greylisting
metermaid.table.greylisr.data_type = string
metermaid.table.greylisr.max_entries = 50000
metermaid.table.greylisr.options = nocase
metermaid.table.greylisr.block_time = pt5m
metermaid.table.greylisr.resubmit_time = pt4h
metermaid.table.greylisr.inactivity_time = p7d
```

Note that the time formats can now be in ISO 8601 duration format.

This table defines a greylisting table with the following characteristics:

- All triplets are rejected during the first 5 minutes, even if multiple attempts occur.
- In order for a triplet to be recognized and permitted, a subsequent attempt must be made after that 5 minute window, but within 4 hours of the initial attempt.
- Once a triplet is permitted, it will remain as permitted in the table for 7 days after its last use.

The corresponding access control mapping table is simpler when using a greylisting table as no special handling by the mapping table is required:

```
ORIG_MAIL_ACCESS

! Check the source IP address, sender, and recipient in MeterMaid's greylisr
table.
! If the call to greylisting() returns success, then Messaging Server should
return
! a temporary rejection. If the call fails, then the greylisting check has passed
! and other access control checks can continue.

TCP|$@*|$@*|*|$@*|SMTP$@*|MAIL|tcp_local|*|1|* \
$C$[IMTA_LIB:check_metermaid.so,greylisr,greylisr,$0|$1|$2]\
$N$X4.5.1|Temporary$ failure$ -$ retry$ later$E
```

(Note that the suggested form of the string being used in a greylisting table is *source-ip|env-sender|env-recipient*.)

After the MTA has received the **MAIL FROM:** and **RCPT TO:** SMTP commands, it will use the envelope addresses as well as the source IP address in a **greylisting** call to MeterMaid. Using the table configuration above, MeterMaid will determine whether or not this particular transaction should be permitted. If the MTA should send a temporary rejection at this point, the call to **check_metermaid.so** will succeed, and the **\$N** part of this entry will be returned indicating a rejection. The **\$X4.5.1** flags this rejection as a temporary condition so that a **4xx** SMTP response will be given.

Enhancing Greylisting Functionality

However greylisting is configured, there are additional steps one can take to make further improvements to its functionality. Several possibilities are listed here. They can be used individually or combined together to make more powerful setups.

This section contains the following topics:

- [Preloading the Greylisting Table with Outbound Transactions](#)

- Matching a Range of IP Addresses
- Simplifying the Sender Address
- Providing an Opt-In Mechanism
- Whitelisting Based on User's Addressbook
- Combining Functionality: A Complex Example

Preloading the Greylisting Table with Outbound Transactions

Since it is often the case that one can expect to receive mail from addresses to which the local users are already sending, it may be useful to preload such address combinations into the greylist table. Since the future source IP address is not known, MeterMaid supports a special address of `*` that will match any other supplied address.

To preload the address combinations, one needs to add an entry to the access control mapping table:

```
X-IS_INTERNAL_CHANNEL
```

```
tcp_intranet      $Y
tcp_submit        $Y
tcp_auth          $Y
*                  $N
```

```
ORIG_MAIL_ACCESS
```

```
! For mail that is coming from a local user and going to an external recipient, we
! can save that user/recipient combination and store it into the greylist table
! for
! future permission.
```

```
TCP|$@*|$@*|$@*|$@*|SMTP$@*|MAIL|*|*|tcp_local|*  $C$|X-IS_INTERNAL_CHANNEL;$0|\
$[IMTA_LIB:check_metermaid.so,store,greylist,*|$2|$1,1]
```

(Note that the **store** routine requires a value although it is not used by a greylisting table. Any value may be specified here and is ignored by MeterMaid.)

This mapping table entry first checks to see whether the source channel is considered a channel used by our local users. If the channel is in the list provided by the **X-IS_INTERNAL_CHANNEL** mapping table, then processing continues with the call to the **store** routine of **check_metermaid.so** to store this new combination into the **greylist** table. The combination is stored so that it will match incoming messages from the current recipient going to the current sender, and these messages may come from any source IP address and be permitted.

Matching a Range of IP Addresses

A complication that can occur with greylisting is dealing with remote MTAs that use several different hosts to process deliveries. This can mean that one attempt may occur from 192.168.12.1, but a subsequent attempt may come from a different host like 192.168.12.5. Additional delays may be introduced until an attempt is repeated from a host that had tried it previously.

One way to help address this is to limit IP address matching to the first three octets, allowing more hosts to be considered to be the same source. This would match addresses coming from the same A.B.C.D/24 (class C) subnet. The mapping table setup would be very similar to the examples above, but with a change to the IP address wildcard matching.

```
ORIG_MAIL_ACCESS
```

```
! When checking the source IP address, only use the first three octets in the
string
```

```
! passed to MeterMaid.
```

```
TCP|$@*|$@*|$D*.$D*.$D*.$@*|$@*|SMTP$@*|MAIL|tcp_local|*|1|* \
$C$[IMTA_LIB:check_metermaid.so,greylisting,greylis,$0.$1.$2|$3|$4)\
$N$X4.5.1|Temporary$ failure$ -$ retry$ later$E
```

Simplifying the Sender Address

Some sender addresses will be more complex than a simple `user@example.com` including such features as subaddresses or VERP (variable envelope return path) notation. For more information see https://en.wikipedia.org/wiki/Variable_envelope_return_path. It may be useful to help greylisting recognize the base form of the address using some basic canonicalization in order to keep track of the basic, simplified address form. This simplification can be done by using a nested mapping table call out to perform the canonicalization.

```
X-CORRESPONDENT
```

```
! Subsidiary mapping for removing any subaddress or VERP style material from
! the local-part of an address.
```

```
$_*[$+=-\~]%^@*    $0@$3$Y
*                  $0$Y
```

```
ORIG_MAIL_ACCESS
```

```
TCP|$@*|$@*|*|$@*|SMTP$@*|MAIL|tcp_local|*|1|* \
$C$[IMTA_LIB:check_metermaid.so,greylisting,greylis,$0|$|X-CORRESPONDENT;$1||$2)\
$N$X4.5.1|Temporary$ failure$ -$ retry$ later$E
```

Here, the `X-CORRESPONDENT` table is used to reconstruct the sender address into the simpler `user@example.com` form. The result from this is then used in the call to the `greylisting` function.

Providing an Opt-In Mechanism

Note: This section requires Messaging Server 7 Update 2 or later for the necessary `INCLUDE_SPARES` option.

It may be desirable to allow users to choose whether to have MeterMaid perform greylisting on their incoming mail. This could be especially useful when considering some local mail recipients who may not wish to be subject to delays in receiving incoming mail from unknown senders, such as recipients like **sales** or **customer service**. For these local users, one can set up additional LDAP attributes to be used in conjunction with the existing `ORIG_MAIL_ACCESS` mapping table processing.

For this example, let us assume that one creates a new LDAP attribute `mailUserGreyListOptIn` that will be set to **true** or **false**. Those users who have this attribute set to **true** will have their incoming mail checked with greylisting, while those who have it set for **false** will skip this check and receive their mail immediately.

In order to have the MTA look at this extra attribute, it must be configured into the **option.dat** configuration file.

```
LDAP_SPARE_5=mailUserGreyListOptIn
! Include LDAP_SPARE_5 in ORIG_MAIL_ACCESS probes by setting bit 22 (counting from
0)
! of INCLUDE_SPARES. Bit 22 has the value 4194304.
INCLUDE_SPARES=4194304
```

This will add the value of the user's **mailUserGreyListOptIn** attribute to the probe string used in the **ORIG_MAIL_ACCESS** mapping table.

```
ORIG_MAIL_ACCESS

! This example assumes INCLUDE_SPARES=4194304 is set, so that probes corresponding
! to submissions from remote (tcp_local) senders to local recipients have the
form:
!
! TCP|host-ip|host-port|source-ip|source-port|SMTP-app-info|MAIL|tcp_local|
! remote-sender-address|l|local-recipient-address|recipient-mailUserGreyListOptIn
!
! TCP|$@*|$@*|*$@*|SMTP$@*|MAIL|tcp_local|$_*|l|$_*|true \
$C$[IMTA_LIB:check_metermaid.so,greylisting,greylis,$0|$1|$2|\
$N$X4.5.1|Temporary$ failure$ -$ retry$ later$E
```

The key difference in this mapping table entry is the addition of **l true** to the matching string. Since the **INCLUDE_SPARES** option will append the content of the **mailUserGreyListOptIn** attribute to the probe string, this mapping entry can match against only those where the recipient's **mailUserGreyListOptIn** attribute has been set to **true**, thus skipping others who may be opting out of greylisting.

Whitelisting Based on User's Addressbook

The goal of greylisting is to allow mail from remote senders to local recipients once they are known. In most cases, this happens when a transaction presents this combination on a subsequent delivery attempt. It is also possible to use the recipient's LDAP-based address book to check for the sender's address to determine whether to bypass greylisting for an already recognized address. In order to do this, another LDAP attribute must be added to the **ORIG_MAIL_ACCESS** probe string by including these values into the **option.dat** file:

```
LDAP_SPARE_6=psroot
! Include LDAP_SPARE_6 in ORIG_MAIL_ACCESS probes by setting bit 23 (counting from
0)
! of INCLUDE_SPARES. Bit 23 has the value 8388608.
INCLUDE_SPARES=8388608
```

Furthermore, if appropriate, the MTA's **LDAP_PAB_xyz** options should be set to the proper values for accessing the PAB LDAP server. (However, the usual **pab.*** (Unified Configuration) or **local.service.pab.*** (legacy configuration) settings are usually adequate, and usually do not need to be overridden for MTA purposes via the MTA-specific **LDAP_PAB_xyz** options.)

```
ORIG_MAIL_ACCESS

!
! This example assumes INCLUDE_SPARES=4194304 is set, so that probes corresponding
! to submissions from remote (tcp_local) senders to local recipients have the
form:
!
```

```

! TCP|host-ip|host-port|source-ip|source-port|SMTP-app-info|MAIL|tcp_local|
! remote-sender-address|l|local-recipient-address|recipient-psroot

!
! Matches on this line mean that the sender was found in the recipient's address
book.
! "Whitelist" those addresses, bypassing the greylisting check.
!
TCP|$@*|$@*|*|$@*|SMTP$@*|MAIL|tcp_local|$_*|l|$_*|* \
$C$]pabldap:///3?piEmail1?sub?(|(piEmail1=$1)(piEmail2=$1)(piEmail3=$1))[$E$Y
!
! Now, for all other senders, do the normal greylisting check.
!
TCP|$@*|$@*|*|$@*|SMTP$@*|MAIL|tcp_local|$_*|l|$_*|* \
$C$[IMTA_LIB:check_metermaid.so,greylisting,greylis,$0|$1|$2]\
$N$X4.5.1|Temporary$ failure$ -$ retry$ later$E

```

This mapping table example shows an LDAP callout being done to check to see whether the sender is already known to the local recipient. This allows users to add their correspondents to their address book as a way to whitelist those entries, allowing those senders to bypass the greylisting when sending mail to these local recipients.

Combining Functionality: A Complex Example

It is possible to combine many of these elements together into a much more comprehensive setup. This example makes use of the preloading, opt-in, and address book whitelisting features together.

First, the two LDAP attributes must be available to the mapping table probe. They can be added with these options in **option.dat**:

```

LDAP_SPARE_5=mailUserGreyListOptIn
LDAP_SPARE_6=psroot
! Include LDAP_SPARE_5 and LDAP_SPARE_6 in ORIG_MAIL_ACCESS probes by
! setting bits 22 and 23 (counting from 0) of INCLUDE_SPARES; that is,
! INCLUDE_SPARES=12582912=4194304+8388608=(1<<22)+(1<<23)
INCLUDE_SPARES=12582912

```

Then, the **mappings** file excerpt below shows how the above elements may be combined.

```

! Subsidiary mapping for checking incoming port and channel against a
! list of "internal submission" channels.
!
! Probe format is
!   port.channel
!
X-INTERNAL-CHANNELS

587.tcp_submit    $Y
25.tcp_auth       $Y
25.tcp_intranet   $Y

! Subsidiary mapping for removing any subaddress or VERP style
! material from the local-part of an address.
! This mapping also performs LDAP URL style quoting of the retained
! portion of the address.
!
X-CORRESPONDENT

```

```

    $_*$[+=\~]%^@*    $=$0@$3$_$Y
    *                  $=$0$_$Y

ORIG_MAIL_ACCESS

! This example assumes INCLUDE_SPARES=12582912 (or some superset of bits) is
! set, so that probes corresponding to submissions from local senders to
! remote recipients have a form of:
!
! TCP|host-ip|host-port|source-ip|source-port|SMTP-app-info|MAIL|source-channel|
! local-sender-address|tcp_local|remote-recipient-address|
! sender-mailUserGreyListOptIn|sender-psroot
!
! while probes corresponding to SMTP MAIL submissions from remote
! (tcp_local) senders to local recipients have the form:
!
! TCP|host-ip|host-port|source-ip|source-port|SMTP-app-info|MAIL|tcp_local|
! remote-sender-address|l|local-recipient-address|
! recipient-mailUserGreyListOptIn|recipient-psroot
!
! The overall logic includes pre-population of the "greylist" table at (1)
! with *|remote-correspondent|local-user on outgoing messages from local users
! who have opted-in to greylisting (have mailUserGreyListOptIn: true) (0),
! and then checks of incoming messages to local users who want greylisting (2)
! against:
!   (i) the local-user's PAB (3)
!   (ii) the pre-populated entries in the "greylist" table (4) or (5)
!   (iii) the "greylist" table tracking "recent" submission attempts from
!         not-otherwise-known (not pre-populated due to local-user sending
!         to them, nor recognized in local-user's PAB) remote senders (4) or (5)
! Note that (ii) and (iii) are done by one probe to the greylist table, as
! MeterMaid first performs the (ii) check automatically due to the format of
! probe. This probe the greylist table is either done at (4) (for IPv4
! source IPs) or at (5) (for IPv6 source IPs).
!
! For outgoing messages, from local users to remote correspondents,
! pre-populate the "greylist" table using the "store" entry point with
! *|simplified-quoted-remote-correspondent|local-user
! This is so that replies from that remote-correspondent (from whatever
! source-IP) to that local-user will be accepted.
! The subsidiary mapping table X-INTERNAL-CHANNELS is used to check
! (based on the host-port and source-channel) whether the message is
! one from a local user to a remote correspondent. The subsidiary
! mapping table X-CORRESPONDENT is used to canonicalize the
! recipient-address.
! (0)
!
TCP|$@*|*|$@*|$@*|SMTP$@*|MAIL|*|*|tcp_local|*|true|* \
    $C$|X-INTERNAL-CHANNELS;$0.$1|PREPOPULATE|$|X-CORRESPONDENT;$3||$2
!
! If the message was indeed from a local user who wants grey-listing,
! then the above entry matched and reset the probe to now be:
! PREPOPULATE|quoted-simplified-remote-correspondent|local-user
! Then the entry below pre-populates the greylist table with an
! entry for
! *|quoted-simplified-remote-correspondent|local-user
! (1)
!
PREPOPULATE|*|* \
    $C$[IMTA_LIB:check_metermaid.so,store,greylis,*|$0|$1,1]$E

```

```

!
! For incoming submission attempts from remote correspondents (tcp_local
! submission attempts):
! (2) Entry matches remote senders to recipients that
! want grey-listing (mailUserGreyListOptIn: true). Entry constructs
! a new GREYLIST... probe retaining relevant fields, namely:
! GREYLIST|source-ip|quoted-simplified-sender|recipient|psroot
! where the quoted-simplified-sender field is processed (simplified
! and LDAP quoted) using the subsidiary X-CORRESPONDENT mapping table
! (3) For the recipients who want grey-listing, the probe is now
! GREYLIST|source-ip|quoted-simplified-sender|recipient|psroot
! Look up the (simplified) sender address in the
! recipients PAB. Accept submission if found, fall through otherwise.
! (4) If the sender address wasn't found at (3), fall-through and now
! attempt a MeterMaid "greylist" table lookup. The probe to this
! "greylist" table will have the form:
! source-IP-subnet|quoted-simplified-sender|recipient
! This entry matches on IPv4 incoming source IPs, and ignores the last
! eight bits to give an IPv4 subnet.
! Because the probe has the form A|B|C, and it is a probe to a
! greylisting endpoint, MeterMaid will automatically initially attempt
! a *|B|C probe, only bothering with the A|B|C probe if its initial,
! automatic probe fails. Thus any pre-populated, generic source IP
! entry will match first, prior to MeterMaid attempting a lookup of
! the specific source IP subnet. If the specific triad is
! found in the "greylist" table as being due to be greylisted (rejected
! temporarily), then the probe "succeeds" -- set a new probe string
! FIRSTATTEMPT and continue so that (6) will match and the greylist
! response will be issued.
! Otherwise, the MeterMaid probe "fails" -- as for the cases
! where the probe matches a "good" (pre-populated, or resubmitted
! after block_time) entry.
! (5) The same as (4), but matching on IPv6 incoming source IPs, ignoring
! the last 64 bits to give an IPv6 subnet.
! (6) Issue the temporary rejection when the greylist probe of (4) or (5)
! "succeeded".
!
! For remote senders (source channel tcp_local) to local recipients with the
! LDAP_SPARE_5 attribute "true", reset the probe to the GREYLIST|...form
! (2)
!
! TCP|$@*|$@*|$@*|SMTP$@*|MAIL|tcp_local|$_*|1|$_*|true|* \
! $CGREYLIST|$0|$X-CORRESPONDENT;$1||$2|$=$3$_
!
! If a recipient wants grey-listing, the probe has been rebuilt to be:
! GREYLIST|ip-source|simplified-sender-address|recipient-address|psroot
! where simplified-sender-address omits any subaddress/VERP-y sorts of fluff
! and has had any LDAP URL required quoting applied, and where psroot has also
! had any LDAP URL required quoting applied.
! So next check whether the simplified-sender-address can be found in
! the recipient-address user's PAB (found under psroot); if the sender is
! found, then accept this message -- this sender is "known".
! (3)
!
! GREYLIST|*|*|* \
! $C$j]pabldap:///3?piEmail1?sub?((piEmail1=$1)(piEmail2=$1)(piEmail3=$1))[$E$Y
!
! Otherwise, if the sender was not known to this recipient, then fall down
! to the subsequent entry which performs the MeterMaid grey-list check.
!
!

```

```

! Match on IPv4 addresses and probe the greylist table.
! If this sender matches an entry in the greylist table, whether
! pre-populated or due to a recent sending attempt, then let their message in.
! If the greylist table probe says the sender needs greylisting,
! (that is, never seen before, or seen before but only within block_time),
! then continue with the probe changed to "FIRSTATTEMPT" so that it'll fall
! through and match the temporary rejection entry below at (6).
! Otherwise, if the sender was in the greylist table but after block_time,
! then MeterMaid "fails" this probe, so the check ends; the table processing
! "falls-through" and, if no other entry matches, the submission is
! permitted.
! (4)
!
GREYLIST|$D*.$D*.$D*.$D*|*|*|* \
$[IMTA_LIB:check_metermaid.so,\
greylisting,greylist,$0.$1.$2|$4|$5]$CFIRSTATTEMPT
!
! (5)
!
GREYLIST|$H*:$H*:$H*:$H*:$@H*:$@H*:$@H*:$@H*|*|*|* \
$[IMTA_LIB:check_metermaid.so,\
greylisting,greylist,$0:$1:$2:$3|$4|$5]$CFIRSTATTEMPT
!
! Must be a "new" sending attempt -- give it a temporary rejection
! (6)
!
FIRSTATTEMPT $N$X4.5.1|Temporary$ failure$ -$ retry$ later

```

Mapping Table Notes

The mapping tables above use several features that may be unfamiliar to casual users of the MTA's mapping table, including these:

Table 15-3 UPDATE TABLE

Strings	Explanation
\$@	Disables saving the following wildcard match that will not be needed for right-hand side substitutions. This permits sufficient saved wildcards (of which there can be at most ten) to be available for matching fields of more interest.
\$_	Specifies "non-greedy" (minimal) matching of the portion of the string; used for the local-part of the sender address prior to the first occurrence of a special character possibly indicating a subaddress or VERP address variation.
\$[+=\-%	Matches an occurrence of any one of the specified characters. Note that the hyphen character must be quoted with a backslash character to be interpreted as a literal hyphen character rather than indicating a character range.
\$D*	Matches only decimal digits; used for parsing IP addresses.
\$H*	Matches only hexadecimal digits; used for parsing IPv6 addresses.

MeterMaid Reference

This document contains MeterMaid reference information and contains the following sections:

- [configutil Options](#)
- [Table Types](#)
- [check_metermaid.so Reference](#)

configutil Options

For MeterMaid-specific **configutil** options, you can search for options containing **metermaid** in the *Messaging Server Reference*.

Table Types

There are two possible table types allowed by the **metermaid.table.*.type** options:

- [greylisting Tables](#)
- [simple Tables](#)
- [throttle Tables](#)

greylisting Tables

Greylisting tables may be used to provide an anti-spam/anti-virus technique. More information about setting up these tables can be found at the "[Implementing Greylisting by Using MeterMaid](#)" page.

simple Tables

A simple table may be used to store arbitrary data referenced by a key. The key data type is defined by **metermaid.table.tablename.data_type**, and the value data type is defined by **metermaid.table.tablename.value_type**. Some operations are only available to those simple tables where the value data type is **integer**.

throttle Tables

Throttle tables are used to specify a particular "hit count" *quota* over **quota_time** seconds to limit connections, transactions, or certain other components of incoming connections. MeterMaid automatically maintains the count over time, decrementing it back down after **quota_time** has passed.

check_metermaid.so Reference

The **check_metermaid.so** shared library is traditionally used to throttle incoming connections in a mapping table such as **PORT_ACCESS** or **MAIL_ACCESS**. We now use a set of new routines to use data stored in **simple** tables (please refer to the **msconfig metermaid.table.*_type** option, or the **configutil metermaid.table.*_type** option). These routines now permit one to store, retrieve, and test arbitrary data stored in MeterMaid's ephemeral data store. Oracle Communications Messaging Server also uses the **greylisting** routine that works with **greylisting** tables.

Table 16–1 shows the routines available in **check_metermaid.so**, which of the two table types, **simple** and/or **throttle**, are supported for those routines, and a brief description of each one. Below that, detailed information about each routine is provided.

Note: If any error should occur during processing such as a failure to communicate with MeterMaid, or if invalid options are provided to these routines, the routines will simply return FALSE and the shared library callout will fail.

Table 16–1 *check_metermaid.so Available Routines*

Routine	Description
adjust	Adds to or subtracts from an integer value in a table
adjust_and_test	Performs an adjust and then returns the result of a test operation
fetch	Returns a value from the table
greylisting	Returns TRUE if we are temporarily rejecting this transaction
remove	Removes an entry from a table
store	Stores a value into the table
test	Tests an integer value with a simple comparison, returning TRUE or FALSE
throttle	Increments a "hit count" and returns TRUE if quota_count is exceeded

The following sections provide a description of each routine, a table showing the options used by the routine, the value returned, if any, to the calling environment, and sample usage.

- [adjust Routine](#)
- [adjust_and_test Routine](#)
- [fetch Routine](#)
- [greylisting Routine](#)
- [remove Routine](#)
- [store Routine](#)
- [test Routine](#)
- [throttle Routine](#)

adjust Routine

This section describes the **adjust** routine.

Description

The **adjust** routine allows you to make a numeric adjustment to an integer value in a **simple** table.

Table 16–2 *adjust Routine Options*

Option	Description
<i>table</i>	Table in which key is found.
<i>key</i>	Key corresponding to the value being adjusted.
<i>adjustment</i>	Positive or negative value to be added to the value.

adjust Routine Return Value

Returns TRUE with the new value for **key** after the adjustment has been made as the resultant string.

adjust works by taking the current integer value for *key* in the table called *table* and adding *adjustment* to it, then storing the resulting value back. *adjustment* can be negative, thus reducing the value which can be negative.

If *key* doesn't exist in *table*, it will be presumed to have an initial value of 0 and *key* will be stored into *table* with a new value of *adjustment*.

Table 16–3 *adjust Routine Supported Table and Value Types*

Table	Value	Supported ?
greylisting	-	-
simple	integer	X
simple	string	-
throttle	-	-

Example

The following example shows that the current value for **fred@example.org** will be increased by 35 after the **adjust** is completed. If **fred@example.org** did not exist in the **scores** table, it would be stored with a value of 35.

```
$/opt/sun/comms/messaging/lib/check_metermaid.so,adjust,scores,fred@example.org,+35]
```

adjust_and_test Routine

This section describes the **adjust_and_test** routine.

Description

The **adjust_and_test** routine allows you to make a numeric adjustment to an **integer** value in a **simple** table, and then test that value against a provided comparator.

Table 16–4 *adjust_and_test Routine Options*

Option	Description
<i>table</i>	Table in which key is found
<i>key</i>	Key corresponding to the value being adjusted
<i>adjustment</i>	Positive or negative value to be added to the value
<i>comparator</i>	Comparison symbol(s) '<', '>', and/or '=', followed by a numeric value

adjust_and_test Routine Return Value

Returns TRUE if the comparison is true, and FALSE otherwise; no resultant string is returned.

adjust_and_test first takes the current integer value for *key* in the table called *table* and adding *adjustment* to it, storing the resulting value back. The routine then compares the resulting value against the *comparator*, returning the result.

If *key* doesn't exist in *table*, it will be presumed to have an initial value of 0 and *key* will be stored into *table* with a new value of *adjustment*, and the comparison will be made against *adjustment*.

Table 16–5 *adjust_and_test Routine Supported Table and Value Types*

Table	Value	Supported?
greylisting	-	-
simple	integer	X
simple	string	-
throttle	-	-

Example

The following example shows that the current value for **fred@example.org** will be decreased by 2 after the **adjust** is completed. If **fred@example.org** did not exist in the **scores** table, it would be stored with a value of -2. Then this new value is checked to see if it is greater than or equal to 20, returning TRUE if it is.

```
$[/opt/sun/comms/messaging/lib/check_metermaid.so,adjust_and_test,scores,fred@example.org,-2,>=20]
```

fetch Routine

This section describes the **fetch** routine.

Description

The **fetch** routine retrieves a value from a **simple** table.

Table 16–6 *fetch Routine Options*

Option	Description
table	Table in which key is found
key	Key corresponding to the value being returned

fetch Routine Return Value

Returns TRUE if *key* exists and returns its value as the resultant string, otherwise returns FALSE.

fetch retrieves the value associated with *key* in *table* and returns it as the resultant string. If *key* does not exist in *table*, then FALSE is returned and no resultant string is available.

Table 16–7 *fetch Routine Supported Table and Value Types*

Table	Value	Supported?
greylisting	-	X
simple	integer	X
simple	string	X
throttle	-	-

Example

The following example retrieves the current score for **fred@example.org** which can then use that information for subsequent processing, such as in a mapping table.

```
$/opt/sun/comms/messaging/lib/check_metermaid.so,fetch,scores,fred@example.org]
```

greylisting Routine

The **greylisting** routine is used to validate entries in the table based on time and resubmission attempts. This is used as part of a Greylisting setup (see "[Implementing Greylisting by Using MeterMaid](#)" for details).

Table 16–8 *greylisting Routine Options*

Option	Description
<i>table</i>	Table in which <i>key</i> is checked/stored
<i>key</i>	Key corresponding to the value being tested for greylisting.

greylisting Routine Return Value

Returns TRUE if this probe should cause a temporary rejection for the submission attempt, otherwise returns FALSE to permit the attempt.

greylisting first probes *table* for *key* to see whether this entry has been previously permitted. If it is found to be allowed, FALSE is returned to permit the submission. Then it checks to see whether *key* exists and is in its resubmission period (specified by the **resubmit_time** table option). If so, *key* is marked as valid and is permitted (returning FALSE). If *key* exists, but is in the **block_time** period, the submission is refused as validation occurs after **block_time** has passed, and **greylisting** returns TRUE to return a temporary rejection for the attempt. Lastly, if *key* does not exist, it is stored into *table* and **greylisting** returns TRUE to return a temporary rejection for this new attempt.

Example

This example shows a greylisting probe for mail from **barney@example.com** going to local user **fred@example.org**. If this returns TRUE, then the mapping code should

return a temporary rejection so that the message submission should be reattempted later.

```
$(IMTA_LIB:check_metermaid.so,greymailing,greymail_
table,192.168.10.34|barney@example.com|fred@example.org]
```

remove Routine

This section describes the **remove** routine.

Description

The **remove** routine removes an entry from table.

Table 16–9 *remove Routine Options*

Option	Description
<i>table</i>	Table in which <i>key</i> is found.
<i>key</i>	Key corresponding to the value being removed.

remove Routine Return Value

Returns TRUE if *key* was removed from *table*, otherwise returns FALSE.

When an entry in a table is no longer needed, it may be removed using *remove*. Subsequent attempts to access *key* will result in its value not being found.

Supported Table and Value Types

Table 16–10 *remove Routine Supported Table and Value Types*

Table	Value	Supported?
greymailing	-	X
simple	<i>integer</i>	X
simple	string	X
throttle	-	X

Example

The following example can be used when the record for **fred@example.org** is no longer needed.

```
$/opt/sun/comms/messaging/lib/check_metermaid.so,remove,scores,fred@example.org]
```

store Routine

This section describes the **store** routine.

Description

The **store** routine is used to store a new value into the table.

Table 16–11 *store Routine Options*

Option	Description
<i>table</i>	Table into which the new value is to be stored.

Table 16–11 (Cont.) store Routine Options

Option	Description
<i>key</i>	Key corresponding to the value.
<i>value</i>	New value.

store Routine Return Value

Returns TRUE if the new value was stored, FALSE otherwise. Returns no resultant string.

store is similar to **adjust** in that it can be used to put data into a table. Unlike **adjust**, however, any previous value that may exist is overwritten by **store**. Also, in addition to integer data, strings may also be stored into those tables that permit it. For a **greylisting** table, *value* is ignored and instead the new *key* is stored into *table* as a valid entry for subsequent queries.

Supported Table and Value Types**Table 16–12 store Routine Supported Table and Value Types**

Table	Value	Supported?
greylisting	-	X
simple	integer	X
simple	string	X
throttle	-	-

Example

The following example sets an initial value into a table that can be used by subsequent **fetch** operations.

```
$/opt/sun/comms/messaging/lib/check_metermaid.so,store,loginhosts,barney@example.org,quarry.example.org]
```

test Routine

This section describes the **test** routine.

Description

The **test** routine allows you to compare an integer value in a **simple** table against a supplied comparator.

Table 16–13 test routine Options

Option	Description
<i>table</i>	Table in which <i>key</i> is found
<i>key</i>	Key corresponding to the value being adjusted
<i>comparator</i>	Comparison symbol(s) '<', '>', and/or '=', followed by a numeric value

test Routine Return Value

Returns TRUE if the comparison is true, and FALSE otherwise; no resultant string is returned.

test takes the current integer value for *key* in the table called *table* and compares the resulting value against the *comparator*, returning the result. If *key* does not exist in *table*, then 0 is used as the value to be compared.

Table 16–14 test Routine Supported Table and Value Types

Table	Value	Supported?
greylisting	-	X
simple	integer	X
simple	string	-
throttle	-	X

Example

The following example tests the number of login attempts made to see whether it exceeds a defined threshold.

```
$[/opt/sun/comms/messaging/lib/check_metermaid.so,test,logins,wilma@example.org,>5]
```

throttle Routine

This section describes the **throttle** routine.

Description

The **throttle** routine is used to count incoming connections or transactions over a period of time enforcing a quota limit.

Table 16–15 throttle Routine Options

Option	Description
<i>table</i>	Table in which is holding the items being counted
<i>key</i>	Key corresponding to the particular "hit count" to be incremented

throttle Routine Return Value

Returns TRUE if **quota** has been exceeded during the past **quota_time** seconds, otherwise returns FALSE.

For more detailed information on setting up **throttle** tables with configuration examples, refer to the *Messaging Server Security Guide*.

Table 16–16 throttle Routine Supported Table and Value Types

Table	Value	Supported?
greylisting	-	-
simple	integer	-
simple	string	-
throttle	-	X

Message Tracking and Recall

Oracle Communications Messaging Server includes a general message tracking and recall facility that conforms to RFCs 3885-3887. This includes support for the SMTP Extension for Message Tracking, Message Tracking Query Protocol (MTQP), and the Submission Future Release Extension (RFC 4865). The client provides tracking ID and a secret and future release delay during which the message can be recalled from the mail queue.

The recall aspect is implemented as a tracking extension. Additionally, some tracking extensions have been implemented to make it possible for a tracking client to track and recall messages submitted by a non-tracking client.

Table 17-1 shows the MTA and channel options used for and relevant to message tracking and recall.

Table 17-1 MTA and Channel Options

Option	Description
tracking_mode <i>integer</i>	The tracking_mode MTA option controls how message tracking information is stored by the MTA. The default value of 0 disables storage of tracking information. Setting tracking_mode to 1 enables storage using a memcached server shared across the deployment. The use of other values is currently restricted. Note that the memcache protocol is implemented by many database systems. Any implementation that supports memcached semantics can be used as an alternative to memcached . Default 0.
tracking_retries <i>integer</i>	Tracking information updates can fail because of simultaneous access attempts to the underlying database. If this happens the update can be retried. The tracking_retries MTA option specifies how many times to retry the update. Default 5.
tracking_retry_delay <i>integer</i>	The tracking_retry_delay MTA option specifies the amount of time to delay between retry attempts in units of centiseconds. Default 10.
tracking_debug <i>integer</i>	The tacking_debug MTA option is used to enable debug output from the MTA's tracking subsystem. The default value of 0 disables debug output. Positive values enable it. The larger the value, the more output that is produced. Default 0.

Table 17-1 (Cont.) MTA and Channel Options

Option	Description
log_tracking <i>integer</i>	<p>The log_tracking MTA option controls logging of message tracking and recall information. Bit 0 (value 1), if set, includes the tracking ID and the current tracking in transaction log entries. The two appear as a single field. The ID appears first and is separated from the timeout value by a colon. Tracking information appears immediately after the envelope ID and before the deferred delivery time. An mt attribute is used in the XML log format. If bit 1 (value 2) is set in the log_tracking MTA option, then the tracking ID and timeout, again separated by a colon, appear in the LOG_ACTION mapping table probe, again immediately after the envelope ID and before the deferred delivery time.</p> <p>Default 0.</p>
log_futurerelease	<p>The log_futurerelease MTA option logs the values associated with the FUTURERELEASE SMTP extension. Setting bit 0 (value 1) causes the future release value, expressed as an offset in seconds from the current time, to be logged immediately after the mailbox UID is logged. An fr attribute is used in the XML log format. If bit 1 (value 2) is set in the log_futurerelease MTA option, then this information appears in the LOG_ACTION mapping table probe immediately after the mailbox UID.</p>
nottrackingclient, nottrackingserver, trackingclient, trackingserver	<p>These channel options control the availability and use of the Message Tracking SMTP extension specified in RFC 3885. Availability of the extension is enabled with the trackingserver source channel option. The nottrackingserver channel option disables the availability of the extension, and is the default.</p> <p>The SMTP client's use of the extension is controlled by the trackingclient and nottrackingclient channel options. The former enables use of the extension. The latter disables it and is the default.</p> <p>Note that the extension must be enabled on all internal SMTP clients and servers throughout a deployment in order for tracking and recall to work across that deployment.</p>
trackingtimeoutdefault <i>integer</i>	<p>The trackingtimeoutdefault option specifies the default timeout that's applied if the MTRK option is issued to the SMTP server without a timeout value.</p> <p>Default 259200 (3 days).</p>
trackingtimeoutmax <i>integer</i>	<p>The trackingtimeoutmax option specifies the maximum allowed timeout value in seconds. Any value greater than the maximum is silently lowered to the maximum.</p> <p>Default 1209600 (14 days).</p>
trackingtimeoutmin <i>integer</i>	<p>The trackingtimeoutmin option specifies the minimum allowed timeout value in seconds. Any value less than the minimum is silently raised to the minimum.</p> <p>Default 86400 (1 day).</p>
trackinginternal, trackingrelayed, trackingdelivered	<p>The handling of messages relayed internal to a deployment, including internal channel hops, needs to be distinguished from the case where messages leave the administrative domain for tracking and recall to work properly. However, the SMTP protocol is commonly used in both cases.</p> <p>Additionally, the case where a successful channel dequeue results in message delivery also needs to be distinguished from dequeues where this does not occur.</p> <p>Three channel options are provided to specify these semantics: trackinginternal, trackingrelayed, and trackingdelivered. trackinginternal, the default, specifies that the message is being transferred internally. trackingrelayed specifies that the channel transfers messages to some external system. trackingdelivered specifies that the channel performs final delivery of the message.</p>

Table 17-1 (Cont.) MTA and Channel Options

Option	Description
trackingmultiple, trackingsingle, trackingfirst	<p>RFC 3885 specifies how the Message Tracking SMTP Extension interacts with aliases and mailing lists. In particular, it says, "MTAs MUST NOT copy MTRK certifiers when a recipient is aliased, forwarded, or otherwise redirected and the redirection results in more than one recipient. However, an MTA MAY designate one of the multiple recipients as the "primary" recipient to which tracking requests shall be forwarded; other addresses MUST NOT receive tracking certifiers. MTAs MUST NOT forward MTRK certifiers when doing mailing list expansion."</p> <p>This arguably makes sense for tracking-only applications where presenting the results of a complex alias expansion process to the end user may be confusing. However, the situation with message recall is different. Users expect recall to work when feasible, including when alias expansion is involved. (Mailing lists are different; a mailing list effectively "owns" its messages once it expands, so recall past a mailing list expansion is inappropriate.)</p> <p>Accordingly, three source channel options are provided to control the MTA's behavior in this regard. trackingmultiple, the default, tells the MTA to pass tracking ID/timeout information to all recipients of an alias expansion. trackingsingle causes tracking ID/time information to pass through only when there is a single recipient. trackingfirst causes tracking information to pass through to the first alias expansion recipient. (Note that in the case of aliases stored in LDAP, the first recipient is unpredictable.)</p>
trackinggenerate <i>n</i>	<p>Message tracking and recall depends on the generation, attachment, and transfer of tracking identifiers. Such identifiers are normally generated and attached to messages by the submitting client. However, essentially no clients currently support the generation of such identifiers, making it impossible to write a separate tracking/recall client to deal with messages submitted by a non-tracking-enabled client. Additionally, a user who elects to use multiple clients, some tracking-enabled and some not, will end up with only a subset of their messages able to be tracked and recalled.</p> <p>Tracking identifiers can also provide, independent of their use for user tracking and recall, a stable identifier that ties MTA log entries across multiple systems together in ways that envelope IDs and Message-ID header fields do not and cannot. As such, automatic assignment of tracking identifiers to message on ingress as well as submission has real utility independent of user tracking and recall functions.</p> <p>The trackinggenerate source channel option addresses these needs. A single required integer option specifies the default tracking timeout. If set, a tracking identifier for the message is generated in one of two ways:</p> <ul style="list-style-type: none"> ■ If authentication has been used and the user has a general recall secret associated with their LDAP entry, a per-message recall secret is generated by computing a SHA-1 hash of the concatenation of the content of Message-ID: header field, the Date: header field (if present), and the general recall secret. The tracking timeout is controlled by the trackinggenerate value. ■ If authentication was not used or no general recall secret is associated with the account, a tracking identifier is created by hashing a unique identifier with an MD4 hash. Note that the security of this process is controlled not by the randomness of the unique identifier or the use of MD4, but rather by the infeasibility of computing X given T, where $SHA1(X) = MD4(T)$.
ldap_attr_auth_recall_secret	<p>The ldap_attr_auth_recall_secret MTA option specifies the name of the LDAP attribute where a user's general recall secret is stored. This option has no default.</p>
memcache_host	<p>Specifies the name of the host running either memcached or a memcached-compatible server where tracking and recall information will be stored.</p> <p>Note that each tracked message creates a single entry with the tracking ID as the key. The size of the entries is highly dependent on the number of recipients. Messages with large numbers of recipients can result in very large entries. This option has no default.</p>
memcache_port	<p>Specifies the memcached port on the host specified by memcache_host. Default 11211.</p>



Administering Event Notification Service

This information describes what you need to do to enable the Event Notification Service Publisher (ENS Publisher) and administer Event Notification Service (ENS) in Unified Configuration.

Topics:

- [ENS Publisher in Messaging Server](#)
- [Configuring the ENS Publisher in Unified Configuration](#)
- [Administering Event Notification Service](#)

For more information on ENS and ENS APIs, see "[Messaging Server Specific Event Notification Service Information](#)" and "[Event Notification Service C API Reference](#)."

ENS Publisher in Messaging Server

The Event Notification Service (ENS) is the underlying publish-and-subscribe service. ENS acts as a dispatcher used by Communications Suite applications as a central point of collection for certain types of *events* that are of interest to them. Events are changes to the value of one or more properties of a resource. Any application that wants to know when these types of events occur registers with ENS, which identifies events in order and matches notifications with subscriptions. ENS and the ENS publisher are bundled with Oracle Communications Messaging Server.

Configuring the ENS Publisher in Unified Configuration

ENS has the following default behavior:

- ENS is enabled by default. The initial configuration sets the **ens.enable** option to 1.
- No configuration is required to load the ENS publisher because the **ms-internal** instance is automatically loaded and configured. Therefore, you do not need to create a separate **ms-internal** instance.
- If you want to configure options for the pre-loaded **ms-internal** default instance, set them with the **ms-internal** instance name. For example, **notifytarget:ms-internal.settings**.
- IMAP IDLE now only works using ENS, because the ability to use IMAP IDLE with JMQ has been removed.
- The **notifytarget:ms-internal.enshost** defaults to **base.listenaddr** if it is not set.

Administering Event Notification Service

Administering ENS consists of starting and stopping the service, and changing the options to control the behavior of the ENS publisher.

Topics in this section:

- [Starting and Stopping ENS](#)
- [Event Notification Service Configuration Options](#)
- [ENS SSL Support](#)

Starting and Stopping ENS

If desired, you can use the **start-msgens** and **stop-messageens** commands to start and stop the ENS server.

Event Notification Service Configuration Options

The **notifytarget:target.*** options control the behavior of the publisher. Use the **msconfig set** command to set these options. For a list of options, see the *Messaging Server Reference*.

- To enable ENS, make sure that the **ens.enable** option is set to **1**, for example:

```
/opt/sun/comms/messaging64/bin/msconfig set ens.enable 1
/opt/sun/comms/messaging64/bin/msconfig show ens.enable
role.ens.enable = 1
```

ENS SSL Support

SSL support has been added to ENS in a separate default port 8997. New options have been added to support this feature.

To enable or disable SSL support for ENS:

- **ens.enablesslport** (Unified Configuration)
- **local.ens.enablesslport** (legacy configuration)

To change the ENS **sslport**:

- **ens.sslport** (Unified Configuration)
- **local.ens.sslport** (legacy configuration)

To add **sslnicknames**:

- **ens.ssnicknames**(Unified Configuration)
- **local.ens.ssnicknames** (legacy configuration)

To make a notification target to use TLS/SSL:

- **notifytarget:target-name.ensusesssl** (Unified Configuration)
- **local.store.notifyplugin.target-name.ensusesssl** (legacy configuration)

Both SSL and non-SSL ports can be enabled for ENS at the same time. The ENS notification targets can use TLS/SSL for its communication with the ENS broker.

The notification targets of type, ENS, can use TLS/SSL to communicate with the ENS broker specified.

The default value of the option **notifytarget:target-name.ensusesssl** will be 1, if **ens.ENABLESSLPORT** is 1 and one of the following conditions is satisfied:

- The **notifytarget** is the **ms-internal** plugin
- The value of the **notifytarget:target-name.enshost** is not set
- The value of **notifytarget:target-name.enshost** is equal to the value of **service.listenaddr**
- The value of **notifytarget:target-name.enshost** is the loopback address, "127.0.0.1" or "::1"

If the **ensusesssl** option of the **notifytarget** is set, then the TLS/SSL will be used to communicate with the host defined by the options **ensHost** and **ensPort**, in the **notifytarget** plugin.

The default value of the option:

- **notifytarget:target-name.ensport** (Unified Configuration)
- **local.store.notifyplugin.target-name.ensport** (legacy configuration)

will be equal to the value of **ens.sslport** if the value of **notifytarget:target-name.ensusesssl** is 1. Otherwise, it will be equal to the value of **ens.port**.

ENS Support for Password Based Authentication

Messaging Server has added options to support password based authentication to the ENS server.

1. Option to Enable/Disable authentication (**ens.mustauthenticate**).
2. Option to change the secret for authentication (**ens.secret**).
3. Option to specify username for the ENS **notifyplugin** (**notifytarget:target-name.ensuser**)
4. Option to specify password for the ENS **notifyplugin** (**notifytarget:target-name.enspwd**)

Password based authentication of the ENS clients to the broker is enabled, by setting the option, **mustauthenticate**.

The option:

local.ens.mustauthenticate (legacy configuration)

or

ens.mustauthenticate (Unified Configuration)

enables or disables whether authentication is required by the ENS broker. The default value of the **ens.mustauthenticate** option is 0.

If **mustauthenticate** option is set, authentication is required by the ENS broker in both SSL and non-SSL ports.

The ENS broker accepts any user name but the password for authentication is set by the option:

local.ens.secret (legacy configuration)

or

ens.secret (Unified Configuration)

There is no default value for **ens.secret**. If **mustauthenticate** is set, authentication is required by the ENS Broker on both SSL and non-SSL ports. The password for authentication can be set with the **ens.secret** option. All connections to the ENS Broker will fail unless the **ens.secret** is set with a password.

The notification targets of type ENS can be made to use password based authentication.

The userid for authentication is set using the option:

local.store.notifyplugin.target-name.ensuser (legacy configuration)

or

notifytarget:target-name.ensuser (Unified Configuration)

The default value of the option, **ensuser** is "guest."

The password for authentication to ENS broker in a notification target is set using the option:

local.store.notifyplugin.target-name.enspwd (legacy configuration)

or

notifytarget:target-name.enspwd (Unified Configuration)

There is no default value for **enspwd**. The value of the option **enspwd** will be equal to the value of option **ens.secret**, if one of the following conditions satisfies:

1. The notification is the **ms-internal** plugin.
2. The value of **notifytarget:target-name.enshost** is NULL.
3. The value of **notifytarget:target-name.enshost** is equal to the value of **service.listenaddr**
4. The value of **notifytarget:target-name.enshost** is the loopback address, "127.0.0.1" or "::1".

If the **ensuser** and **enspwd** are provided, then the **notifytarget** figures out whether the ENS broker that it connects to require password based authentication or not. If the ENS broker that the notify target connects requires a password, then the password provided will be used or else it won't be used.

With the newer version of the ENS broker that uses authentication with **ens.mustauthenticate** set to 1, you must set a password using the **ens.secret** option. Otherwise all connections to the ENS broker will fail. If authentication is disabled with **ens.mustauthenticate** set to 0, the older version of the ENS broker which does not have authentication will be used. By default, authentication is disabled.

Note: If you use the older ENS Client APIs with the newer ENS broker (i.e. authentication enabled), it will not work. When authentication is enabled, using ENS will require setting of the option, **local.ens.secret** and use of newer API, **ens_sopen**.

Messaging Server Specific Event Notification Service Information

This information describes the Oracle Communications Messaging Server specific items you need to use the ENS APIs.

Topics:

- [Event Notification Types and Options](#)
- [Implementation Notes](#)

Event Notification Types and Options

For Messaging Server, there is only one event reference, which can be composed of several options. There are various types of event notifications. [Table 19–1](#) lists the event types supported by Messaging Server and gives a description of each. Event notifications are also generated when a user creates, deletes, or renames a folder.

Event Types

Table 19–1 *Event Types*

Event Types	Description
ChangeFlag	Shows change status as "1" add, "2" remove, or "3" replace.
DeleteMsg	Messages marked as "Deleted" are removed from the mailbox. This is the equivalent to IMAP expunge.
Login	User logged in from IMAP, HTTP, or POP.
Logout	User logged out from IMAP, HTTP, or POP.
NewMsg	New message was received by the system into the user's mailbox. Can have a payload of message headers and body.
OverQuota	Operation failed because the user's mailbox exceeded one of the quotas (diskquota, msgquota). The MTA channel holds the message until the quota changes or the user's mail box count goes below the quota. If the message expires while it is being held by the MTA, it will be expunged.
PurgeMsg	Message expunged (as a result of an expired date) from the mailbox by the server process imexpire. This is a server side expunge, whereas DeleteMsg is a client side expunge. This is not a purge in the true sense of the word.
ReadMsg	Message in the mailbox was read (in the IMAP protocol, the message was marked Seen).

Table 19–1 (Cont.) Event Types

Event Types	Description
TrashMsg	Message was marked for deletion by IMAP or HTTP. The user may still see the message in the folder, depending on the mail client's configuration. The messages are to be removed from the folder when an expunge is performed.
UnderQuota	Quota went back to normal from OverQuota state.
UpdateMsg	Message was appended to the mailbox (other than by NewMsg). for example, the user copied an email message to the mailbox. Can have a payload of message headers and body.

The following applies to the above supported event types:

- For **NewMsg** and **UpdateMsg**, message pay load is turned off by default to prevent overloading ENS. For information on how to enable the payload, see "Payload." No other event types support a payload.
- Event notifications can be generated for changes to the **INBOX** alone, or to the **INBOX** and all other folders. The following configuration variable allows for **INBOX** only (value = 0), or for both the **INBOX** and all other folders (value = 1):

```
local.store.notifyplugin.noneInbox.enable
```

The default setting is for **INBOX** only (value = 0).

Note: There is no mechanism to select folders; all folders are included when the variable is enabled (value = 1).

- The **NewMsg** notification is issued only after the message is deposited in the user mailbox (as opposed to "after it was accepted by the server and queued in the message queue").
- Every notification carries several pieces of information (called options) depending on the event type, for example, **NewMsg** indicates the IMAP **uid** of the new message. For details on the options each event type takes, see "Available Options for Each Event Type."
- Events are not generated for POP3 client access.
- All event types can be suppressed by issuing **XNOTNOTIFY**. For example, an IMAP script used for housekeeping only (the users are not meant to be notified) might issue it to suppress all events.

Options

iBiff uses the following format for the ENS event reference:

```
enp://127.0.0.1/store_?param=_value&param1=_value1&param2=_value2_
```

The event key **enp://127.0.0.1/store** has no significance other than its uniqueness as a string. For example, the hostname portion of the event key has no significance as a hostname. It is simply a string that is part of the URI. However, the event key is user configurable. The list of iBiff event reference options is listed in tables "Mandatory Event Reference Options" and "Optional Event Reference Options" that follow.

The second part of the event reference consists of option-value pairs. This part of the event reference is separated from the event key by a question mark (?). The option and value are separated by an equals sign (=). The option-value pairs are separated by an ampersand (&). Note that there can be empty values, for which the value simply does not exist.

Table 19–2 describes the mandatory event reference options that need to be included in every notification.

Mandatory Event Reference Options

Table 19–2 Mandatory Event Reference Options

Option	Data Type	Description
evtType	string	Specifies the event type.
hostname	string	The hostname of the machine that generated the event.
mailboxName	string	Specifies the mailbox name in the message store. The mailboxName has the format uid@domain , where uid is the user's unique identifier, and domain is the domain the user belongs to. The @domain portion is added only when the user does not belong to the default domain (i.e. the user is in a hosted domain).
pid	integer	ID of the process that generated the event.
process	string	Specifies the name of the process that generated the event.
timestamp	64-bit integer	Specifies the number of milliseconds since the epoch (midnight GMT, January 1, 1970).

Table 19–3 describes optional event reference options, which might be seen in the event depending on the event type (see "Available Options for Each Event Type.")

Optional Event Reference Options

Table 19–3 Optional Event Reference Options

Option	Data Type	Description
client	IP address	The IP address of the client logging in or out.
diskQuota	signed 32-bit integer	Specifies the disk space quota in kilobytes. The value is set to -1 to indicate no quotas.
diskUsed	signed 32-bit integer	Specifies the amount of disk space used in kilobytes.
hdrLen	unsigned 32-bit integer	Specifies the size of the message header. Note that this might not be the size of the header in the payload, because it might have been truncated.
imapUid	unsigned 32-bit integer	Specifies the IMAP uid option.
lastUid	unsigned 32-bit integer	Specifies the last IMAP uid value that was used.
numDel	unsigned 32-bit integer	Specifies the number of messages marked as deleted in the mailbox.
numMsgs	unsigned 32-bit integer	Specifies the number of total messages in the mailbox.
numMsgsMax	signed 32-bit integer	Specifies the quota for the maximum number of messages. The value is set to -1 to indicate no quotas.

Table 19–3 (Cont.) Optional Event Reference Options

Option	Data Type	Description
numSeen	unsigned 32-bit integer	Specifies the number of messages in the mailbox marked as seen (read).
size	unsigned 32-bit integer	Specifies the size of the message. Note that this may not be the size of payload, since the payload is typically a truncated version of the message.
uidValidity	unsigned 32-bit integer	Specifies the IMAP uid validity option.

Note: Subscribers should allow for undocumented options when parsing the event reference. This allows for future compatibility when new options are added.

Table 19–4 shows the options that are available for each event type. For example, to see which options apply to a **TrashMsg** event, look in the column header for "ReadMsg, TrashMsg" and then note that these events can use **numDel**, **numMsgs**, **numSeen**, and **userValidity**.

Available Options for Each Event Type

Note: Oracle reserves the right to change "no" to "yes" at any time needed.

Table 19–4 Available Options for Each Event Type

Option	NewMsg, UpdateMsg	ReadMsg, TrashMsg	DeleteMsg, PurgeMsg	Login, Logout	OverQuota, UnderQuota
client	No	No	No	Yes	No
diskQuota	No	No	No	No	Yes
diskUsed	No	No	No	No	Yes
hdrLen	Yes	No	No	No	No
imapUid	Yes	No	Yes	No	No
lastUid	No	No	Yes	No	No
numDel	No	Yes	No	No	No
numMsgs	Yes	Yes	Yes	No	Yes
numMsgsMax	No	No	No	No	Yes
numSeen	No	Yes	No	No	No
size	Yes	No	No	No	No
uidValidity	Yes	Yes	Yes	No	No
userid	No	No	No	Yes	No

Payload

ENS allows a payload for two event types: **NewMsg**, and **UpdateMsg**; the other event types do not carry a payload. The payload portion of these two notifications can contain any of the following data:

- No header or body data (default setting)
- Message header data only
- Message body data only
- Both message header and body data

The amount and type of data sent as the payload of the ENS event is determined by the configuration options found in "[Payload Configuration Options](#)."

Payload Configuration Options

Table 19–5 describes the payload configuration options.

Table 19–5 Payload Configuration Options

Configuration Options	Description
local.store.notifyplugin.*.maxbodysize	Specifies the maximum size (in bytes) of the body that will be transmitted with the notification. Syntax: uint32 Default: 0
local.store.notifyplugin.*.maxheadersize	Specifies the maximum size (in bytes) of the header that will be transmitted with the notification. Syntax: uint32 Default: 0

Note that both options are set to zero as the default so that no header or body data is sent with ENS notifications.

Examples

The following example shows a **NewMsg** event reference (it is actually a single line that is broken up to several lines for readability):

```
enp://127.0.0.1/store?evtType=NewMsg&timestamp=1047488403000&
hostname=eman&process=imta&pid=476&mailboxName=testuser&numMsgs=16
&uidValidity=1046993605&imapUid=62&size=877&hdrLen=814
```

In this example, for the **DeleteMsg** event. Messages marked as deleted by IMAP or HTTP were expunged. The user would not see the message in the folder any more.

```
enp://127.0.0.1/store?evtType=DeleteMsg&timestamp=1047488588000&
hostname=eman&process=imapd&pid=419&mailboxName=testuser&
numMsgs=6&uidValidity=1046993605&imapUid=61&lastUid=62
```

And a third example shows a **ReadMsg** event. Message was marked as Seen by IMAP or HTTP.

```
enp://127.0.0.1/store?evtType=ReadMsg&timestamp=1047488477000&
hostname=eman&process=imapd&pid=419&mailboxName=testuser&
uidValidity=1046993605&numSeen=11&numDel=9&numMsgs=16
```

Implementation Notes

The current implementation does not provide security on events that can be subscribed to. Thus, a user could register for all events, and portions of all other users' mail. Because of this it is strongly recommended that the ENS subscriber be on the "safe" side of the firewall at the very least.

The ENS server supports two options to control TCP access to the ENS server (**service.ens.domainallowed** and **service.ens.domainnotallowed**). These options work the same way as the equivalent options for POP, IMAP, and HTTP. See the discussion on configuring client access to POP, IMAP, and HTTP services in the *Messaging Server Security Guide*. These options replace the functionality of the `ENS_ACCESS` environment variable that was included in the legacy ENS server.

Event Notification Service C API Reference

This information details the new ENS C API.

Topics:

- [API Overview](#)
- [API Basic Usage](#)
- [API Usage Notes](#)

API Overview

The ENS C API, *ens.h*, is located in the *MessagingServer_home/examples/enssdk/* directory. The **ens_pub.c** sample publisher and **ens_sub.c** sample subscriber demonstrate use of the ENS C API.

Here is the API header (**ens.h**):

```
====
// ens.h -- ENS C client API
//
// Copyright (c) 2009, 2010, Oracle and/or its affiliates. All rights reserved.

#ifndef ENS_HEADER_INCLUDED
#define ENS_HEADER_INCLUDED 1

#ifdef __cplusplus
extern "C" {
#endif

//
// Connecting
//

// an ENS client
struct ensclient_t;
typedef struct ensclient_t ensclient_t;

// callback invoked if the ENS connection dies
// subscriptions are no longer valid but must not be unsubscribed
typedef void (*lost_cnx_cb_t)(void*);

ensclient_t* ens_open(const char* host, int port, lost_cnx_cb_t
lost_cnx_cb, void* lost_cnx_arg);

// automatically cleans up all existing subscription handles
void ens_close(ensclient_t*);
```

```
//
// Publishing events
//

typedef void (*destructor_t)(void*);           // cleanup function

void ens_publish(ensclient_t*, const char* evt, char* body, size_t
bodysz, destructor_t body_delete);

//
// Receiving events
//

// a subscription handle
struct sub_t;
typedef struct sub_t subscription_t;

// handler called when a subscribed event is received
typedef void (*notify_cb_t)(void *rock, char *event, char *body, size_t
body_len);

subscription_t* ens_subscribe(ensclient_t*, const char* evt, notify_cb_t
cb, void* rock);
void ens_unsubscribe(ensclient_t*, subscription_t*);

#ifdef __cplusplus
}
#endif

#endif // ENS_HEADER_INCLUDED
====
```

API Basic Usage

The client calls **ens_open()** to start a connection. If reliability across ENS connection outages is important, the client should provide a lost connection handler callback. The lost connection handler normally marks any client-specific subscription information as invalid, calls **ens_close**, and triggers a task to attempt a reconnect by using **ens_open** (possibly after a delay).

The client calls **ens_close()** on shut down.

The client calls **ens_subscribe()** to subscribe to events and gets a callback when a matching event is received. The client calls **ens_unsubscribe()** to unsubscribe from an event (usually not necessary as a client can just call **ens_close**).

To publish an event, use **ens_publish**. (In general, you do not need to do so and the sample code should be sufficient.)

To build the sample programs, link against the **libens** library, which is normally installed in **/opt/sun/comms/messaging64/lib/libens.so**.

The **ens_sub.c** sample program is helpful to see what events are generated and how the event strings and message payloads are formatted.

Both Oracle Communications Messaging Server publishers (that is, **imapd**) and the ENS server (**enpd**) are designed to drop events if an overload situation occurs.

New Client API `ens_sopen`

New Client API, `ens_sopen`, can connect to the ENS broker with authentication.

Customers who want to connect to the ENS broker use the API `ens_open` declared in `ens.h` to create a new client connection to the ENS broker. The API, `ens_open`, does not support authentication and TLS/SSL.

So we added a new API, `ens_sopen`, to create a secure connection to the ENS broker that supports authentication and TLS/SSL. The arguments to the API, `ens_sopen`, include all the arguments to the API, `ens_open`, and also includes arguments to accept an username and a password for authentication. It also includes an argument to specify whether to use TLS/SSL while making an connection to the ENS Broker specified.

The new API declared in `lib/ens/ens.h` is:

```
ensclient_t* ens_sopen (const char* host, int port, int use_ssl, const char* auth_user,  
const char* auth_secret, lost_cnx_cb_t lost_cnx_cb, void* lost_cnx_arg);
```

In order to connect to the ENS brokers that support TLS/SSL and authentication, you must use the new API, `ens_sopen`. The API accepts the arguments `auth_user` and `auth_secret`, but it may or may not use them depending on whether or not the ENS broker it is connecting to requires authentication.

If the connection is made to the SSL port, then the value for the argument, `use_ssl`, must be 1. If the connection is made to the non-SSL port, then the value for the argument, `use_ssl`, must be 0.

API Usage Notes

The ENS C API is presently the recommended API for C-based software that needs to subscribe to Messaging Server events. Use of the Glassfish Message Queue, OpenMQ, or Java Enterprise System Message Queue C API is not recommended.

Configuring IMAP IDLE

Topics:

- [Benefits of Using IMAP IDLE](#)
- [Configuring IMAP IDLE with ENS in Unified Configuration](#)

Benefits of Using IMAP IDLE

The IMAP IDLE extension to the IMAP specification, defined in RFC 2177, allows an IMAP server to notify the mail client when new messages arrive and other updates take place in a user's mailbox. The IMAP IDLE feature has the following benefits:

- Mail clients do not have to poll the IMAP server for incoming messages. Eliminating client polling reduces the workload on the IMAP server and enhances the server's performance. Client polling is most wasteful when a user receives few or no messages; the client continues to poll at the configured interval, typically every 5 or 10 minutes.
- A mail client displays a new message to the user much closer to the actual time it arrives in the user's mailbox. A change in message status is also displayed in near-realtime. The IMAP server does not have to wait for the next IMAP polling message before it can notify the client of a new or updated mail message. Instead, the IMAP server receives a notification as soon as a new message arrives or a message changes status. The server then notifies the client through the IMAP protocol.

Configuring IMAP IDLE with ENS in Unified Configuration

IMAP IDLE with ENS has the following default behavior:

- ENS is enabled by default. The initial configuration sets the **ens.enable** option to 1.
- Every message store has its own **enpd** server.
- The **imapd** process, store delivery channels, and store utilities report changes to the **enpd** server on the local store.
- Some additional configuration is helpful for improved security, HA, and flag updates, as explained in "[To Configure IMAP IDLE with ENS.](#)"
- IMAP IDLE does not require that events be aggregated to a single **enpd** server and the IDLE event distribution is more efficient if each store uses its own **enpd** server.

Prerequisites for Configuring IMAP IDLE with ENS

Make sure ENS is enabled by setting the **ens.enable** option to 1:

```
msconfig set ens.enable 1
```

To Configure IMAP IDLE with ENS

1. Configure the **enpd** server to allow (or restrict) connections only from the hosts running the message stores by configuring the **ens.domainallowed** and **ens.domainnotallowed** options as necessary. For example, the following command allows access to the local host only:

```
msconfig set ens.domainallowed enpd:127.0.0.1
```

The following command allows access to the local host and all IP addresses 192.168.0.* except 192.168.0.17:

```
msconfig set ens.domainallowed '"enpd:192.168.0.0/255.255.255.0,127.0.0.1 EXCEPT 192.168.0.17"'
```

These options work the same way as the equivalent options for POP, IMAP, and HTTP. These options replace the functionality of the **ENS_ACCESS** environment variable that was included in the legacy ENS server.

2. Stop, then restart Oracle Communications Messaging Server.

```
cd /opt/sun/comms/messaging/bin
stop-msg
start-msg
```

3. Verify that the IMAP services now include the IDLE feature. Use telnet to connect to the IMAP host and port.

```
telnet <IMAP_hostname> <port>
```

Example:

```
telnet myhost imap
trying 192.18.01.44 ...
connected to myhost.example.com
* OK [CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ NAMESPACE UIDPLUS
CHILDREN BINARY UNSELECT SORT LANGUAGE STARTTLS IDLE XSENDER X-NETSCAPE
XSERVERINFO X-SUN-SORT X-SUN-IMAP X-ANNOTATEMORE AUTH=PLAIN]
myhost.example.com IMAP4 service (Oracle Communications Messaging Server
7u5-4.07 64bit (built Mar 21 2012)
```

To Disable IMAP IDLE

- To disable IMAP IDLE, set the **ens.enable** option to 0 (default is 1). For example:

```
msconfig set ens.enable 0
```

Lemonade Profile 1 Support

Absent from Oracle Communications Messaging Server's support for Lemonade Profile 1 is SMTP BINARYMIME. In addition, Messaging Server's CONDSTORE and ANNOTATE implementations might cause performance issues, so use caution when working with these features. See the appropriate sections in the following information for details.

Topics:

- [Introduction to Lemonade](#)
- [Lemonade Features](#)
- [Support for BURL](#)
- [IMAP URLAUTH Support](#)
- [IMAP CATENATE Support](#)
- [IMAP Conditional Store Operation Support](#)
- [IMAP ANNOTATE Support](#)
- [Controlling IMAP CAPABILITIES Vector](#)
- [Support for SMTP Submission Service Extension for Future Message Release](#)

Introduction to Lemonade

Lemonade refers to an IETF working group formed to address the requirements of supporting standards-based email in a mobile or other resource-constrained environment. A "resource-constrained" environment is one where any or all of the following might be encountered:

- Low bandwidth, high latency networks
- Intermittent network connectivity
- Scarce power and compute cycles
- Minimizing data usage is a goal

The Lemonade Profile (RFC 4550, <http://tools.ietf.org/html/rfc4550>) defines a set of IMAP and SMTP extensions that address these constraints. Messaging Server implements most of the extensions defined in RFC 4550 (Lemonade Profile 1) and some of the extensions defined in RFC 5550 (Lemonade Profile 2). This information describes the configurable extensions.

Note: The Lemonade standard is mostly intended for mobile clients. Its goal is to reduce network traffic (both in volume and number of interactions) and to move the CPU load from the client to the server. Clients must have support for Lemonade built into them.

Lemonade Features

Some of the more interesting features of Lemonade include the following:

- Forward a message without download (enabled by CATENATE, URLAUTH, and BURL)
- Quick resync (enabled by CONDSTORE and QRESYNC)
- Persistent sort and search (enabled by CONTEXT)
- Conversion

The following sections describe these features in more detail.

Support for BURL

Note: Refer to the discussion on BURL support for SMTP SUBMIT in the *Messaging Server Reference* for details on configuring and using BURL.

Messaging Server supports the BURL command, which extends the SMTP submission profile by adding a new command to fetch submission data from an IMAP server. This permits a mail client to inject content from an IMAP server into the SMTP infrastructure without downloading it to the client and uploading it back to the server. Thus, you could forward an email message without first downloading it to the client.

For more information, see <http://www.ietf.org/rfc/rfc4468.txt>.

Support is enabled in Messaging Server by the BURL_ACCESS mapping. The mapping receives two different probe strings:

```
port_access-probe-info|channel|uid|
port_access-probe-info|channel|uid|url
```

Here **port_access-probe-info** consists of all the information usually included in a PORT_ACCESS mapping table probe. It will be blank if BURL is being used in a "disconnected" context such as batch SMTP. The **channel** is the current source channel and **uid** is the user's authenticatd UID. The **uid** will be blank if no authentication has been performed. The **:\$S** input flags will be set if SASL authentication has been performed and **:\$T** will be set if TLS is in use.

The first probe is done when responding to EHLO. In order to offer BURL support the mapping must set **\$Y** and optionally provide a space-separated list of supported URL types. The mapping assumes **imap** if no string is returned.

The second probe is performed when a BURL command is actually sent by the submit client. It includes the URL specified in the BURL command. Additionally, **:\$I** will be set if the URL contains any vertical bars (which if present could possibly confuse some sorts of access checks). The mapping must set **\$Y** for the URL to be accepted for processing. If **\$D** is also set the string result of the mapping replaces the originally specified URL.

At an absolute minimum the mapping must verify that a proper type of URL has been specified. Typically only **imap:** URLs should be allowed. Additionally, in the case of "submit" IMAP URLs, a check needs to be made to insure that the URL belongs to the user, that is, the access user in the URL matches the authenticated UID for the submit session. Additionally, it is almost always essential to restrict access to an appropriate set of IMAP servers.

The default BURL settings for Unified Configuration are the following:

BURL_ACCESS

```
*|tcp_*|%*|                                imap$Y
*|@*|imap://*;URLAUTH=submit+$1*:*   $:A$M$Y
```

The SMTP server has to have the ability to log in to the IMAP server as the submit user. The **imap_username** and **imap_password** MTA options are used to accomplish this. **imap_username** specifies the submit user and defaults to the setting of the **imap.submituser** option if not specified. The **imap_password** option specifies the password which of course much match the value set for the submit user account. The **imap_password** option has no default value.

IMAP URLAUTH Support

Messaging Server supports the URLAUTH extension to IMAP and the IMAP URL Scheme (IMAPURL). This extension provides a means by which an IMAP client can use URLs carrying authorization to access limited message data on the IMAP server. An IMAP server that supports this extension indicates this with a capability name of "URLAUTH."

For more information, see <http://www.ietf.org/rfc/rfc4467.txt>.

IMAP CATENATE Support

Messaging Server supports the CATENATE extension to IMAP, which extends the APPEND command to allow clients to create messages on the IMAP server that may contain a combination of new data along with parts of (or entire) messages already on the server. Using this extension, the client can catenate parts of an already existing message onto a new message without having to first download the data and then upload it back to the server.

For more information, see <http://www.ietf.org/rfc/rfc4469.txt>.

IMAP Conditional Store Operation Support

Messaging Server supports IMAP Conditional Store Operations (CONDSTORE). IMAP CONDSTORE enables clients to coordinate changes to a common IMAP mailbox, for example, when multiple users are accessing shared mailboxes. The Conditional Store facility provides a protected update mechanism for message state information that can detect and resolve conflicts between multiple writing mail clients. The Conditional Store facility also allows a client to quickly resynchronize mailbox flag changes.

Warning: Use caution when enabling CONDSTORE with Convergence, as there might be degradation of IMAP performance. Our hope is to fix this problem in a future release. When this happens, we will remove this caution.

For more information, see <http://www.ietf.org/rfc/rfc4551.txt>.

IMAP ANNOTATE Support

Messaging Server supports the ANNOTATE extension to IMAP, which permits clients and servers to maintain "meta data" for messages, or individual message parts, stored in a mailbox on the server. For example, you could use IMAP ANNOTATE to attach comments and other useful information to a message, or to attach annotations to specific parts of a message, marking them as seen or important, or a comment added.

Warning: Use caution when enabling ANNOTATE, as there might be degradation of IMAP performance. Our hope is to fix this problem in a future release. When this happens, we will remove this caution.

For more information, see <http://www.ietf.org/rfc/rfc5257.txt>. Of note in this document is Section 3.4, "Access Control," which summarizes access control restrictions, including the new ACL "n" right.

Controlling IMAP CAPABILITIES Vector

When migrating a multi-system deployment from Messaging Server 6.3 to 7, it is important that the systems advertise consistent IMAP extension sets, especially with respect to **CONDSTORE**. During migration you can configure Messaging Server 7 to display the same capability set as the older Messaging Server version. You can also turn on the new features on all back systems simultaneously. This feature is only of real significance with **CONDSTORE** and if you have lemonade-aware clients.

You can also filter the initial capability vector advertised in the IMAP banner in the MMP.

Set the IMAP capability options to **1** to control the **CAPABILITIES** vector. To see a list of these options, run the following command:

```
msconfig
msconfig> help option capability_*
```

You can also refer to the *Messaging Server Reference* to see the IMAP capability options.

The default values for most of these options is **1**. The exceptions are **IMAP4** and **CONDSTORE**. The default for **IMAP4** is 0 unless **obsoleteimap** is set, in which case it is **1**. These options only affect whether a particular feature is advertised, except for **imap.capability_condstore** which also enables the feature.

Turning on (or not turning off) a capability does not necessarily mean that the feature will be advertised. **IDLE**, **STARTTLS**, and **XREFRESH** are only advertised if enabled by these options and other condition exist that make it appropriate for them to be advertised.

Support for SMTP Submission Service Extension for Future Message Release

Messaging Server supports Lemonade Profile 1, which is an extension to the SMTP submission protocol for a client to indicate a future time for the message to be released for delivery. This extension permits a client to use server-based storage for a message that should be held in queue until an appointed time in the future. This is useful for

clients that do not have local storage or are otherwise unable to release a message for delivery at an appointed time. This functionality is useful for sending announcements to be read at the beginning of a work day, to send birthday greetings a day or so ahead, or to use as a lightweight facility to build a personal reminder service.

For more information, see <http://tools.ietf.org/rfc/rfc4865.txt>.

Support is enabled in Messaging Server by placing the **futurerelease** channel option on the source channel used for initial message submission. The option takes a single integer argument: the maximum number of seconds a message can be held.

Managing Logging

This information provides overview information on the logging facilities for the Oracle Communications Messaging Server MTA, the Message Store, and services. This information also provides procedures for how to manage these logging facilities.

Topics:

- [Overview of Logging](#)
- [Managing MTA Message and Connection Logs](#)
- [Managing Message Store, Admin, and Default Service Logs](#)
- [Using Message Store Log Messages](#)
- [MMP Logging](#)

Overview of Logging

This section contains the following subsections:

- [What Is Logging and How Do You Use it?](#)
- [Types of Logging Data](#)
- [Types of Messaging Server Log Files](#)
- [Tools for Managing Logging](#)
- [Tracking a Message Across the Various Log Files](#)

What Is Logging and How Do You Use it?

Logging is the means by which a system provides you with time-stamped and labeled information about the system's services. Logging provides both a current snapshot of the system as well as a historical view.

By understanding and using Messaging Server log files, you can:

- Gather message statistics, such as message size, rate of message delivery, and how many messages are passing through the MTA
- Perform trend determination
- Correlate capacity planning
- Troubleshoot problems

For example, if your site needs to add more disk storage due to an increase in the number of users, you can use the Messaging Server log files to see what percentage

your system demand has increased by and plan for the amount of new disk storage you need.

You can also use Messaging Server logs to understand what your messaging pattern looks like across one day. Understanding when your daily peak loads occur helps you conduct capacity planning.

Logging is also helpful for troubleshooting user problems. For example, if a user isn't receiving expected mail messages, you can use the Messaging Server logging facilities to trace the user's mail messages. In so doing, you might find out that the messages didn't arrive because they were automatically filtered and sent to a SPAM folder.

Types of Logging Data

In general, logging provides you with two types of information:

- Operational data
- Error conditions, also known as event logging

For the most part, Messaging Server logging provides operational data. This operational data contains information such as: the date and time a message entered the system; the sender and recipient of the message; when the message was written to disk; and at a later point in time, when the message was removed from disk and inserted into user's mailbox.

However, Messaging Server logging does also provide some event logging data. To obtain event logging data, you need to pull together multiple items from different log files. You could then use a unique constant, such as message ID, to search and correlate the life cycle of a message as it passed from point to point through the system.

Types of Messaging Server Log Files

Messaging Server logging consists of three types of log files:

1. *MTA logs*. These logs provide operational data previously described for the Message Transfer Agent.
2. *Error logs*. These are the MTA debug logs, and the MTA subcomponent logs (that is, job controller, dispatcher and so on).
3. *Message Store and Service logs*. These logs provide messages from the http server, mshttpd, imap, and pop services, as well as the Admin service. The format of these logs differs from that of the first two types of logs.

Table 23–1 lists the different types of log files. By default, log files are located in the `MessagingServer_home/data/log` directory. You can customize and view each type of log file individually.

Table 23–1 Messaging Server Log Files

Type of Log File	Log File Description	Default Name
Message Transfer Agent	Show information about message traffic through the MTA including date and time information, enqueue and dequeue information, and so on.	mail.log_current , mail.log_yesterday , mail.log
Connections	Contains remote machines (MTAs) that connect to this system to send email.	connection.log

Table 23–1 (Cont.) Messaging Server Log Files

Counters	Contains message trends in terms of messages sent and received on a per channel basis.	counters
Job Controller	Contains data on the master, job controller, sender, and dequeue channel programs.	job_controller.log
Dispatcher	Contains errors pertaining to the dispatcher. Turning on dispatcher debugging will increase the information.	dispatcher.log
Channel	Records errors pertaining to the channel. Channel options master_debug and slave_debug turn on channel debugging, which increases the verbosity of the channel log files. Level and type of information is controlled with the various *_DEBUG MTA options.	<i>channel-name_master.log*</i> (example: tcp_local_master.log*) <i>channel-name_slave.log*</i> (example: tcp_local_slave.log*)
IMAP	Contains logged events related to IMAP4 activity of this server	imap, imap.sequenceNum.timeStamp
POP	Contains logged events related to POP3 activity of this server	pop, pop.sequenceNum.timeStamp
HTTP	Contains logged events related to HTTP activity of this server	http, http.sequenceNum.timeStamp
Default	Contains logged events related to other activity of this server, such as command-line utilities and other processes	default, default.sequenceNum.timeStamp
msgtrace	Contains trace information for the Message Store. File can grow very large very quickly. Monitor accordingly.	msgtrace
watcher	Monitors process failures and unresponsive services (see Table 3–6, "Services Monitored by watcher and msprobe") and will log error messages indicating specific failures.	watcher

where:

sequenceNum - Specifies an integer that specifies the order of creation of this log file compared to others in the log-file directory. Log files with higher sequence numbers are more recent than those with lower numbers. Sequence numbers do not roll over. They increase monotonically for the life of the server (beginning at server installation).

timeStamp - Specifies a large integer that specifies the date and time of file creation. (Its value is expressed in standard UNIX time: the number of seconds since midnight January 1, 1970.)

For example, a log file named **imap.63.915107696** would be the 63rd log file created in the directory of IMAP log files, created at 12:34:56 PM on December 31, 1998.

The combination of open-ended sequence numbering with a timestamp gives you more flexibility in rotating, expiring, and selecting files for analyzing. For more specific suggestions, see "Defining and Setting Service Logging Options."

Tools for Managing Logging

You can customize the policies for creating and managing Messaging Server log files by using the **msconfig** command.

For Message Store, the settings you specify affect which and how many events are logged. You can use those settings and other characteristics to refine searches for logged events when you are analyzing log files.

The MTA uses a separate logging facility you configure MTA logging by specifying information in configuration files.

For log analysis and report generation beyond the capabilities of Messaging Server, you need to use other tools. You can manipulate log files on your own with text editors or standard system tools.

With a scriptable text editor supporting regular-expression parsing, you can potentially search for and extract log entries based on any of the criteria discussed in this information, and possibly sort the results or even generate sums or other statistics.

In UNIX environments you might also be able to modify and use existing report-generation tools that were developed to manipulate UNIX **syslog** files. If you wish to use a public-domain **syslog** manipulation tool, remember that you might need to modify it to account for the different date/time format and for the two extra components (*facility* and *logLevel*) that appear in Messaging Server log entries but not in **syslog** entries.

Tracking a Message Across the Various Log Files

The following describes how a message flows through the system, and at what point information gets written to the various log files. This description is meant to aid you in your understanding of how to use Message Server's log files to troubleshoot and resolve problems. See [Figure 5-2, "MTA Architecture"](#) to follow along.

1. A remote host makes a connection to the TCP socket on your messaging host, requesting SMTP service.
2. The MTA dispatcher responds to the request, and hands off the connection to your messaging host's SMTP service. As the MTA is modular in design, it consists of a set of processes, including the job controller and the SMTP service dispatcher. The dispatcher takes the incoming TCP connection and sends it to the SMTP service. The SMTP service writes the message to disk to a channel area. The SMTP service understands the message's envelope options, such as sender and recipient. Configuration entries in the system tell what destination channel it belongs to.
3. The dispatcher writes to the **dispatcher.log** file that it forked a thread and made the thread available to incoming connection from a certain IP address.
4. The SMTP server writes to its **tcp_smtp_server.log** file, recording the dialog of what happens when the remote host connected to it and sent a message. This log file gets created when dispatcher hands off to SMTP server on the host's IP.
5. The SMTP server writes the message to a queue area on disk for a channel program such as **tcp_intranet**, and informs the job controller.
6. The job controller contacts the channel program.
7. The channel program delivers the message. Each channel has its own log file. However, these logs usually show the starting and stopping of the channel. To get more information, you need to enable debug level for the channel. However, as this can slow down your system and actually make problems more obscure if left on, you should only enable debug level when an actual problem is occurring.

Note: For efficiency, if a channel is already running for an existing process, and a new message comes in, the system does not spawn a new channel process. The currently running process picks up the new message.

8. The message is delivered to its next hop, which could be another host, another TCP connection, and so forth. This information is written to the **connection.log** file when **SEPARATE_CONNECTION_LOG** is enabled. At the same time that the SMTP server writes the message to a queue area on disk, the channel responsible for the message writes a record in the **mail.log_current** file. The record shows such information as the date and time the message was enqueued, the sender, the recipient, so forth. See "MTA Message Logging Examples" for more information. The most useful file for tracing the message is the **mail.log_current** file.

Managing MTA Message and Connection Logs

The MTA provides facilities for logging each message as it is enqueued and dequeued. It also provides dispatcher error and debugging output.

This section consists of the following subsections:

- [Understanding the MTA Log Entry Format](#)
- [Enabling MTA Logging](#)
- [Specifying Additional MTA Logging Options](#)
- [MTA Message Logging Examples](#)
- [Enabling Dispatcher Debugging](#)

You can control logging on a per-channel basis or you can specify that message activity on all channels be logged. In the initial configuration, logging is disabled on all channels.

See "Enabling MTA Logging" for more information.

Enabling logging causes the MTA to write an entry to the *MessagingServer_home/data/log/mail.log_current* file each time a message passes through an MTA channel. Such log entries can be useful for gathering statistics on how many messages are passing through the MTA (or through particular channels). You can also use these log entries to investigate other issues, such as whether and when a message was sent or delivered.

The message return job, which runs every night around midnight, appends any existing **mail.log_yesterday** to the cumulative log file, **mail.log**, renames the current **mail.log_current** file to **mail.log_yesterday**, and then begins a new **mail.log_current** file. The message return job also performs the analogous operations for any **connection.log*** files.

While the MTA performs automatic rollovers to maintain the current file, you must manage the cumulative **mail.log** file by determining policies for tasks such as backing up the file, truncating the file, deleting the file, and so on.

When considering how to manage the log files, note that the MTA periodic return job will execute a site-supplied *MessagingServer_home/data/site-programs/bin/daily_cleanup* script, if one exists. Thus some sites might choose to supply their own cleanup procedure that, for instance, renames the old **mail.log** file once a week (or once a month), and so on.

Note: With logging is enabled, the **mail.log** file steadily grows and, if left unchecked, consumes all available disk space. Monitor the size of this file and periodically delete unnecessary contents. You can also delete the entire file as another version will be created as needed.

Understanding the MTA Log Entry Format

The MTA log file is written as ASCII text. By default, each log file entry contains eight or nine fields as shown in the example below.

```
16-Feb-2007 14:54:13.72 tcp_local ims-ms EE 1 adam@example.com
rfc822;marlowe@example.org marlowe@ims-ms-daemon
```

The log entry shows:

1. The date and time the entry was made (in the example, 16-Feb-2007 14:54:13.72).
2. The channel name for the source channel (in the example, **tcp_local**).
3. The channel name for the destination channel (in the example, **ims-ms**). For SMTP channels, when **LOG_CONNECTION** is enabled, a plus (+) indicates inbound to the SMTP server; a minus (-) indicates outbound via the SMTP client.
4. The type of entry (in the example, **EE**). Entries can consist of a single action code (see [Table 23–2, "Logging Entry Action Codes"](#)) or an action code and one or more modifier codes (see [Table 23–3, "Logging Entry Modifier Codes"](#)). The format for entries is as follows

<action_code><zero or more optional modifiers>

For example a logging entry code of **EEC** means that the email was Enqueued (action-code **E**) using ESMTP (modifier **E**) and SMTP Chunking (modifier **C**). Please refer to the tables below for details on the currently used action and modifier codes.

5. The size of the message (in the example, **1**). This is expressed in kilobytes by default, although this default can be changed by using the **BLOCK_SIZE** MTA option. The SMS channel can be configured to log a page count rather than file size in this field. See [Table 27–5, "SMS Channel Options"](#) for information on **LOG_PAGE_COUNT**.
6. The envelope **From:** address (in the example, **adam@example.com**). Note that for messages with an empty envelope **From:** address, such as notification messages, this field is blank.
7. The original form of the envelope **To:** address (in the example, *marlowe@example.org*).
8. The active (current) form of the envelope **To:** address (in the example, **marlowe@ims-ms-daemon**).
9. The delivery status (SMTP channels only).

The following three tables describe the logging entry codes.

Table 23–2 Logging Entry Action Codes

Entry	Description
-------	-------------

Table 23–2 (Cont.) Logging Entry Action Codes

B	Bad command sent to the SMTP server. The recipient address field will contain the command that was rejected while the diagnostic field will contain the response the SMTP server gave. MTA channel option, <code>MAX_B_ENTRIES</code> , controls how many bad commands will be logged in a given session. Default is 10.
D	Successful dequeue
E	Successful enqueue
J	Rejection of attempted enqueue (rejection by slave channel program)
K	Recipient message rejected. If the sender requests <code>NOTIFY=NEVER</code> DSN flag set or if the message times out or if the message is manually returned (for example: <code>imsimta qm "delete"</code> command always generates a "K" record for each recipient, while a <code>qm "return"</code> command will generate a "K" record rather than an "R" record). This indicates that there was no notification sent to the sender per the sender's own request. This can be compared with "R" records, which are the same sort of rejection/time-out, but where a new notification message (back to the original sender) is also generated regarding this failed message.
Q	Temporary failure to dequeue
R	Recipient address rejected on attempted dequeue (rejection by master channel program), or generation of a failure/bounce message
S	LMTP deposit into the message store. This action code is used on the LMTP server side.
V	Warning message that will appear whenever a transaction is abnormally aborted. There will be one "V" record per enqueued recipient address.
W	Warning message sent to notify original sender that the message has not been delivered yet, but it is still in the queue being retried.
Z	Some successful recipients, but this recipient was temporarily unsuccessful; the original message file of all recipients was dequeued, and in its place a new message file for this and other unsuccessful recipients will be immediately enqueued

Table 23–3 describes the logging entry modifier codes.

Table 23–3 Logging Entry Modifier Codes

Entry	Description
A	SASL authentication used.
B	SMTP BINARYMIME extension used (RFC 3030).
C	Chunking was used. Note that ESMTP has to be used for chunking to work, so you'll typically see field values like <code>EEC</code> or <code>DEC</code> .
E	An EHLO command was issued/accepted and therefore ESMTP was used.
L	LMTP was used.
Q	SMTP PIPELINING extension used (RFC 2920).

Table 23-3 (Cont.) Logging Entry Modifier Codes

S	TLS/SSL used. S transaction log entries now increment the various submitted message counters associated with the channel.
U	BURL used (RFC 4468).

If **LOG_CONNECTION** is enabled, then an additional set of action codes will be used. See the discussion on **LOG_CONNECTION** in the *Messaging Server Reference*. These are described below.

Table 23-4 SMTP LOG_CONNECTION Action Codes + or - Entries

Entry	Description
C	Connection closed. A diagnostic field will follow. Written to connection.log_current (or {mail.log_current if a single log file is being used). Used to record the reason why the connection was closed. In particular, if the connection was closed due to some session disconnect limit being reached, that fact will show up in the diagnostics field.
O	Connection opened.
T	PORT_ACCESS log entry. Further details available in the "PORT_ACCESS Mapping Table."
U	Logs SMTP authentication successes and failures. Format is the same as other O and C entries. In particular, the same application and transport information fields appear in same order. The username will be logged in the username field if it is known. Bit 7 (value 128) of the LOG_CONNECTION MTA option controls this.
X	Connection rejected.
Y	Connection attempt failed before being established.
I	ETRN command received.

With **LOG_CONNECTION**, **LOG_FILENAME**, **LOG_MESSAGE_ID**, **LOG_NOTARY**, **LOG_PROCESS**, and **LOG_USERNAME** MTA options all enabled, the format becomes as shown in the example below. (The sample log entry line has been wrapped for typographic reasons; the actual log entry would appear on one physical line.)

```
16-Feb-2007 15:04:01.14 2bbe.5.3 tcp_local ims-ms
EE 1 service@example.org rfc822;adam@example.com
adam@ims-ms-daemon 20 /opt/SUNWmsgsr/data/queue/ims-ms/000/ZZf0r2i0HIaY1.01
&lt;0JDJ00803FAON200@mailstore.example.org> mailsrv
example.org (example.org [192.160.253.66])
```

Where the additional fields, beyond those already discussed above, are:

1. The process ID (expressed in hexadecimal), followed by a period (dot) character and a count. If this had been a multithreaded channel entry (that is, a **tcp_*** channel entry), there would also be a thread ID present between the process ID and the count. In the example, the process ID is 2bbe.5.3.
2. The **NOTARY** (delivery receipt request) flags for the message, expressed as an integer (in the example, 20).
3. The file name in the MTA queue area (in the example, **/opt/SUNWmsgsr/data/queue/ims-ms/000/ZZf0r2i0HIaY1.01**).

4. The message ID (in the example, `0JDJ00803FAON200@mailstore.example.org`).
5. The name of the executing process (in the example, `mailsrv`). On UNIX, for dispatcher processes such as the SMTP server, this will usually be `mailsrv` (unless SASL was used, in which case it will be the authenticated user name, for example, `*service@example.org`).
6. The connection information (in the example, `example.org (example.org [192.160.253.66])`). The connection information consists of the sending system or channel name, such as the name presented by the sending system on the HELO/EHLO line (for incoming SMTP messages), or the enqueueing channel's official host name (for other sorts of channels). In the case of TCP/IP channels, the sending system's real name, that is, the symbolic name as reported by a DNS reverse lookup and/or the IP address, can also be reported within parentheses as controlled by the `ident*` channel options. See the discussion on the IDENT Lookups in the *Messaging Server Reference* for an instance of the default `identnone` option, that selects display of both the name found from the DNS and IP address.

Enabling MTA Logging

To gather statistics for just a few particular MTA channels, enable the logging channel option on just those MTA channels of interest. Many sites prefer to enable logging on all MTA channels. In particular, if you are trying to track down problems, the first step in diagnosing some problems is to notice that messages are not going to the channel you expected or intended, and having logging enabled for all channels can help you investigate such problems.

To Enable MTA Logging on a Specific Channel

1. Run the `msconfig edit channels` command.
2. To enable logging for a particular channel, add the `logging` option to the channel definition. For example:

```
channel-name option1 option2 logging
```

In addition, you can also set a number of configuration options such as directory path for log files, log levels, and so on. See "Managing Message Store, Admin, and Default Service Logs."

Note: The message return job, which runs every night around midnight, appends any existing `mail.log_yesterday` to the cumulative log file, `mail.log`, renames the current `mail.log_current` file to `mail.log_yesterday`, and then begins a new `mail.log_current` file. It also performs the analogous operations for any `connection.log*` files. It is possible that `mail.log_yesterday` contains time stamps which have already passed over rotation time.

To Enable MTA Logging on All Channels

1. Run the `msconfig edit channels` command.
2. Add the `logging` option to your `defaults` channel configuration file. For example:

```
defaults notices 1 2 4 7 copywarnpost copysendpost postheadonly
noswitchchannel \
immonurgent maxjobs 7 defaulthost example.org example.org logging
```

!

```
! delivery channel to local /var/mail store
l subdirs 20 viaaliasrequired maxjobs 7
mailhost.example.org
```

Specifying Additional MTA Logging Options

In addition to the basic information always provided when logging is enabled, you can specify that additional, optional information fields be included by setting various **LOG_*** MTA options.

To Send MTA Logs to syslog

1. Enter the following command to set **log_messages_syslog** to 1:

```
msconfig set log_messages_syslog 1
```

To write MTA message log file entries to syslog, you need to set the **log_messages_syslog** option to a non-zero value. The absolute value of the non-zero value sets the syslog priority and facility mask. Negative values disable the generation of the regular **mail.log*** entries. Positive values mean that the syslog entries are generated in addition to the regular **mail.log*** entries. 0 is the default and means no syslog or event logging is performed.

Facility and priority numbers are located in the **/usr/include/sys/syslog.h** file.

To Control Formatting of Log Entries

1. Set the **LOG_FORMAT** option by running **msconfig set log_format *n*** where *n* corresponds to one of the settings below
 - 1 (default) the standard format.
 - 2 requests non-null formatting: empty address fields are converted to the string "<>".
 - 3 requests counted formatting: all variable length fields are preceded by **N**, where **N** is a count of the number of characters in the field.
 - 4 causes log entries to be written in an XML-compatible format. Entry log entry appears as a single XML element containing multiple attributes and no sub-elements. Three elements are currently defined, **en** for enqueue/dequeue entries, **co** for connection entries, and **he** for header entries.

Enqueue/dequeue (**en**) elements can have the following attributes:

```
ts - time stamp (always present)
no - node name (present if LOG_NODE=1)
pi - process id (present if LOG_PROCESS=1)
sc - source channel (always present)
dc - destination channel (always present)
ac - action (always present)
sz - size (always present)
so - source address (always present)
od - original destination address (always present)
de - destination address (always present)
rf - recipient flags (present if LOG_NOTARY=1)
fi - filename (present if LOG_FILENAME=1)
ei - envelope id (present if LOG_ENVELOPE_ID=1)
mi - message id (present if LOG_MESSAGE_ID=1)
us - username (present if LOG_USERNAME=1)
ss - source system (present if bit 0 of LOG_CONNECTION
is set and source system information is available)
```

se - sensitivity (present if LOG_SENSITIVITY=1)
 pr - priority (present if LOG_PRIORITY=1)
 in - intermediate address (present if LOG_INTERMEDIATE=1)
 ia - initial address (present if bit 0 of LOG_INTERMEDIATE is set and intermediate address information is available)
 fl - filter (present if LOG_FILTER=1 and filter information is available)
 re - reason (present if LOG_REASON=1 and reason string is set)
 di - diagnostic (present if diagnostic info available)
 tr - transport information (present if bit 5 of LOG_CONNECTION is set and transport information is available)
 ap - application information (present if bit 6 of LOG_CONNECTION is set and application information is available)
 qt - the number of seconds the message has spent in the queue (LOG_QUEUE_TIME=1)

Here is a sample **en** entry:

```

<en ts="2004-12-08T00:40:26.70" pi="0d3730.10.43" sc="tcp_local"
dc="l" ac="E" sz="12" so="info-E8944AE8D033CB92C2241E@whittlesong.com"
od="rfc822;ned+2Bcharsets@mauve.sun.com"
de="ned+charsets@mauve.sun.com" rf="22"
fi="/path/ZZ01LI4XPX0DTM00IKA8.00" ei="01LI4XPQR2EU00IKA8@mauve.sun.com"
mi="&lt;11a3b401c4dd01$7c1c1ee0$1906fad0@elara>" us=""
ss="elara.whittlesong.com ([208.250.6.25])"
in="ned+charsets@mauve.sun.com" ia="ietf-charsets@innosoft.com"
fl="spamfilter1:rvLiXh158xWdQKa9iJ0d7Q==, addheader, keep"/>
  
```

Note that this entry has been wrapped for clarity; actual log file entries always appear on a single line. Connection (**co**) entries can have the following attributes:

ts - time stamp (always present, also used in en entries)
 no - node name (present if LOG_NODE=1, also used in en entries)
 pi - process id (present if LOG_PROCESS=1, also used in en entries)
 sc - source channel (always present, also used in en entries)
 dr - direction (always present)
 ac - action (always present, also used in en entries)
 tr - transport information (always present, also used in en entries)
 ap - application information (always present, also used in en entries)
 mi - message id (present only if message id info available, also used in en entries)
 us - username (present only if username information available, also used in en entries)
 di - diagnostic (present only if diagnostic information available, also used in en entries)
 ct - the length of the connection, in seconds. (LOG_QUEUE_TIME=1, also used in en entries)

Here is a sample **co** entry:

```

<co ts="2004-12-08T00:38:28.41" pi="1074b3.61.281" sc="tcp_local" dr="+
ac="0" tr="TCP|209.55.107.55|25|209.55.107.104|33469" ap="SMTP"/>
  
```

Header (**he**) entries have the following attributes:

ts - time stamp (always present, also used in en entries)
 no - node name (present if LOG_NODE=1, also used in en entries)
 pi - process id (present if LOG_PROCESS=1, also used in en entries)
 va - header line value (always present)

Here is a sample **he** entry:

```
<he ts="2004-12-08T00:38:31.41" pi="1074b3.61.281" va="Subject: foo"/>
```

To Correlate Log Message Entries

- Set the **LOG_MESSAGE_ID** option to 1 by running **msconfig set log_message_id 1**.

A value of 0 is the default and indicates that message IDs are not saved in the **mail.log*** files.

To Log Amount of Time Messages Have Spent in the Queue

- Set the **LOG_QUEUE_TIME** option to 1 by running **msconfig set log_queue_time 1**. This option logs the amount of time messages spent in the queue. The queue time is logged as an integer value in seconds. It appears immediately after the application information string in non-XML format logs. The attribute name in XML formatted logs for this value is **qt**.

To Identify Message Delivery Retries

- Set the **LOG_FILENAME** option to 1 by running **msconfig set log_filename 1**. This option makes it easier to immediately spot how many times the delivery of a particular message file has been retried. This option can also be useful in understanding when the MTA does or does not split a message to multiple recipients into separate message file copies on disk.

To Log TCP/IP Connections

- Set the **LOG_CONNECTION** option by running **msconfig set log_connection 1**. This option causes the MTA to log TCP/IP connections, as well as message traffic. The connection log entries are written to the **mail.log*** files by default. Optionally, the connection log entries can be written to **connection.log*** files. See the **SEPARATE_CONNECTION_LOG** option for more information.

To Write Entries to the connection.log File

- Set the **SEPARATE_CONNECTION_LOG** option to 1 by running **msconfig set separate_connection_log 1**. Use this option to specify that connection log entries instead be written to **connection.log** files. The default value of 0 causes the connection logging to be stored in the MTA log files.

To Correlate Log Messages by Process ID

- Set the **LOG_PROCESS** option by running **msconfig set log_process 1**. When used in conjunction with **LOG_CONNECTION**, this option enables correlation by process ID of which connection entries correspond to which message entries.

To Save User Names Associated with a Process That Enqueues Mail to the mail.log File

- Set the **LOG_USERNAME** option by running **msconfig set log_username 1**. This option controls whether or not the user name associated with a process that enqueues mail is saved in the **mail.log** file. For SMTP submissions where SASL (SMTP AUTH) is used, the user name field will be the authenticated user name (prefixed with an asterisk character).

MTA Message Logging Examples

The exact field format and list of fields logged in the MTA message files vary according to the logging options set. This section shows a few examples of interpreting typical sorts of log entries.

For a description of additional, optional fields, see "[Specifying Additional MTA Logging Options.](#)"

Note: For typographic reasons, log file entries will be shown folded onto multiple lines. Actual log file entries are one line per entry.

When reviewing a log file, keep in mind that on a typical system many messages are being handled at once. Typically, the entries relating to a particular message will be interspersed among entries relating to other messages being processed during that same time. The basic logging information is suitable for gathering a sense of the overall numbers of messages moving through the MTA.

If you wish to correlate particular entries relating to the same message to the same recipient(s), enable **LOG_MESSAGE_ID**. To correlate particular messages with particular files in the MTA queue area, or to see from the entries how many times a particular not-yet-successfully-dequeued message has had delivery attempted, enable **LOG_FILENAME**. For SMTP messages (handled via a TCP/IP channel), if you want to correlate TCP connections to and from remote systems with the messages sent, enable **LOG_PROCESS** and some level of **LOG_CONNECTION**.

MTA Logging Example: User Sends an Outgoing Message

The following example shows a basic example of the sort of log entries one might see if a local user sends a message out an outgoing TCP/IP channel, for example, to the Internet. In this example, **LOG_CONNECTION** is enabled. The lines marked with (1) and (2) are one entry---they would appear on one physical line in an actual log file. Similarly, the lines marked with (3) - (7) are one entry and would appear on one physical line.

Example MTA Logging: A Local User Sends An Outgoing Message

```
16-Feb-2007 15:41:32.36 tcp_intranet tcp_local EE 1 (1)
adam@example.com rfc822;marlowe@example.org marlowe@example.org (2)
example.org (example.org [192.160.253.66])

16-Feb-2007 15:41:34.73 tcp_local DE 1 (3)
adam@example.com rfc822;marlowe@example.org marlowe@example.org (4)
thor.example.org dns;thor.example.org

(TCP|206.184.139.12|2788|192.160.253.66|25) (5)

(thor.example.org ESMTP Sendmail ready Thu 15 Feb 2007 21:37:29 -0700 [MST]) (6)

smtp;250 2.1.5 &lt;marlowe@example.org>... Receipt ok (7)
```

1. This line shows the date and time of an enqueue with ESMTP (EE) from the **tcp_intranet** channel to the **tcp_local** channel of a one (1) block message.
2. This is part of the same physical line of the log file as (1), presented here as a separate line for typographical convenience. It shows the envelope **From:** address, in this case **adam@example.com**, and the original version and current version of the envelope **To:** address, in this case **marlowe@example.org**.
3. This shows the date and time of a dequeue with ESMTP (DE) from the **tcp_local** channel of a one (1) block message that is, a successful send by the **tcp_local** channel to some remote SMTP server.
4. This shows the envelope **From:** address, the original envelope **To:** address, and the current form of the envelope **To:** address.

5. This shows that the actual system to which the connection was made is named **thor.example.org** in the DNS, that the local sending system has IP address 206.184.139.12 and is sending from port 2788, that the remote destination system has IP address 192.160.253.66 and the connection port on the remote destination system is port 25.
6. This shows the SMTP banner line of the remote SMTP server.
7. This shows the SMTP status code returned for this address; 250 is the basic SMTP success code and in addition, this remote SMTP server responds with extended SMTP status codes and some additional text.

MTA Logging Example: Including Optional Logging Fields

This example shows a logging entry similar to that shown in "Example MTA Logging: Sending to a List" with **LOG_FILENAME=1** and **LOG_MESSAGE_ID=1** showing the file name (1 and 3 below) and message ID (2 and 4 below). The message ID in particular can be used to correlate which entries relate to which message.

Example MTA Logging: Including Optional Logging Fields

```
16-Feb-2007 15:41:32.36 tcp_intranet tcp_local EE 1
adam@example.com rfc822;marlowe@example.org marlowe@example.org
/opt/SUNWmsgsr/data/queue/tcp_local/002/Zzf0r4i0Wdy51.01 (1)
<0JDJ00D02IBWDX00@example.com> (2)
example.org (example.org [192.160.253.66])

16-Feb-2007 15:41:34.73 tcp_local DE 1
adam@example.com rfc822;marlowe@example.org marlowe@example.org
/opt/SUNWmsgsr/data/queue/tcp_local/002/Zzf0r4i0Wdy51.01 (3)
<0JDJ00D02IBWDX00@example.com> (4)
thor.example.org dns;thor.example.org
(TCP|206.184.139.12|2788|192.160.253.66|25)
(thor.example.org ESMTTP Sendmail ready at Thu, 15 Feb 2007 21:37:29 -0700 [MST])
smtp;250 2.1.5 <marlowe@example.org>... Recipient ok
```

MTA Logging Example: Sending to a List

This example illustrates sending to multiple recipients with **LOG_FILENAME=1**, **LOG_MESSAGE_ID=1**, and **LOG_CONNECTION=1** enabled. Here user *adam@example.com* has sent to the MTA mailing list *test-list@example.com*, which expanded to *bob@example.com*, *carol@example.org*, and *david@example.org*. Note that the original envelope **To:** address is **test-list@example.com** for each recipient, though the current envelope **To:** address is each respective address. Note how the message ID is the same throughout, though two separate files (one for the **I** channel and one going out the **tcp_local** channel) are involved.

Example MTA Logging: Sending to a List

```
20-Feb-2007 14:00:16.46 tcp_local tcp_local EE 1
adam@example.com rfc822;test-list@example.com carol@example.org
/opt/SUNWmsgsr/data/queue/tcp_local/004/Zzf0r2D0yuej4.01
<0JDQ00706R0FX100@example.com>
example.org (example.org [192.160.253.66])

20-Feb-2007 14:00:16.47 tcp_local tcp_local EE 1
adam@example.com rfc822;test-list@example.com david@example.org
/opt/SUNWmsgsr/data/queue/tcp_local/004/Zzf0r2D0yuej4.01
<0JDQ00706R0FX100@example.com>
example.org (example.org [192.160.253.66])

20-Feb-2007 14:00:16.48 tcp_local ims-ms EE 1
```

```

adam@example.com rfc822;test-list@example.com bob@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/008/Zzf0r2D0yuej6.01
&lt;0JDQ00706R0FX100@example.com>
example.org (example.org [192.160.253.66])

20-Feb-2007 14:00:16.68 ims-ms D 1
adam@example.com rfc822;test-list@example.com bob@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/008/Zzf0r2D0yuej6.01
&lt;0JDQ00706R0FX100@example.com>

20-Feb-2007 14:00:17.73 tcp_local DE 1
adam@example.com rfc822;test-list@example.com carol@example.org
/opt/SUNWmsgsr/data/queue/tcp_local/004/Zzf0r2D0yuej4.01
&lt;0JDQ00706R0FX100@example.com>
gw.example.org dns;gw.example.org (TCP|206.184.139.12|2788|192.160.253.66|25)
(gw.example.org -- SMTP Sendmail)
smtp;250 2.1.5 &lt;carol@example.org >... Recipient ok

20-Feb-2007 14:00:17.75 tcp_local DE 1
adam@example.com rfc822;test-list@example.com david@example.org
/opt/SUNWmsgsr/data/queue/tcp_local/004/Zzf0r2D0yuej4.01
&lt;0JDQ00706R0FX100@example.com>
gw.example.org dns;gw.example.org (TCP|206.184.139.12|2788|192.160.253.66|25)
(gw.example.org -- SMTP Sendmail)
smtp;250 2.1.5 &lt;david@example.org>... Recipient ok

```

MTA Logging Example: Sending to a Nonexistent Domain

This example illustrates an attempt to send to a nonexistent domain (here **very.bogus.com**); that is, sending to a domain name that is not noticed as nonexistent by the MTA's rewrite rules and that the MTA matches to an outgoing TCP/IP channel. This example assumes the MTA option settings of **LOG_FILENAME=1** and **LOG_MESSAGE_ID=1**.

When the TCP/IP channel runs and checks for the domain name in the DNS, the DNS returns an error that no such name exists. Note the "rejection" entry (R), as seen in (5), with the DNS returning an error that this is not a legal domain name, as seen in (6).

Because the address is rejected after the message has been submitted, the MTA generates a bounce message to the original sender. The MTA enqueues the new rejection message to the original sender (1), and sends a copy to the postmaster (4) before deleting the original outbound message (the R entry shown in (5)).

Notification messages, such as bounce messages, have an empty envelope From: address--as seen, for instance, in (2) and (8)--in which the envelope From: field is shown as an empty space. The initial enqueue of a bounce message generated by the MTA shows the message ID for the new notification message followed by the message ID for the original message (3). (Such information is not always available to the MTA, but when it is available to be logged, it allows correlation of the log entries corresponding to the outbound failed message with the log entries corresponding to the resulting notification message.) Such notification messages are enqueued to the process channel, which in turn enqueues them to an appropriate destination channel (7).

Example MTA Logging: Sending to a Nonexistent Domain

```

20-Feb-2007 14:17:07.77 tcp_intranet tcp_local E 1
adam@example.com rfc822;user@very.bogus.com user@very.bogus.com
/opt/SUNWmsgsr/data/queue/tcp_local/008/Zzf0r2D0CVaL0.00
&lt;0JDQ00903RS89T00@example.com>
example.org (example.org [192.160.253.66])

```

```

20-Feb-2007 14:17:08.24 tcp_local process E 1 (1)
rfc822;adam@example.com adam@example.com (2)
/opt/SUNWmsgsr/data/queue/process/ZZf0r2D0CVbR0.00
&lt;0JDQ00904RSK9Z00@example.com>,&lt;0JDQ00903RS89T00@example.com> (3)
tcp-daemon.mailhost.example.com

20-Feb-2007 14:17:08.46 tcp_local process E 1 (4)
rfc822;postmaster@example.com postmaster@example.com
/opt/SUNWmsgsr/data/queue/process/ZZf0r2D0CVbR1.00
&lt;0JDQ00906RSK9Z00@example.com>,&lt;0JDQ00903RS89T00@example.com>
tcp-daemon.mailhost.example.com

20-Feb-2007 14:17:08.46 tcp_local R 1 (5)
adam@example.com rfc822;user@very.bogus.com user@very.bogus.com
/opt/SUNWmsgsr/data/queue/tcp_local/008/ZZf0r2D0CVaL0.00
&lt;0JDQ00903RS89T00@example.com>
Illegal host/domain name found (6)
(TCP active open: Failed gethostbyname() on very.bogus.com, resolver errno = 1)

20-Feb-2007 14:17:09.21 process ims-ms E 3 (7)
rfc822;adam@example.com adam@ims-ms-daemon (8)
/opt/SUNWmsgsr/data/queue/ims-ms/018/ZZf0r2D0CVbS1.00
&lt;0JDQ00904RSK9Z00@example.com>
process-daemon.mailhost.example.com

20-Feb-2007 14:17:09.72 process ims-ms E 3
rfc822;postmaster@example.com postmaster@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/014/ZZf0r2D0CVbS2.00
&lt;0JDQ00906RSK9Z00@example.com>
process-daemon.mailhost.example.com

20-Feb-2007 14:17:09.73 ims-ms D 3
rfc822;adam@example.com adam@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/018/ZZf0r2D0CVbS1.00
&lt;0JDQ00904RSK9Z00@example.com>

20-Feb-2007 14:17:09.84 ims-ms D 3
rfc822;postmaster@example.com postmaster@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/014/ZZf0r2D0CVbS2.00
&lt;0JDQ00906RSK9Z00@example.com>

```

MTA Logging Example: Sending to a Nonexistent Remote User

This example illustrates an attempt to send to a bad address on a remote system. This example assumes MTA option settings of **LOG_FILENAME=1** and **LOG_MESSAGE_ID=1**, and channel option settings of **LOG_BANNER=1** and **LOG_TRANSPORTINFO=1**. Note the rejection entry (R), seen in (1). But in contrast to the rejection entry in "Example MTA Logging: Sending to a Nonexistent Domain," note that the rejection entry here shows that a connection to a remote system was made, and shows the SMTP error code issued by the remote SMTP server, (2) and (3). The inclusion of the information shown in (2) is due to setting the channel options **LOG_BANNER=1** and **LOG_TRANSPORTINFO=1**.

Example MTA Logging: Sending to a Nonexistent Remote User

```

26-Feb-2007 13:56:35.16 tcp_intranet tcp_local EE 1
adam@example.com rfc822;nonesuch@example.org nonesuch@example.org
/opt/SUNWmsgsr/data/queue/tcp_local/000/ZZf0s690a3mf2.01
&lt;0JE100J08UU24H00@example.com>
example.org (example.org [192.160.253.66])

```

```
26-Feb-2007 13:56:35.19 tcp_local process E 1
rfc822;adam@example.com adam@example.com
/opt/SUNWmsgsr/data/queue/process/ZZf0s690a3m12.00
&lt;0JE100J09UUB4N00@example.com>,&lt;0JE100J08UU24H00@example.com>
tcp-daemon.mailhost.example.com
```

```
26-Feb-2007 13:56:35.20 tcp_local process E 1
rfc822;postmaster@example.com postmaster@example.com
/opt/SUNWmsgsr/data/queue/process/ZZf0s690a3m13.00
&lt;0JE100J0BUUB4N00@example.com>,&lt;0JE100J08UU24H00@example.com>
tcp-daemon.mailhost.example.com
```

```
26-Feb-2007 13:56:35.20 tcp_local RE 1 (1)
adam@example.com rfc822;nonesuch@example.org nonesuch@example.org

/opt/SUNWmsgsr/data/queue/tcp_local/000/ZZf0s690a3mf2.01
&lt;0JE100J08UU24H00@example.com>
thor.example.org dns;thor.example.org
(TCP|206.184.139.12|2788|192.160.253.66|25) (2)
(thor.example.org -- Server ESMTTP [Sun Java System Messaging
Server 6.2-8.01 [built Feb 16 2007]])
smtp;550 5.1.1 unknown or illegal alias: nonesuch@example.org (3)
```

```
26-Feb-2007 13:56:35.62 process ims-ms E 4
rfc822;adam@example.com adam@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/003/ZZf0s690a3mm5.00
&lt;0JE100J09UUB4N00@example.com>
process-daemon.mailhost.example.com
```

```
26-Feb-2007 13:56:36.07 process ims-ms E 4
rfc822;postmaster@example.com postmaster@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/016/ZZf0s690a3nm7.01
&lt;0JE100J0BUUB4N00@example.com>
process-daemon.mailhost.example.com
```

```
26-Feb-2007 13:56:35.83 ims-ms D 4
rfc822;adam@example.com adam@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/003/ZZf0s690a3mm5.00
&lt;0JE100J09UUB4N00@example.com>
```

```
26-Feb-2007 13:56:36.08 ims-ms D 4
rfc822;postmaster@example.com postmaster@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/016/ZZf0s690a3nm7.01
&lt;0JE100J0BUUB4N00@example.com>
```

MTA Logging Example: Rejecting a Remote Side's Attempt to Submit a Message

This example illustrates the sort of log file entry resulting when the MTA rejects a remote side's attempt to submit a message. (This example assumes that no optional **LOG_*** options are enabled, so only the basic fields are logged in the entry. Note that enabling the **LOG_CONNECTION** option, in particular, would result in additional informative fields in such J entries.) In this case, the example is for an MTA that has set up SMTP relay blocking (see "[Configuring SMTP Relay Blocking](#)") with an **ORIG_SEND_ACCESS** mapping, including:

```
ORIG_SEND_ACCESS
```

```
! ...numerous entries omitted...
!
```

```
tcp_local|*|tcp_local|* $NRelaying$ not$ permitted
```

and where **alan@very.bogus.com** is not an internal address. Hence the attempt of the remote user **harold@example.org** to relay through the MTA system to the remote user **alan@very.bogus.com** is rejected.

Example MTA Logging: Rejecting a Remote Side's Attempt to Submit a Message

```
26-Feb-2007 14:10:06.89 tcp_local JE 0 (1)
harold@example.org rfc822; alan@very.bogus.com (2)
530 5.7.1 Relaying not allowed: alan@very.bogus.com (3)
```

1. This log shows the date and time the MTA rejects a remote side's attempt to submit a message. The rejection is indicated by a J record. (Cases where an MTA channel is attempting to send a message which is rejected is indicated by R records, as shown in "Example MTA Logging: Sending to a Nonexistent Domain" and "MTA Logging Example: Sending to a Nonexistent Remote User").

Note: The last J record written to the log will have an indication stating that it is the last for a given session. Also, the current version of Messaging Server does not place a limit on the number of J records.

2. The attempted envelope **From:** and **To:** addresses are shown. In this case, no original envelope **To:** information was available so that field is empty.
3. The entry includes the SMTP error message the MTA issued to the remote (attempted sender) side.

MTA Logging Example: Multiple Delivery Attempts

This example illustrates the sort of log file entries resulting when a message cannot be delivered upon the first attempt, so the MTA attempts to send the message several times. This example assumes option settings of **LOG_FILENAME=1** and **LOG_MESSAGE_ID=1**.

Example MTA Logging: Multiple Delivery Attempts

```
26-Feb-2007 14:38:16.27 tcp_intranet tcp_local EE 1 (1)
adam@example.com rfc822;user@some.org user@some.org
/opt/SUNWmsgsr/data/queue/tcp_local/001/Zzf0s690kN_y0.00
&lt;0JE100L05WRJIC00@example.com>
```

```
26-Feb-2007 14:38:16.70 tcp_local Q 1 (2)
adam@example.com rfc822;user@some.org user@some.org
/opt/SUNWmsgsr/data/queue/tcp_local/001/Zzf0s690kN_y0.00 (3)
&lt;0JE100L05WRJIC00@example.com>
TCP active open: Failed connect() 192.1.1.1:25 Error: no route to host (4)
```

...several hours worth of entries...

```
26-Feb-2007 16:58:11.20 tcp_local Q 1 (5)
adam@example.com rfc822;user@some.org user@some.org
/opt/SUNWmsgsr/data/queue/tcp_local/001/ZYf0s690kN_y0.01 (6)
&lt;0JE100L05WRJIC00@example.com>
TCP active open: Failed connect() 192.1.1.1:25 Error: no route to host
```

...several hours worth of entries...

```
26-Feb-2007 19:15:12.11 tcp_local Q 1
adam@example.com rfc822;user@some.org user@some.org
```

```
/opt/SUNWmsgsr/data/queue/tcp_local/001/ZXf0s690kN_y0.00 (7)
<0JE100L05WRJIC00@example.com>
TCP active open: Failed connect() 192.1.1.1:25 Error: Connection refused (8)
```

...several hours worth of entries...

```
26-Feb-2007 22:41:12.63 tcp_local DE 1 (9)
adam@example.com rfc822;user@some.org user@some.org
/opt/SUNWmsgsr/data/queue/tcp_local/001/ZXf0s690kN_y0.00
<0JE100L05WRJIC00@example.com>
host.some.org dns;host.some.org (TCP|206.184.139.12|2788|192.1.1.1|25)
(All set, fire away)
smtp;250 2.1.5 <user@some.org >... Recipient ok
```

1. The message comes in the **tcp_internal** channel---perhaps from a POP or IMAP client, or perhaps from another host within the organization using the MTA as an SMTP relay; the MTA enqueues it to the outgoing **tcp_local** channel.
2. The first delivery attempt fails, as indicated by the Q entry.
3. That this is a first delivery attempt can be seen from the **ZZ*** filename.
4. This delivery attempt failed when the TCP/IP package could not find a route to the remote side. As opposed to "[MTA Logging Example: Sending to a Nonexistent Domain](#)," the DNS did not object to the destination domain name, **some.org**; rather, the "no route to host" error indicates that there is some network problem between the sending and receiving side.
5. The next time the MTA periodic job runs it reattempts delivery, again unsuccessfully.
6. The file name is now **ZY***, indicating that this is a second attempt.
7. The file name is **ZX*** for this third unsuccessful attempt.
8. The next time the periodic job reattempts delivery the delivery fails, though this time the TCP/IP package is not complaining that it cannot get through to the remote SMTP server, but rather the remote SMTP server is not accepting connections. (Perhaps the remote side fixed their network problem, but has not yet brought their SMTP server back up---or their SMTP server is swamped handling other messages and hence was not accepting connections at the moment the MTA tried to connect.)
9. Finally the message is dequeued.

MTA Logging Example: Incoming SMTP Message Routed Through the Conversion Channel

This example illustrates the case of a message routed through the conversion channel. The site is assumed to have a CONVERSIONS mapping table such as:

```
CONVERSIONS
```

```
IN-CHAN=tcp_local;OUT-CHAN=ims-ms;CONVERT Yes
```

This example assumes option settings of **LOG_FILENAME=1** and **LOG_MESSAGE_ID=1**.

Example MTA Logging: Incoming SMTP Message Routed Through the Conversion Channel

```
26-Feb-2007 15:31:04.17 tcp_local conversion EE 1 (1)
amy@example.edu rfc822;bert@example.com bert@ims-ms-daemon
```

```
/opt/SUNWmsgsr/data/queue/conversion/ZZf0s090wFwx2.01
&lt;0JE100206Z7J5F00@example.edu>
```

```
26-Feb-2007 15:31:04.73 conversion ims-ms E 1 (2)
amy@example.edu rfc822;bert@example.com bert@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/007/ZZf0s090wMwq1.00
&lt;0JE100206Z7J5F00@example.edu>
```

```
26-Feb-2007 15:31:04.73 conversion D 1 (3)
amy@example.edu rfc822;bert@example.com bert@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/conversion/ZZf0s090wFwx2.01
&lt;0JE100206Z7J5F00@example.edu>
```

```
26-Feb-2007 15:31:04.73 ims-ms D 1 (4)
amy@example.edu rfc822;bert@example.com bert@ims-ms-daemon
/opt/SUNWmsgsr/data/queue/ims-ms/007/ZZf0s090wMwq1.00
&lt;0JE100206Z7J5F00@example.edu>
```

1. The message from external user amy@example.edu comes in addressed to the **ims-ms** channel recipient bert@example.com. The **CONVERSIONS** mapping entry, however, causes the message to be initially enqueued to the conversion channel (rather than directly to the **ims-ms** channel).
2. The conversion channel runs and enqueues the message to the **ims-ms** channel.
3. Then the conversion channel can dequeue the message (delete the old message file).
4. And finally the **ims-ms** channel dequeues (delivers) the message.

MTA Logging Example: Outbound Connection Logging

This example illustrates log output for an outgoing message when connection logging is enabled, via **LOG_CONNECTION=3**. **LOG_PROCESS=1**, **LOG_MESSAGE_ID=1** and **LOG_FILENAME=1** are also assumed in this example. The example shows the case of user adam@example.com sending the same message (note that the message ID is the same for each message copy) to three recipients, bobby@hosta.example.com, carl@hosta.example.com, and dave@hostb.example.com. This example assumes that the message is going out a **tcp_local** channel marked (as such channels usually are) with the **single_sys** channel option. Therefore, a separate message file on disk will be created for each set of recipients to a separate host name, as seen in (1), (2), and (3), where the bobby@hosta.example.com and carl@hosta.example.com recipients are stored in the same message file, but the dave@hostb.example.com recipient is stored in a different message file.

Example MTA Logging: Outbound Connection Logging

```
28-Feb-2007 09:13:19.18 409f.3.1 tcp_intranet tcp_local EE 1
adam@example.com rfc822;bobby@hosta.example.com bobby@hosta.example.com
/opt/SUNWmsgsr/data/queue/tcp_local/000/ZZf0s4g0G2Zt0.00 (1)
&lt;0JE500C0371HRJ00@example.com>
example.org (example.org [192.160.253.66])
```

```
28-Feb-2007 09:13:19.18 409f.3.1 tcp_intranet tcp_local EE 1
adam@example.com rfc822;carl@hosta.example.com carl@hosta.example.com
/opt/SUNWmsgsr/data/queue/tcp_local/000/ZZf0s4g0G2Zt0.00 (2)
&lt;0JE500C0371HRJ00@example.com>
example.org (example.org [192.160.253.66])
```

```
28-Feb-2007 09:13:19.19 409f.3.2 tcp_intranet tcp_local EE 1
adam@example.com rfc822;dave@hostb.example.com dave@hostb.example.com
```

```

/opt/SUNWmsgsr/data/queue/tcp_local/004/Zzf0s4g0G2Zt1.00 (3)
&lt;0JE500C0371HRJ00@example.com>
example.org (example.org [192.160.253.66])

28-Feb-2007 09:13:19.87 40a5.2.0 tcp_local - O (4)
TCP|206.184.139.12|5900|206.184.139.66|25
SMTP/hostb.example.com/mailhub.example.com (5)

28-Feb-2007 09:13:20.23 40a5.3.4 tcp_local - O (6)
TCP|206.184.139.12|5901|206.184.139.70|25
SMTP/hosta.example.com/hosta.example.com (7)

28-Feb-2007 09:13:20.50 40a5.2.5 tcp_local DE 1
adam@example.com rfc822;bobby@hosta.example.com bobby@hosta.example.com
/opt/SUNWmsgsr/data/queue/tcp_local/000/Zzf0s4g0G2Zt0.00
&lt;0JE500C0371HRJ00@example.com>
hosta.example.com dns;hosta.example.com (8)
(TCP|206.184.139.12|5901|206.184.139.70|25)
(hosta.example.com -- Server ESMTP [Sun Java System Messaging Server
6.2-8.01 [built Feb 16 2007]])
smtp;250 2.1.5 bobby@hosta.example.com and options OK.

28-Feb-2007 09:13:20.50 40a5.2.5 tcp_local DE 1
adam@example.com rfc822;carl@hosta.example.com carl@hosta.example.com
/opt/SUNWmsgsr/data/queue/tcp_local/000/Zzf0s4g0G2Zt0.00
&lt;0JE500C0371HRJ00@example.com>
hosta.example.com dns;hosta.example.com
(TCP|206.184.139.12|5901|206.184.139.70|25)
(hosta.example.com -- Server ESMTP [Sun Java System Messaging Server
6.2-8.01 [built Feb 16 2007]])
smtp;250 2.1.5 carl@hosta.example.com and options OK.

28-Feb-2007 09:13:20.50 40a5.2.6 tcp_local - C (9)
TCP|206.184.139.12|5901|206.184.139.70|25
SMTP/hosta.example.com/hosta.example.com

28-Feb-2007 09:13:21.13 40a5.3.7 tcp_local DE 1
adam@example.com rfc822;dave@hostb.example.com dave@hostb.example.com
/opt/SUNWmsgsr/data/queue/tcp_local/004/Zzf0s4g0G2Zt1.00
&lt;0JE500C0371HRJ00@example.com>
mailhub.example.com dns;mailhub.example.com
(TCP|206.184.139.12|5900|206.184.139.66|25)
(mailhub.example.com ESMTP Sendmail ready at Tue, 27 Feb 2007 22:19:40 GMT)
smtp;250 2.1.5 &lt;dave@hostb.example.com>... Recipient ok

28-Feb-2007 09:13:21.33 40a5.3.8 tcp_local - C (10)
TCP|206.184.139.12|5900|206.184.139.66|25
SMTP/hostb.example.com/mailhub.example.com

```

1. The message is enqueued to the first recipient...
2.and to the second recipient...
3.and to the third recipient.
4. Having **LOG_CONNECTION=3** set causes the MTA to write this entry. The minus, -, indicates that this entry refers to an outgoing connection. The **O** means that this entry corresponds to the opening of the connection. Also note that the process ID here is the same, 40a5, since the same process is used for the multithreaded TCP/IP channel for these separate connection opens, though this open is being performed by thread 2 vs. thread 3.

5. As there are two separate remote systems to which to connect, the multithreaded SMTP client in separate threads opens up a connection to each---the first in this entry, and the second shown in 7. This part of the entry shows the sending and destination IP numbers and port numbers, and shows both the initial host name, and the host name found by doing a DNS lookup. In the **SMTP/initial-host/dns-host** clauses, note the display of both the initial host name, and that used after performing a DNS MX record lookup on the initial host name: **mailhub.example.com** is apparently an MX server for **hostb.example.com**.
6. The multithreaded SMTP client opens up a connection to the second system in a separate thread (though the same process).
7. As there are two separate remote systems to which to connect, the multithreaded SMTP client in separate threads opens up a connection to each---the second in this entry, and the first shown above in 5. This part of the entry shows the sending and destination IP numbers and port numbers, and shows both the initial host name, and the host name found by doing a DNS lookup. In this example, the system **hosta.example.com** apparently receives mail directly itself.
8. Besides resulting in specific connection entries, **LOG_CONNECTION=3** also causes inclusion of connection related information in the regular message entries, as seen here for instance.
9. Having **LOG_CONNECTION=3** causes the MTA to write this entry. After any messages are dequeued, (the bobby and carl messages in this example), the connection is closed, as indicated by the **C** in this entry.
10. The connection mailhub.example.com is closed now that the delivery of the message (dave in this example) is complete.

MTA Logging Example: Inbound Connection Logging

This example illustrates log output for an incoming SMTP message when connection logging is enabled, via **LOG_CONNECTION=3**.

Example MTA Logging: Inbound Connection Logging

```
28-Feb-2007 11:50:59.10 tcp_local + O (1)
TCP|206.184.139.12|25|192.160.253.66|1244 SMTP (2)

28-Feb-2007 11:51:15.12 tcp_local ims-ms EE 1
service@example.org rfc822;adam@example.com adam@ims-ms-daemon
THOR.EXAMPLE.ORG (THOR.EXAMPLE.ORG [192.160.253.66]) (3)

28-Feb-2007 11:51:15.32 ims-ms D 1
service@example.org rfc822;adam@example.com adam@ims-ms-daemon

28-Feb-2007 11:51:15.66 tcp_local + C (4)
TCP|206.184.139.12|25|192.160.253.66|1244 SMTP
```

1. The remote system opens a connection. The **O** character indicates that this entry regards the opening of a connection; the **+** character indicates that this entry regards an incoming connection.
2. The IP numbers and ports for the connection are shown. In this entry, the receiving system (the system making the log file entry) has IP address 206.184.139.12 and the connection is being made to port 25; the sending system has IP address 192.160.253.66 and is sending from port 1244.
3. In the entry for the enqueue of the message from the incoming TCP/IP channel (**tcp_local**) to the **ims-ms** channel recipient, note that information beyond the default is included since **LOG_CONNECTION=3** is enabled. Specifically, the

name that the sending system claimed on its HELO or EHLO line, the sending system's name as found by a DNS reverse lookup on the connection IP number, and the sending system's IP address are all logged.

4. The inbound connection is closed. The **C** character indicates that this entry regards the closing of a connection; the **+** character indicates that this entry regards an incoming connection.

Enabling Dispatcher Debugging

Dispatcher error and debugging output (if enabled) are written to the file **dispatcher.log** in the MTA log directory. A default dispatcher configuration is created at installation time and can be used without any changes made. However, if you want to modify the default configuration for security or performance reasons, you can do so by running the **msconfig** command.

Table 23–5 Dispatcher Debugging Bits

Bit	Hexadecimal value	Decimal value	Usage
0	x 00001	1	Basic Service Dispatcher main module debugging.
1	x 00002	2	Extra Service Dispatcher main module debugging.
2	x 00004	4	Service Dispatcher configuration file logging.
3	x 00008	8	Basic Service Dispatcher miscellaneous debugging.
4	x 00010	16	Basic service debugging.
5	x 00020	32	Extra service debugging.
6	x 00040	64	Process related service debugging.
7	x 00080	128	Not used.
8	x 00100	256	Basic Service Dispatcher and process communication debugging.
9	x 00200	512	Extra Service Dispatcher and process communication debugging.
10	x 00400	1024	Packet level communication debugging.
11	x 00800	2048	Not used.
12	x 01000	4096	Basic Worker Process debugging.
13	x 02000	8192	Extra Worker Process debugging.
14	x 04000	16384	Additional Worker Process debugging, particularly connection hand-offs.
15	x 08000	32768	Not used.
16	x 10000	65536	Basic Worker Process to Service Dispatcher I/O debugging.
17	x 20000	131072	Extra Worker Process to Service Dispatcher I/O debugging.
20	x 100000	1048576	Basic statistics debugging.
21	x 200000	2097152	Extra statistics debugging.
24	x 1000000	16777216	Log PORT_ACCESS denials to the dispatcher.log file.

To Enable Dispatcher Error Debugging Output

1. Run the **msconfig set dispatcher.debug -1**. (See "Restricted Options.") You can also set the logical or environmental variable **IMTA_DISPATCHER_DEBUG** (UNIX),

which defines a 32-bit debug mask in hexadecimal, to the value **FFFFFFFF**. The table above describes the meaning of each bit.

To Set Dispatcher options (Oracle Solaris)

The dispatcher services offered in the dispatcher configuration affect requirements for various system options. The system's heap size (**datasize**) must be enough to accommodate the dispatcher's thread stack usage.

1. To display the heap size (that is, default **datasize**), use one of the following: The **cs** command:

```
limit
```

The **ksh** command:

```
ulimit -a
```

The Solaris utility:

```
sysdef
```

2. For each dispatcher service compute **STACKSIZE*MAX_CONNS**, and then add up the values computed for each service. The system's heap size needs to be at least twice this number.

Managing Message Store, Admin, and Default Service Logs

This section describes logging for the Message Store (POP, IMAP, and HTTP), Admin, and Default services. (See [Table 23–1, "Messaging Server Log Files"](#) for more information.)

For these services, you specify log settings and to view logs. The settings you specify affect which and how many events are logged. You can use those settings and other characteristics to refine searches for logged events when you are analyzing log files.

This section contains the following subsections:

- [msconfig Logging Options](#)
- [Understanding Service Log Characteristics](#)
 - [Logging Levels](#)
 - [Categories of Logged Events](#)
 - [Service Log File Directories](#)
- [Understanding Service Log File Format](#)
 - [Store and Administration Log File Components](#)
- [Defining and Setting Service Logging Options](#)
 - [Flexible Logging Architecture](#)
 - [Planning the Options You Want](#)
 - [Understanding Logging Options](#)
- [Searching and Viewing Service Logs](#)
 - [Search Options](#)
- [Working With Service Logs](#)
 - [To Send Service Logs to syslog](#)

- To Disable HTTP Logging
- To Set the Server Log Level
- To Specify a Directory Path for Server Log Files
- To Specify a Maximum File Size for Each Service Log
- To Specify a Service Log Rotation Schedule
- To Specify a Maximum Number of Service Log Files Per Directory
- To Specify a Storage Limit
- To Specify the Minimum Amount of Free Disk Space to Reserve
- To Specify an Age for Logs at Which They Expire
- Implementing and Configuring Message Store Transaction Logging
 - Overview of Message Store Transaction Logging
 - Message Store Transaction Logging Log Entries
 - Configuring Message Store Transaction Logging
 - Message Store Transaction Log Examples
- Other Message Store Logging Features
- Message Store Logging Examples
 - Message Store Logging Example: Bad Password
 - Message Store Logging Example: Account Disabled
 - Message Store Logging Example: Message Appended
 - Message Store Logging Example: Message Retrieved by a Client
 - Message Store Logging Example: Message Removed from a Folder
 - Message Store Logging Example: Duplicate Login Messages

msconfig Logging Options

To control the location of log files, use the following options for specifying directory paths:

Note: The location of MTA log files, which are in the **msg-install-path/data/log** directory, cannot be modified, but you can change the *log* subdirectory to symbolically link to another location. To separate the MTA logs from the rest of the log files, use **msconfig** options to specify non-default locations for non-MTA log files.

Table 23–6 msconfig Directory Paths for Log Files

Option	Description
*.logfile.logdir	Directory path for log files. If this is not specified, log files will be placed in the msg-install-path/data/log directory. For the MTA, this option is only used by Message Store insertion tasks Directory path to the imta log file used for Message Store insertion (ims_master, LMTP). It is not used by other parts of the MTA which always log to the default location. The default location is msg-install-path/data/log . Changing that path to a soft-link is supported. (Restart of all services required). Syntax: dirpath
base.logfile.logdir	
dispatcher.logfile.logdir	
ens.logfile.logdir	
http.logfile.logdir	
imap.logfile.logdir	
imaproxy.logfile.logdir	
mta.logfile.logdir	
job_controller.logfile.logdir	
metermaid.logfile.logdir	
mmp.logfile.logdir	
msadmin.logfile.logdir	
messageTrace.logfile.logdir	
pop.logfile.logdir	
popproxy.logfile.logdir	
snmp.logfile.logdir	
submitproxy.logfile.logdir	
tcp_lmtp_server.logfile.logdir	

Understanding Service Log Characteristics

This section describes the following log characteristics for the message store and administration services: logging levels, categories of logged events, filename conventions for logs, and log-file directories.

Logging Levels

The level, or priority, of logging defines how detailed, or verbose, the logging activity is to be. A higher priority level means less detail; it means that only events of high priority (high severity) are logged. A lower level means greater detail; it means that more events are recorded in the log file.

You can set the logging level separately for each service---POP, IMAP, HTTP, Admin, and Default by setting the *service.logfile.loglevel* configuration option (see "[Defining and Setting Service Logging Options](#)"). You can also use logging levels to filter searches for log events. [Table 23–7](#) describes the available levels. These logging levels are a subset of those defined by the UNIX **syslog** facility.

Table 23–7 Logging Levels for Store and Administration Services

Level	Description
Critical	The minimum logging detail. An event is written to the log whenever a severe problem or critical condition occurs---such as when the server cannot access a mailbox or a library needed for it to run.
Error	An event is written to the log whenever an error condition occurs---such as when a connection attempt to a client or another server fails.

Table 23–7 (Cont.) Logging Levels for Store and Administration Services

Level	Description
Warning	An event is written to the log whenever a warning condition occurs---such as when the server cannot understand a communication sent to it by a client.
Notice	An event is written to the log whenever a notice (a normal but significant condition) occurs---such as when a user login fails or when a session closes. This is the default log level.
Information	An event is written to the log with every significant action that takes place---such as when a user successfully logs on or off or creates or renames a mailbox.
Debug	The most verbose logging. Useful only for debugging purposes. Events are written to the log at individual steps within each process or task, to pinpoint problems.

When you select a particular logging level, events corresponding to that level and to all higher (less verbose) levels are logged. The default level of logging is **Notice**.

Note: The more verbose the logging you specify, the more disk space your log files will occupy; for guidelines, see "[Defining and Setting Service Logging Options](#)."

Categories of Logged Events

Within each supported service or protocol, Messaging Server further categorizes logged events by the facility, or functional area, in which they occur. Every logged event contains the name of the facility that generated it. These categories aid in filtering events during searches. [Table 23–8](#) lists the categories that Messaging Server recognizes for logging purposes.

Table 23–8 Categories in Which Log Events Occur

Facility	Description
General	Undifferentiated actions related to this protocol or service
LDAP	Actions related to Messaging Server accessing the LDAP directory database
Network	Actions related to network connections (socket errors fall into this category)
Account	Actions related to user accounts (user logins fall into this category)
Protocol	Protocol-level actions related to protocol-specific commands (errors returned by POP, IMAP, or HTTP functions fall into this category)
Stats	Actions related to the gathering of server statistics
Store	Low-level actions related to accessing the message store (read/write errors fall into this category)

For examples of using categories as filters in log searches, see "[Searching and Viewing Service Logs](#)."

Service Log File Directories

Every logged service is assigned a single directory, in which its log files are stored. All IMAP log files are stored together, as are all POP log files, and log files of any other service. You define the location of each directory, and you also define how many log files of what maximum size are permitted to exist in the directory.

Make sure that your storage capacity is sufficient for all your log files. Log data can be voluminous, especially at lower (more verbose) logging levels.

It is important also to define your logging level, log rotation, log expiration, and server-backup policies appropriately so that all of your log-file directories are backed up and none of them become overloaded; otherwise, you may lose information. See "[Defining and Setting Service Logging Options](#)."

Understanding Service Log File Format

All message store and administration service log files created by Messaging Server have identical content formats. Log files are multiline text files, in which each line describes one logged event. All event descriptions, for each of the supported services, have the general format:

```
<dateTime> <hostName> <processName>[<pid>]: <category> <logLevel>: <eventMessage>
```

Table 23–9 lists the log file components. Note that this format of event descriptions is identical to that defined by the UNIX **syslog** facility, except that the date/time format is different and the format includes two additional components (*category* and *logLevel*).

Store and Administration Log File Components

Table 23–9 lists the log files components for POP and IMAP.

Table 23–9 POP and IMAP Log File Formats

Component	Definition
<i>dateTime</i>	The date and time at which the event was logged, expressed in <i>dd/mm/yyyy hh:mm:ss</i> format, with a time-zone field expressed as <i>+/-hhmm</i> from GMT. For example: 02/Jan/1999:13:08:21 -0700
<i>hostName</i>	The name of the host machine on which the server is running: for example, showshoe . Note: If there is more than one instance of Messaging Server on the host, you can use the process ID (pid) to separate logged events of one instance from another.
<i>processName</i>	The name of the process that generated the event: for example, cgi_store .
<i>pid</i>	The process ID of the process that generated the event: for example, 18753 .
<i>category</i>	The category that the event belongs to: for example, General (see " Example MTA Logging: Sending to a Nonexistent Remote User ").
<i>logLevel</i>	The level of logging that the event represents: for example, Notice (see " Example MTA Logging: Sending to a Nonexistent Domain ").
<i>eventMessage</i>	An event-specific explanatory message that may be of any length: for example, Log created (894305624) .

Here are three examples of logged events:

```
02/May/1998:17:37:32 -0700 showshoe cgi_store[18753]: General Notice:
Log created (894155852)
```

```
04/May/1998:11:07:44 -0400 xyzmail cgi_service[343]: General Error:
function=getserverhello|port=2500|error=failed to connect

03/Dec/1998:06:54:32 +0200 ExamplePost imapd[232]: Account Notice: close
[127.0.0.1]
[unauthenticated] 1998/12/3 6:54:32 0:00:00 0 115 0
```

IMAP and POP event entries may end with three numbers. The example above has 0 115 0. The first number is bytes sent by client, the second number is the bytes sent by the server, and third number is mailboxes selected (always 1 for POP).

When viewing a log file in the Log Viewer window, you can limit the events displayed by searching for any specific component in an event, such as a specific logging level or category, or a specific process ID. For more information, see ["Searching and Viewing Service Logs."](#)

The event message of each log entry is in a format specific to the type of event being logged, that is, each service defines what content appears in any of its event messages. Many event messages are simple and self-evident; others are more complex.

Defining and Setting Service Logging Options

You can define the message store and administration service logging configurations that best serve your administration needs. This section discusses issues that may help you decide on the best configurations and policies, and it explains how to implement them.

Flexible Logging Architecture

The naming scheme for log files (*service.sequenceNum.timeStamp*) helps you to design a flexible log-rotation and backup policy. The fact that events for different services are written to different files makes it easier for you to isolate problems quickly. Also, because the sequence number in a filename is ever-increasing and the timestamp is always unique, later log files do not simply overwrite earlier ones after a limited set of sequence numbers is exhausted. Instead, older log files are overwritten or deleted only when the more flexible limits of age, number of files, or total storage are reached.

Messaging Server supports automatic rotation of log files, which simplifies administration and facilitates backups. You are not required to manually retire the current log file and create a new one to hold subsequent logged events. You can back up all but the current log file in a directory at any time, without stopping the server or manually notifying the server to start a new log file.

In setting up your logging policies, you can set options (for each service) that control limits on total log storage, maximum number of log files, individual file size, maximum file age, and rate of log-file rotation.

Planning the Options You Want

Keep in mind that you must set several limits, more than one of which might cause the rotation or deletion of a log file. Whichever limit is reached first is the controlling one. For example, if your maximum log-file size is 3.5 MB, and you specify that a new log be created every day, you may actually get log files created faster than one per day if log data builds up faster than 3.5 MB every 24 hours. Then, if your maximum number of log files is 10 and your maximum age is 8 days, you may never reach the age limit on log files because the faster log rotation may mean that 10 files will have been created in less than 8 days.

The following default values, provided for Messaging Server administration logs, may be a reasonable starting point for planning:

Maximum number of log files in a directory: 10

Maximum log-file size: 2 MB

Total maximum size permitted for all log files: 20 MB

Minimum free disk space permitted: 5 MB

Log rollover time: 1 day

Maximum age before expiration: 7 days

Level of logging: Notice

You can see that this configuration assumes that server-administration log data is predicted to accumulate at about 2 MB per day, backups are weekly, and the total space allotted for storage of admin logs is at least 25 MB. (These settings may be insufficient if the logging level is more verbose.)

For POP, IMAP or HTTP logs, the same values might be a reasonable start. If all services have approximately the same log-storage requirements as the defaults shown here, you might expect to initially plan for about 150 MB of total log-storage capacity. (Note that this is meant only as a general indication of storage requirements; your actual requirements may be significantly different.)

Understanding Logging Options

You can set options that control the message store logging configuration by the command line.

The optimal settings for these options depend on the rate at which log data accumulates. It may take between 4,000 and 10,000 log entries to occupy 1 MB of storage. At the more verbose levels of logging (such as Notice), a moderately busy server may generate hundreds of megabytes of log data per week. Here is one approach you can follow:

- Set a level of logging that is consistent with your storage limits—that is, a level that you estimate will cause log-data accumulation at approximately the rate you used to estimate the storage limit.
- Define the log file size so that searching performance is not impacted. Also, coordinate it with your rotation schedule and your total storage limit. Given the rate at which log entries accumulate, you might set a maximum that is slightly larger than what you expect to accumulate by the time a rotation automatically occurs. And your maximum file size times your maximum number of files might be roughly equivalent to your total storage limit. For example, if your IMAP log rotation is daily, your expected accumulation of IMAP log data is 3 MB per day, and your total storage limit for IMAP logs is 25 MB, you might set a maximum IMAP log-file size of 3.5 MB. (In this example, you could still lose some log data if it accumulated so rapidly that all log files hit maximum size and the maximum number of log files were reached.)
- If server backups are weekly and you rotate IMAP log files daily, you might specify a maximum number of IMAP log files of about 10 (to account for faster rotation if the individual log-size limit is exceeded), and a maximum age of 7 or 8 days.
- Pick a total storage limit that is within your hardware capacity and that coordinates with the backup schedule you have planned for the server. Estimate the rate at which you anticipate that log data will accumulate, add a factor of

safety, and define your total storage limit so that it is not exceeded over the period between server backups. For example, if you expect to accumulate an average of 3 MB of IMAP log-file data per day, and server backups are weekly, you might specify on the order of 25 - 30 MB as the storage limit for IMAP logs (assuming that your disk storage capacity is sufficient).

- For safety, pick a minimum amount of free disk space that you will permit on the volume that holds the log files. That way, if factors other than log-file size cause the volume to fill up, old log files will be deleted before a failure occurs from attempting to write log data to a full disk.

Searching and Viewing Service Logs

The log files provide the basic interface for viewing message store and administration log data. For a given service, log files are listed in chronological order. Once you have chosen a log file to search, you can narrow the search for individual events by specifying search options.

Search Options

These are useful search options you can specify for viewing log data:

- *A time period.* You can specify the beginning and end of a specific time period to retrieve events from, or you can specify a number of days (before the present) to search. You might typically specify a range to look at logged events leading up to a server crash or other occurrence whose time you know of. Alternatively, you might specify a day range to look at only today's events in the current log file.
- *A level of logging.* You can specify the logging level (see "[Logging Levels](#)" example, Critical to see why the server went down, or Error to locate failed protocol calls.
- *A facility.* You can specify the facility (see "[Categories of Logged Events](#)" that contains the problem; for example, Store if you believe a server crash involved a disk error, or Protocol if the problem lies in an IMAP protocol command error.
- *A text search pattern.* You can provide a text search pattern to further narrow the search. You can include any component of the event (see "[Understanding Service Log File Format](#)" search, such as event time, process name, process ID, and any part of the event message (such as remote host name, function name, error number, and so on) that you know defines the event or events you want to retrieve. Your search pattern can include the following special and wildcard characters:

* Any set of characters (example: *.com)

? Any single character (example: 199?)

[*nnn*] Any character in the set *nnn* (example: [aeiou])

[] Any character not in the set *nnn* (example: [Managing Logging^aeiou])

[] Any character in the range *n-m* (example: [A-Z])

[*Managing Logging^n-m*] Any character not in the range *n-m* (example: [Managing Logging^0-9])

\ Escape character: place before *, ?, [, or] to use them as literals

Note: Searches are case-sensitive.

Examples of combining logging level and facility in viewing logs might include the following:

- Specifying Account facility (and Notice level) to display failed logins, which may be useful when investigating potential security breaches
- Specifying Network facility (and all logging levels) to investigate connection problems
- Specifying all facilities (and Critical logging level) to look for basic problems in the functioning of the server

Working With Service Logs

This section describes how to work with service logs by using the **msconfig** command for searching and viewing logs.

To Send Service Logs to syslog

1. Run the **msconfig** command with the **syslogfacility** option:

```
msconfig logfile.service.syslogfacilityvalue
```

where *service* is **admin**, **pop**, **imap**, **mta**, **base** or **http** and *value* is **user**, **mail**, **daemon**, **local0** to **local7**, or none.

Once the value is set, messages are logged to the **syslog** facility corresponding to the set value and all the other log file service options are ignored. When the option is not set or the value is **none**, logging uses the Messaging Server log files.

To Disable HTTP Logging

If your system does not support HTTP message access, that is, Webmail, you can disable HTTP logging by setting the following variables. Do not set these variables if your system requires Webmail support (for example, Messenger Express).

1. Run the following **msconfig** commands:

```
msconfig  
msconfig> set http.enable 0  
msconfig set http.enablesslport 0  
msconfig write
```

To Set the Server Log Level

1. Run the following **msconfig** command:

```
msconfig set <service>.logfile.loglevel <loglevel>
```

where *service* is **admin**, **pop**, **imap**, **mta**, **base** or **http** and *loglevel* is **Nolog**, **Critical**, **Error**, **Warning**, **Notice**, **Information**, or **Debug**.

To Specify a Directory Path for Server Log Files

1. Run the following **msconfig** command:

```
msconfig <service>.logfile.logdir <dirpath>
```

To Specify a Maximum File Size for Each Service Log

1. Run the following **msconfig** command:

```
msconfig set <service>.logfile.maxlogfilesize <size>
```

where <size> specifies a number of bytes.

To Specify a Service Log Rotation Schedule

1. Run the following **msconfig** command:

```
msconfig set <service>.logfile.rollovertime <number>
```

where *number* specifies a number of seconds.

To Specify a Maximum Number of Service Log Files Per Directory

1. Run the following **msconfig** command:

```
msconfig <service>.logfile.maxlogsize <number>
```

where *number* specifies the maximum number of log files.

To Specify a Storage Limit

1. Run the following **msconfig** command:

```
msconfig _service_.logfile.maxlogsize _number_
```

where *number* specifies a number in bytes.

To Specify the Minimum Amount of Free Disk Space to Reserve

1. Run the following **msconfig** command:

```
msconfig <service>.logfile.minfreediskspace <number>
```

where *number* specifies a number in bytes.

To Specify an Age for Logs at Which They Expire

```
msconfig -o <service>.logfile.expirytime <number>
```

where *number* specifies a number in seconds.

Implementing and Configuring Message Store Transaction Logging

This section contains the following topics:

- [Overview of Message Store Transaction Logging](#)
- [Message Store Transaction Logging Log Entries](#)
- [Configuring Message Store Transaction Logging](#)
- [Message Store Transaction Log Examples](#)

Overview of Message Store Transaction Logging

You can use Message Store transaction logging to record Messaging Server user actions and events, tracing messages similar to the way the MTA traces messages. Tracing messages in this fashion allows you to track critical events of a message's life cycle.

Message Store transaction logging uses the same XML format used in the MTA **mail.log*** and **connection.log*** files. This XML format is the default. It provides the same transaction information previously logged in the process log at the **notice** and **information** log levels as well as information in **msgtrace** logging. It also includes IMAP and POP context numbers in all IMAP and POP log messages. You do not set log levels for Message Store transaction logs. You instead configure settings to determine which events and attributes to log. Message Store transaction logging uses a log roll over daemon which rolls over all logs previously written using **nslogger**. Message Store transaction logging provides a backward-compatibility mode that

behaves the way the previous store logging worked, combining process and action logging together. It provides a **msgtrace** mode also for backward compatibility.

Message Store Transaction Logging Log Entries

Message Store transaction logging supports the following user actions or events.

- li - login action
- ma - append message action
- ex - expunge message action
- fe - fetch or retrieve message action
- qc - quota change action
- qe - quota exceeded limit action
- mc - create mailbox action
- md - delete mailbox action
- mr - rename mailbox action
- ms - subscribe to mailbox action
- mu - unsubscribe to mailbox action
- ac - set mailbox access control list action
- sl: select a mailbox
- co: connect to a mail account

Message Store transaction logging actions can support the following attributes.

- ts - time stamp
- pi - process id
- si - session id
- us - canonical user name, post authentication canonical user name
- mi - message id
- ii - folder id
- tr - transport information
- if - imap flags
- fi - file name
- sz - message size
- no - node name
- tg - RFC 5423 tag
- ac - action flags
- ct - time cost
- ma - mailbox id, mailbox, URI encoded, plus uidvalidity
- om - old mailbox id
- ut - RFC 5423 uidset
- ua - sasl authentication id, RFC 5423 "admin"
- bs - body structure, RFC 5423
- dq - disk quota, RFC 5423
- du - disk used
- qt - over quota trigger
- qr - rule with over quota
- ev - message envelope
- mm - maximum messages
- mc - message count
- sq - module sequence
- sn - service name, imap, pop, admincli
- un - next uid
- uv - uidvalidity
- ui - uid
- cx - context number
- ur - IMAP URL
- uo - SASL original user identity
- at - SSL authorization type
- cs - SSL cipher suite if SSL is used
- cn - SSL server certificate nickname

cc - SSL client certificate subject
 su - mailbox specialuse flag
 ai - setacl identifier
 ar - setacl rights
 hl - message header size
 sd - number of deleted messages
 mq - message quota
 mu - module sequence in milliseconds
 uc - unchanged since in seconds
 uu - unchanged since in milliseconds
 lu - old uid
 ls - old uidset
 nt - anything worthy noting

Note: Not all attributes are available to each action. The following attributes are available for all actions:

ts: time stamp,
 pi: process id,
 sn: service name, imap, pop, admincli

An action log entry is an XML element with the following characteristics:

- The element tag is the two character action label, which follows XML syntax to start and end: <tag ... />
- The element has many attributes defined previously, and each action has own attributes.
- The attribute starts with a two character attribute label, and its value is quoted.
- The following attributes are required for each action entry: timestamp (**ts**), service name (**sn**), and process id (**pi**).

A typical log entry resembles the following:

```
<co ts="09/Feb/2015:13:44:54.059 -0800" sn="imapd" pi="7149" ac="C"
tr="[127.0.0.1:49487]" us="admin" nt="2015/2/9 13:43:52 0:01:02 158 1032 1"/>
```

where:

- **co** indicates the action label, which is connection closed.
- Three mandatory attributes, **ts** (time stamp), **sn** (service name), and **pi** (process id) are present.
- The action specific attributes are **us** (user name) and **ac** (action flags).

Configuring Message Store Transaction Logging

Transaction logging configuration is applied globally to IMAP, POP, and delivery Message Store components.

Enabling Message Store Transaction Logging

To trace transactions in the Message Store transaction log, you configure message tracing in addition to the logging configuration.

To enable Message Store transaction logging, run the following command:

```
msconfig set messagetrace.activate transactlog
```

To show the configuration, run the following command:

```
msconfig show messageTrace.activate
role.messageTrace.activate = transactlog
```

Enabling Message Store Transaction Log Actions and Attributes

The syntax to enable store actions and attributes is a matrix that you configure separately using actions and action attributes.

Table 23–10 shows the actions and their descriptions.

Table 23–10 Message Store Transaction Log Actions

Action	Description
all	All actions are enabled. This is the default.
+((li lo ma ...))	Only actions in the list are enabled. Actions are separated by a space.
-((li lo ma ...))	All actions are enabled except those in the list. Actions are separated by a space.

To enable all Message Store actions except for login and logout actions, run the following command.

```
msconfig set imap.actions '-(li lo)'
```

Running the **msconfig show imap.actions** command shows the **imap.actions** settings.

```
msconfig show imap.actions
role.imap.actions = -(li lo)
```

Table 23–11 shows the attributes and their descriptions.

Table 23–11 Message Store Attributes

Attributes	Description
all	All attributes are enabled. This is the default.
+((us mi ii ...))	Only attributes in the list are enabled. Attributes are separated by a space.
-((us mi ii ...))	All attributes are enabled except those in the list. Attributes are separated by a space.

To enable only the canonical user name and session ID attributes, run the following command.

```
msconfig set imap.actionattributes '+ (us si)'
```

Running the **msconfig show imap.actionattributes** shows the **imap.actionattributes** settings.

```
msconfig show imap.actionattributes
role.imap.actionattributes = + (us si)
```

Message Store Transaction Log Examples

The examples in this section use the following Messaging Store transaction logging settings:

```
msconfig show messageTrace
role.messageTrace.activate = transactlog
```

```
msconfig show imap.actions
role.imap.actions = all
```

```
msconfig show imap.actionattributes
role.imap.actionattributes = all
```

Example Message Store Transaction Logs for IMAP

Example 1: Log in as Joe.

```
<li ts="24/Jul/2014:13:33:04.306 -0700" sn="imapd" pi="26435" us="joe"
tr="[127.0.0.1:45046]" at="plaintext" cs="noSSL" si="0"/>
```

Example 2: Create mailbox Folder1 and Folder2.

```
<mc ts="24/Jul/2014:13:33:40.902 -0700" sn="imapd" pi="26435" us="joe"
ma="Folder1"/>
<mc ts="24/Jul/2014:13:34:05.299 -0700" sn="imapd" pi="26435" us="joe"
ma="Folder2"/>
```

Example 3: Delete mailbox Folder1.

```
<md ts="24/Jul/2014:13:34:12.515 -0700" sn="imapd" pi="26435" us="joe"
ma="Folder1"/>
```

Example 4: Append message to Folder2.

```
<ma ts="24/Jul/2014:13:46:25.271 -0700" sn="imapd" pi="27080" us=""
ma="user/joe/Folder2" sz="163" ui="1" uv="1406234045" cx="0"/>
```

Example 5: Expunge 2 messages in the Inbox.

```
<ex ts="25/Jul/2014:13:50:27.164 -0700" sn="imapd" pi="30247" ma="user/joe"
mi="0N8S00A0B007SS00@host.example.com" no="[127.0.0.1:52383]"/>
<ex ts="25/Jul/2014:13:52:43.163 -0700" sn="imapd" pi="30247" ma="user/joe"
mi="0N8R00A06ZLSS00@host.example.com" no="[127.0.0.1:52383]"/>
```

Example 6: Log out.

```
<co ts="09/Feb/2015:13:44:54.059 -0800" sn="imapd" pi="7149" ac="C"
tr="[127.0.0.1:49487]" us="admin" nt="2015/2/9 13:43:52 0:01:02 158 1032 1"/>
```

Example 7: Connect to an email account.

```
<co ts="05/Feb/2015:11:57:41.896 -0800" sn="imapd" pi="27433"
tr="[192.0.2.1:3984]" at="ssl"/>
```

Example 8: Select a folder.

```
<sl ts="05/Feb/2015:11:57:42.005 \-0800" sn="imapd" pi="27433"
tr="[192.0.2.1:3984]" ma="user/admin" us="admin"/>
```

Example 9: Rename a folder.

```
imap command: A01 RENAME Folder1 Folder2
<mr ts="05/Feb/2015:12:03:52.672 \-0800" sn="imapd" pi="27433" us="admin"
ma="Folder2" om="Folder1"/>
```

Example 10: Subscribe to a folder.

```
imap command: A01 SUBSCRIBE INBOX
<ms ts="05/Feb/2015:12:08:21.589 \-0800" sn="imapd" pi="27433" us="admin"
ma="inbox"/>
```

Example 11: Unsubscribe from a folder.

```
imap command: A01 UNSUBSCRIBE INBOX
<mu ts="05/Feb/2015:12:14:39.187 \-0800" sn="imapd" pi="3613" us="admin"
ma="inbox"/>
```

Example 12: Set acl to a mailbox.

```
imap command: A01 SETACL folder1 david lrstwiead
<ac ts="05/Feb/2015:12:06:34.287 \-0800" sn="imapd" pi="27433" us="admin"
ma="user/admin/folder1" nt="admin lrstwiepkxandc :admin lrstwiepkxandc David
lrstwiead "/>
```

Example 13: Fetch rfc822.text.

```
imap command: A01 FETCH 1 RFC822.TEXT
<fe ts="05/Feb/2015:14:00:54.865 \-0800" sn="imapd" pi="3613" us="admin"
ma="user/admin" sz="-1:0" mi="<54C145D2.7010202@shenmail.com"/>
```

Example 14: Change quota.

```
imap command: A01 SETQUOTA user/admin/Folder1 (STORAGE 512)
<qc ts="05/Feb/2015:14:15:36.059 \-0800" sn="imapd" pi="3613" us="admin"
ur="user/admin/Folder1" dq="512"/>
```

Example 15: Quota exceeds limit.

```
imquotacheck command: ./imquotacheck \-n \-l \-r rulefile
<qe ts="05/Feb/2015:15:22:54.087 \-0800" sn="imquotacheck" pi="26490" us="admin"
dq="2" du="43" mq="0" mc="0" qt="90" qr="rule3"/>
```

For more information on syntax and options, see ["imquotacheck."](#)

Example Message Store Transaction Logs for POP

Example 1: Log in as Joe.

```
<li ts="24/Jul/2014:16:59:43.187 \-0700" sn="popd" pi="27522" us="joe"
tr="[127.0.0.1:29611]" at="plaintext" cs="noSSL" si="21"/>
```

Example 2: Fetch message.

```
<fe ts="24/Jul/2014:17:00:08.205 \-0700" sn="popd" pi="27522" us="joe"
om="user/joe" mi="<0N8S00A0E09CSS00@host.example.com"/>
```

Example 3: Delete message.

```
<ex ts="24/Jul/2014:17:02:04.565 \-0700" sn="popd" pi="27522" ma="user/joe"
mi="<0N8S00A0E09CSS00@host.example.com" no="sc11136733"/>
```

Example 4: Connect to an email account.

```
<co ts="05/Feb/2015:12:01:36.719 -0800" sn="popd" pi="27443"
tr="[10.133.141.196:60395]" at=""/>
```

Example Message Store Transaction Logs for Message Delivery

Example 1: Send five messages to Joe.

```
<ma ts="24/Jul/2014:17:05:13.950 \-0700" sn="ims_master" pi="9236" us=""
ma="user/joe" sz="404652" mi="<0N9800L0UQWP9600@host.example.com" ui="5"
uv="1392927327" cx="0"/>
<ma ts="24/Jul/2014:17:05:13.954 \-0700" sn="ims_master" pi="9236" us=""
ma="user/joe" sz="404652" mi="<0N9800L0WQWP9600@host.example.com" ui="6"
```

```

uv="1392927327" cx="0"/>
<ma ts="24/Jul/2014:17:05:13.974 \-0700" sn="ims_master" pi="9236" us=""
ma="user/joe" sz="404652" mi="<0N9800L0YQWP9600@host.example.com>" ui="7"
uv="1392927327" cx="0"/>
<ma ts="24/Jul/2014:17:05:14.010 \-0700" sn="ims_master" pi="9240" us=""
ma="user/joe" sz="404652" mi="<0N9800L10QWP9600@host.example.com>" ui="8"
uv="1392927327" cx="0"/>
<ma ts="24/Jul/2014:17:05:14.043 \-0700" sn="ims_master" pi="9240" us=""
ma="user/joe" sz="404652" mi="<0N9800L12QWQ9600@host.example.com>" ui="9"
uv="1392927327" cx="0"/>

```

Other Message Store Logging Features

Messaging Server provides a feature called telemetry that can capture a user's entire IMAP or POP session into a file. This feature is useful for debugging client problems. For example, if a user complains that their message access client is not working as expected, this feature can be used to trace the interaction between the access client and Messaging Server. See in "[Checking User IMAP/POP/Webmail Session by Using Telemetry.](#)"

Message Store Logging Examples

The exact field format and list of fields logged in the Message Store log files vary according to the logging options set. This section shows a few examples of interpreting typical sorts of log entries.

Message Store Logging Example: Bad Password

When a user types an invalid password, "authentication" failure is logged, as opposed to a "user not found" message. The message "user not found" is the text passed to the client for security reasons, but the real reason (invalid password) is logged.

Example Message Store Logging: Invalid Password

```
[31/Aug/2004:16:33:14 \-0700] vadar imapd[13027]: Account Notice: badlogin:
[192.18.126.64:40718] plaintext user1 authentication failure
```

Message Store Logging Example: Account Disabled

The following example shows why a user cannot log in due to a disabled account. Furthermore, the disabled account is clarified as "(inactive)" or "(hold)."

Example Message Store Logging: Account Disabled

```
[31/Aug/2004:16:33:14 \-0700] vadar imapd[13027]: Account Notice: badlogin:
[192.18.126.64:40720] plaintext user3 account disabled (hold)
```

Message Store Logging Example: Message Appended

The following example shows an append message, which occurs when whenever a message is appended to a folder. The Message Store log records all messages entering the Message Store through the **ims_master** and **lmtip** channels. Records the "append" of user ID, folder, message size, and message ID.

Example Message Store Logging: Append

```
[31/Aug/2004:16:33:14 \-0700] vadar ims_master[13822]: Store Information:append:
user1:user/user1:659:<Roam.SIMC.2.0.6.1093995286.11265.user1@vadar.example.org>
```

Message Store Logging Example: Message Retrieved by a Client

The Message Store log writes a "fetch" message when a client retrieves a message. The Message Store log records all client fetches of at least one body part. Records the "fetch" of user ID, folder, and message-ID.

Example Message Store Logging: Message Retrieved by a Client

```
[31/Aug/2004:16:33:14 \-0700]] vadar imapd[13729]: Store Information:  
fetch:user1:user/user1:<Roam.SIMC.2.0.6.1093051161.3655.user1@vad.example.org>
```

Message Store Logging Example: Message Removed from a Folder

The following example shows how to remove a message from a folder.

The Message Store writes an "expunge" message when an IMAP or POP message is removed from a folder (but not removed from the system). It is logged whether it is expunged by the user or a utility. Records the "expunge" of folder and message ID.

Example Message Store Logging Example: Message Removed from a Folder

```
[31/Aug/2004:16:33:14 \-0700]] vadar imexpire[13923]: Store Information:  
expunge:user/user1:<Roam.SIMC.2.0.6.1090458838.2929.user1@vadar.example.org>
```

Message Store Logging Example: Duplicate Login Messages

If you configure message trace for one msgtrace log file, the normal "login" messages, which appear in the imap and pop log files, are duplicated in the msgtrace file. The following is a normal login message:

Example Message Store Logging: Login

```
[31/Aug/2004:16:33:14 \-0700]] vadar imapd[13027]: Account Information: login  
[192.18.126.64:40718] user1 plaintext
```

Using Message Store Log Messages

Use the following links to significant Message Store logging and error message pages:

- ["Overview of Logging."](#) Provides an introduction to Messaging Server logging concepts.
- ["Tools for Managing Logging."](#) Lists available tools for managing logs.
- See the discussion on logging options in the *Messaging Server Reference*.
Search for **logfile**. The information lists the configuration options that you set for Messaging Server by using the **configutil** command. You can view the documentation for the Unified Configuration equivalent of a **configutil** option by clicking on the **configutil** name link.
- ["Managing Message Store, Admin, and Default Service Logs."](#) Describes logging for the Message Store (POP, IMAP, and HTTP), Admin, and Default services.

MMP Logging

The logging format used by the MMP is unstable and subject to change at any time. It uses the **nslog** model and syntax used by the store, with similar configuration settings. The most interesting line in the MMP log is:

```
[04/Feb/2016:10:42:24 -0800] myhostname ImapProxyASvc.cfg[17242]: General  
Notice: (id 3) Proxy connect for client 10.0.0.5:55159 (cleartext) canonical  
user 'admin@host.example.com' original user 'admin' auth 'PLAIN' via MMP  
127.0.0.1:55162 to backend 127.0.0.1 (cleartext)
```

which can be used to correlate a connection from a client to a back end connection.

For debugging the MMP, the loglevel may be raised and the **debugkeys** option can be helpful.

For additional information, see the discussion on MMP in the *Messaging Server Reference*.

Monitoring Messaging Server

In most cases, a well-planned, well-configured server will perform without extensive intervention from an administrator. As an administrator, however, it is your job to monitor the server for signs of problems.

Topics:

- [Automatic Monitoring and Restart](#)
- [Daily Monitoring Tasks](#)
- [Utilities and Tools for Monitoring](#)

See also:

- [Monitoring User Access to the Message Store](#)
- [Monitoring System Performance](#)
- [Monitoring Disk Space](#)
- [Monitoring the MTA](#)
- [Monitoring LDAP Directory Server](#)
- [Monitoring the Message Store](#)
- [SNMP Support](#)

Automatic Monitoring and Restart

Oracle Communications Messaging Server provides a way to transparently monitor services and automatically restart them if they crash or become unresponsive (the services hangs or freeze up). It can monitor all message store, MTA, and MMP services including the IMAP, POP, HTTP, job controller, dispatcher, and MMP servers. It does not monitor other services such as SMS or TCP/SNMP servers. (TCP/SNMP is monitored by the job controller.) Refer to "[Automatic Restart of Failed or Unresponsive Services](#)" and "[Monitoring Using msprobe and watcher Functions](#)."

Daily Monitoring Tasks

The most important tasks you should perform on a daily basis are checking postmaster mail, monitoring the log files, and setting up the **stored** utility. These tasks are described below.

Checking Postmaster Mail

Messaging Server has a predefined administrative mailing list set up for postmaster email. Any users who are part of this mailing list will automatically receive mail addressed to postmaster.

The rules for postmaster mail are defined in RFC822, which requires every email site to accept mail addressed to a user or mailing list named postmaster and that mail sent to this address be delivered to a real person. All messages sent to postmaster@host.domain are sent to a postmaster account or mailing list.

Typically, the postmaster address is where users should send email about their mail service. As postmaster, you might receive mail from local users about server response time, from other server administrators who are encountering problems sending mail to your server, and so on. You should check postmaster mail daily.

You can also configure the server to send certain error messages to the postmaster address. For example, when the MTA cannot route or deliver a message, you can be notified via email sent to the postmaster address. You can also send exception condition warnings (low disk space, poor server response) to postmaster.

Monitoring and Maintaining the Log Files

Messaging Server creates a separate set of log files for each of the major protocols or services it supports including SMTP, IMAP, POP, and HTTP. These are located in *MessagingServer_home/data/log*. You should monitor the log files on a routine basis--especially if you are having problems with the server.

Be aware that logging can impact server performance. The more verbose the logging you specify, the more disk space your log files will occupy for a given amount of time. You should define effective but realistic log rotation, expiration, and backup policies for your server. For information about defining logging policies for your server, see "[Managing Logging](#)."

Setting Up the msprobe Utility

The **msprobe** utility automatically performs monitoring and restart functions. For further information see "[Monitoring Using msprobe and watcher Functions](#)."

Utilities and Tools for Monitoring

The following tools are available for monitoring:

- [immonitor-access](#)
- [imcheck](#)
- [Log Files](#)
- [imsimta counters](#)
- [imsimta qm counters](#)
- [MTA Monitoring Using SNMP](#)
- [Monitoring Using msprobe and watcher Functions](#)

immonitor-access

"[immonitor-access](#)" monitors the status of the following Messaging Server components/processes: Mail Delivery (SMTP server), Message Access and Store (POP and IMAP servers), Directory Service (LDAP server) and HTTP server. This utility

measures the response times of the various services and the total round trip time taken to send and retrieve a message. The Directory Service is monitored by looking up a specified user in the directory and measuring the response time. Mail Delivery is monitored by sending a message (SMTP) and the Message Access and Store is monitored by retrieving it. Monitoring the HTTP server is limited to finding out whether or not it is up and running.

For complete instructions, refer to ["immonitor-access."](#)

imcheck

Use ["imcheck"](#) to monitor database statistics including logs and transactions.

counterutil

This utility provides statistics acquired from different system counters. For details, see ["Gathering Message Store Counter Statistics by Using counterutil."](#)

Log Files

Messaging server logs event records for SMTP, IMAP, POP, and HTTP. The policies for creating and managing the Messaging Server log files are customizable.

Since logging can affect the server performance, logging should be considered very carefully before the burden is put on the server. Refer to ["Managing Logging"](#) for more information.

imsimta counters

The MTA accumulates message traffic counters based upon the Mail Monitoring MIB, RFC 1566 for each of its active channels. The channel counters are intended to help indicate the trend and health of your email system. Channel counters are not designed to provide an accurate accounting of message traffic. For precise accounting, instead see MTA logging as discussed in ["Managing Logging."](#)

The MTA channel counters are implemented using the lightest weight mechanisms available so that they cause as little impact as possible on actual operation. Channel counters do not try harder: if an attempt to map the section fails, no information is recorded. If one of the locks in the section cannot be obtained almost immediately, no information is recorded. When a system is shut down, the information contained in the in-memory section is lost forever.

The **imsimta counters -show** command provides MTA channel message statistics (see below). These counters need to be examined over time noting the minimum values seen. The minimums may actually be negative for some channels. A negative value means that there were messages queued for a channel at the time that its counters were zeroed (for example, the cluster-wide database of counters created). When those messages were dequeued, the associated counters for the channel were decremented and therefore leading to a negative minimum. For such a counter, the correct "absolute" value is the current value less the minimum value that counter has ever held since being initialized.

```
Channel Messages Recipients Blocks
-----
tcp_local
Received 29379 79714 982252 (1)
Stored 61 113 -2004 (2)
Delivered 29369 79723 983903 (29369 first time) (3)
Submitted 13698 13699 18261 (4)
Attempted 0 0 0 (5)
Rejected 1 10 0 (6)
```

```
Failed 104 104 4681 (7)
Queue time/count 16425/29440 = 0.56 (8)
Queue first time/count 16425/29440 = 0.56 (9)
Total In Assocs 297637
Total Out Assocs 28306
```

1)Received is the number of messages enqueued to the channel named **tcp_local**. That is, the messages enqueued (E records in the **mail.log*** file) to the **tcp_local** channel by any other channel.

2)Stored is the number of messages stored in the channel queue to be delivered.

3)Delivered is the number of messages which have been processed (dequeued) by the channel **tcp_local**. (That is, D records in the **mail.log*** file.) A dequeue operation may either correspond to a successful delivery (that is, an enqueue to another channel), or to a dequeue due to the message being returned to the sender. This will generally correspond to the number **Received** minus the number **Stored**.

The MTA also keeps track of how many of the messages were dequeued upon first attempt; this number is shown in parentheses.

4)Submitted is the number of messages enqueued (E records in the **mail.log** file) by the channel **tcp_local** to any other channel.

5)Attempted is the number of messages which have experienced temporary problems in dequeuing, that is, Q or Z records in the **mail.log*** file.

6)Rejected is the number of attempted enqueues which have been rejected, that is, J records in the **mail.log*** file.

7)Failed is the number of attempted dequeues which have failed, that is, R records in the **mail.log*** file.

8)Queue time/count is the average time-spent-in-queue for the delivered messages. This includes both the messages delivered upon the first attempt, see (9), and the messages that required additional delivery attempts (hence typically spent noticeable time waiting fallow in the queue).

9)Queue first time/count is the average time-spent-in-queue for the messages delivered upon the first attempt.

Note that the number of messages submitted can be greater than the number delivered. This is often the case, since each message the channel dequeues (delivers) will result in at least one new message enqueued (submitted) but possibly more than one. For example, if a message has two recipients reached via different channels, then two enqueues will be required. Or if a message bounces, a copy will go back to the sender and another copy may be sent to the postmaster. Usually that will be two submissions (unless both are reached through the same channel).

More generally, the connection between **Submitted** and **Delivered** varies according to type of channel. For example, in the conversion channel, a message would be enqueued by some other arbitrary channel, and then the conversion channel would process that message and enqueue it to a third channel and mark the message as dequeued from its own queue. Each individual message takes a path:

```
elsewhere -> conversion E record Received
conversion -> elsewhere E record Submitted
conversion D record Delivered
```

However, for a channel such as **tcp_local** which is not a "pass through," but rather has two separate pieces (slave and master), there is no connection between **Submitted** and **Delivered**. The **Submitted** counter has to do with the SMTP server portion of the **tcp_**

local channel, whereas the **Delivered** counter has to do with the SMTP client portion of the **tcp_local** channel. Those are two completely separate programs, and the messages travelling through them may be completely separate.

Messages submitted to the SMTP server:

tcp_local -> elsewhere E record Submitted

Messages sent out to other SMTP hosts via the SMTP client:

```
elsewhere -> tcp_local E record Received
tcp_local D record Delivered
```

Channel dequeues (delivers) will result in at least one new message enqueued (submitted) but possibly more than one. For example, if a message has two recipients reached via different channels, then two dequeues will be required. Or if a message bounces, a copy will go back to the sender and another copy may be sent to the postmaster. Usually that will be reached through the same channel.

imsimta counters Implementation

For performance reasons, a node running the MTA keeps a cache of channel counters in memory using a shared memory section. As processes on the node enqueue and dequeue messages, they update the counters in this in-memory cache. If the in-memory section does not exist when a channel runs, the section will be created automatically. (The **imta start** command also creates the in-memory section, if it does not exist.)

The command **imta counters -clear** or the **imta qm** command **counters clear** may be used to reset the counters to zero.

imsimta qm counters

The **imsimta qm counters** utility displays MTA channel queue message counters. You must be root or **mailsrv** to run this utility. The output fields are the same as those described in "[imsimta counters](#)."

Example:

```
imsimta counters -create
imsimta qm counters show
Channel Messages Recipients Blocks
-----
tcp_intranet
Received 13077 13859 264616
Stored 92 91 -362
Delivered 12985 13768 264978
Submitted 2594 2594 3641
...
```

Every time you restart the MTA, you must run: **imsimta counters -create**

imsimta qm summarize

The **imsimta qm summarize** utility displays a summary of the number of messages and their status in the MTA channel queues.

For more details of the various switches available, see the **summarize** sub-command in "[imsimta qm](#)" and the **imsimta qm help summarize** command.

qm summarize modes

Like many of the **qm** sub-commands, **summarize** has two modes of operation: The **-directory_tree** mode examines the message files in the MTA queue directories on disk. The **-database** mode queries the **job_controller** process's in-memory database structures. The directory mode creates a heavier load on the IO system and may not reflect what **job_controller** is actually working on, but it can be useful to know if there is a difference between the two. The job controller makes the decisions about which messages are tried next, so the database mode will be more useful.

```

imsimta qm
qm.maint> sum -directory_tree
Channel Messages Size (Mb)
-----
conversion 0 0.0
hold 0 0.0
ims-ms (stopped) 2 0.0
process 0 0.0
reprocess 0 0.0
tcp_intranet (stopped) 0 0.0
tcp_local (stopped) 2 0.0
-----
Totals 4 0.0

```

Notice that the **-database** mode breaks down the number messages into three categories. **Active** messages are currently being tried by a worker process. **Pending** messages are ready to be tried by a worker as soon as thread/slot is available. **Delayed** messages have been tried before and are waiting for a specified time to be tried again as per the **backoff** option for that channel.

```

qm.maint> sum -database
Total Total
Channel Messages = Active + Pending + Delayed Size (Mb)
-----
conversion 0 0 0 0 0.0
hold 0 0 0 0 0.0
ims-ms (stopped) 2 0 2 0 0.0
l 0 0 0 0 0.0
process 0 0 0 0 0.0
reprocess 0 0 0 0 0.0
tcp_intranet (stopped) 0 0 0 0 0.0
tcp_local (stopped) 2 0 2 0 0.0
-----
Totals 4 0 4 0 0.0

```

Note: In these examples, some channels had been stopped using **imsimta qm stop <channel>** to provide some data to look at.

Held messages

A **.HELD** message file is a message which has encountered a loop or otherwise been *sidelined* and requires administrative intervention for some reason. You can see such messages using the **-held** switch. Note that **job_controller** will have no knowledge of held messages, therefore the **-database** and **-held** switches are mutually exclusive. For more information about **.HELD** messages see "[Diagnosing and Cleaning up .HELD Messages.](#)"

```

qm.maint> sum -held -database
%QM-E-CMDERR, Conflicting parameters and/or qualifiers: (DATABASE AND HELD)

qm.maint> sum -held

```

```

Held Held
Channel Messages Size (Mb) Oldest Queued Messages Size (Mb) Oldest Held
-----
conversion 0 0.0 0 0.0
hold 0 0.0 1 0.0 23 Apr, 21:35:16
ims-ms (stopped) 2 0.0 6 Apr, 13:24:00 0 0.0
process 0 0.0 0 0.0
reprocess 0 0.0 0 0.0
tcp_intranet (stopped) 0 0.0 0 0.0
tcp_local (stopped) 2 0.0 5 May, 10:16:08 0 0.0
-----
Totals 4 0.0 6 Apr, 13:24:00 1 0.0 23 Apr, 21:35:16

```

Displaying Summary by Destination Host

The **-hosts** switch displays a breakdown of the messages in the queue by destination host for channels where that is meaningful. This information is stored in the **job_controller** process in-memory queue cache database. Therefore **-hosts** implies **-database**.

```

qm.maint> sum -hosts
Total Total
Channel Host Messages = Active + Pending + Delayed Size (Mb)
-----
conversion 0 0 0 0 0.0
hold 0 0 0 0 0.0
ims-ms (stopped) 2 0 2 0 0.0
l 0 0 0 0 0.0
process 0 0 0 0 0.0
reprocess 0 0 0 0 0.0
tcp_intranet (stopped) 0 0 0 0 0.0
tcp_local (stopped) 2 0 2 0 0.0
aol.com 1 0 1 0 0.0
sun.com 1 0 1 0 0.0
-----
Totals 4 0 4 0 0.0

```

imsimta qm jobs

After starting the tcp_local channel:

```

tcp_local 1 1 0 0 0.0
aol.com 1 1 0 0 0.0

```

And to see what processes are working on what jobs:

```

qm.maint> jobs tcp_local
tcp_local channel:

Pending: 0 jobs
Active: 1 jobs, 1 messages (0.00 Mb), 1 recipients
Current jobs have delivered 1 messages, requeued 0 messages

Active jobs and messages:

22157: 1 messages (0.00 Mb), 1 recipients
1 messages processed and 0 requeued

```

Active hosts:

aol.com

Active messages:

ZZg0u410_P_~1.01 (1.0 Kb)

MTA Monitoring Using SNMP

Messaging Server supports system monitoring through the Simple Network Management Protocol (SNMP). Using an SNMP client (sometimes called a *network manager*) such as Sun Net Manager or HP OpenView (not provided with this product), you can monitor certain parts of the Messaging Server. Refer to "[SNMP Support](#)" for details.

Monitoring Using msprobe and watcher Functions

Messaging Server provides two processes, **watcher** and **msprobe** to monitor various system services. **watcher** watches for server crashes and restarts them as necessary. **msprobe** monitors server hangs (unresponsiveness). Specifically **msprobe** monitors the following:

- *Server Response Time.* **msprobe** connects to the enabled servers using their protocol commands and measures their response times. If the response time exceeds the alarm warning threshold, an alarm message is sent (see "[Alarm Messages](#)") to a server, or the server response time exceeds a specified timeout period, the server is restarted. Server response times are recorded in a counter database and is logged to the default log file. **counterutil** can be used to display the server response time statistics ("[counterutil](#)").

The following servers are monitored by **msprobe**: **imap**, **pop**, **http**, **cert**, **job_controller**, **smtp**, **lmtpl**, **mmp** and **ens**. When **smtp** or **lmtpl** are not responding, the dispatcher is restarted. **ens** cannot be automatically restarted.

- *Disk usage.* **msprobe** checks the disk availability and usage for every message store partition. Specifically it checks the message store **mboxlist** database directory and the MTA queue directory. If disk usage exceeds a configured threshold, an alarm message is sent. The disk sizes and usages are recorded in a counter database and is logged to the default log file. Administrators can use the **counterutil** utility (see "[counterutil](#)") to display the disk usage statistics.
- *Message Store mboxlist Database Log File Accumulation.* Log file accumulation is an indication of an **mboxlist** database error. **msprobe** counts the number of active log files and if the number of active log files is larger than the threshold, **msprobe** logs a critical error message to the **default** log file to inform the admin to restart the server. If the autorestart is enabled (**local.autorestart** to **yes**), the store daemon is restarted.

watcher and **msprobe** are controlled by the **msconfig** options shown in [Table 24–1](#). Further information can be found in "[Automatic Restart of Failed or Unresponsive Services](#)."

Table 24–1 *msprobe and watcher msconfig Options*

Options	Description
base.autorestart.enable	Enable automatic server restart. Automatically restarts failed or hung services. Default: 1

Table 24–1 (Cont.) msprobe and watcher msconfig Options

Options	Description
base.autorestart.timeout	Failure retry time-out. If a server fails more than twice in this designated amount of time, then the system stops trying to restart the server. The value (set in seconds) should be set to a period value longer than the msprobe interval (schedule.task:msprobe). Default: 600 seconds
msprobe.probe:service.timeout	Timeout for a specific server before restart. <i>service</i> can be imap, pop, http, cert, job_controller, smtp, lmtp, mmp or ens. Default: use msprobe.timeout
msprobe.probe:service.warningthreshold	Number of seconds of a specific server's non-response before a warning message is logged to default log file. <i>service</i> can be imap, pop, http, cert, job_controller, smtp, lmtp, mmp or ens. Default: Use msprobe.warningthreshold
msprobe.warningthreshold	Number of seconds of server non-response before a warning message is logged to default log file. Default: 25 secs
msprobe.queuedir	MTA queue directory to check if queue size exceeded threshold defined by alarm.system:diskavail.thresholddirection . Default: none
msprobe.timeout	Period of server non-response before restarting that server. See schedule.task:msprobe.crontab . Default: 30 seconds
schedule.task:msprobe.crontab	msprobe run schedule. A crontab style schedule string (see schedule.task:expire.enable in "Expire and Purge Log and Scheduling Options"). Note that by default, this is automatically set. See "Pre-defined Automatic Tasks." To disable: set schedule.task:msprobe.enable to 0.
watcher.enable	Enable watcher which monitors service failures. (IMAP, POP, HTTP, job controller, dispatcher, message store (stored), imsched , and MMP. (LMTP/SMTP servers are monitored by the dispatcher and LMTP/SMTP clients are monitored by the job_controller.) Logs error messages to the default log file for specific failures. Default: 1

Alarm Messages

msprobe can issue alarms in the form of email messages to the postmaster (see "To Monitor **imapd**, **popd** and **httpd**") warning of a specified condition. A sample email alarm sent when a certain threshold is exceeded is shown below:

```
Subject: ALARM: server response time in seconds of "ldap_example.com_389" is 10
Date: Tue, 17 Jul 2001 16:37:08 -0700 (PDT)
From: postmaster@example.com
To: postmaster@example.com
Server instance: /opt/SUNWmsgsr
Alarmid: serverresponse
Instance: ldap_example_europa.com_389
Description: server response time in seconds
Current measured value (17/Jul/2001:16:37:08 -0700): 10
Lowest recorded value: 0
Highest recorded value: 10
Monitoring interval: 600 seconds
Alarm condition is when over threshold of 10
Number of times over threshold: 1
```

You can specify how often **msprobe** monitors disk and server performance, and under what circumstances it sends alarms. This is done by using the **msconfig** command to set the alarm options. Table 24–2 shows some useful alarm options along with their

default setting. See the *Messaging Server Reference* for a complete list of all options.

Table 24–2 Useful Alarm Message *msconfig* Options

Option	Description (Default in Parenthesis)
alarm.noticehost	(localhost) Machine to which you send warning messages.
alarm.noticeport	(25) The SMTP port to which to connect when sending alarm message.
alarm.noticercpt	(Postmaster@localhost) Whom to send alarm notice.
alarm.noticesender	(Postmaster@localhost) Address of sender the alarm.
alarm.system:diskavail.description	(Percentage mail partition disk space available.) Text for description field for disk availability alarm.
alarm.system:diskavail.statinterval	(3600) Interval in seconds between disk availability checks. Set to 0 to disable checking of disk usage.
alarm.system:diskavail.threshold	(10) Percentage of disk space availability below which an alarm is sent.
alarm.system:diskavail.thresholddirection	(-1) Specifies whether the alarm is issued when disk space availability goes below threshold (-1) or above it (1).
alarm.system:diskavail.warninginterval	(24). Interval in hours between subsequent repetition of disk availability alarms.
alarm.system:serverresponse.description	(Server response time in seconds.) Text for description field for servers response alarm.
alarm.system:serverresponse.statinterval	(600) Interval in seconds between server response checks. Set to 0 to disable checking of server response.
alarm.system:serverresponse.threshold	(10) If server response time in seconds exceeds this value, alarm issued.
alarm.system:serverresponse.thresholddirection	(1) Specifies whether alarm is issued when server response time is greater than (1) or less than (-1) the threshold.
alarm.system:serverresponse.warninginterval	(24) Interval in hours between subsequent repetition of server response alarm.

Monitoring the MTA

This chapter describes how to monitor the Message Transfer Agent (MTA).

Topics:

- [Monitoring the Size of the Message Queues](#)
- [Checking for Held messages](#)
- [Monitoring Rate of Delivery Failure](#)
- [Monitoring Inbound SMTP Connections](#)
- [Monitoring the Dispatcher and Job Controller Processes](#)

Monitoring the Size of the Message Queues

Excessive message queue growth may indicate that messages are not being delivered, are being delayed in their delivery, or are coming in faster than the system can deliver them. This may be caused by a number of reasons such as a denial of service attack caused by huge numbers of messages flooding your system, or the Job Controller not running.

See "[Channel Message Queues](#)," "[Messages Are Not Dequeued](#)," and "[MTA Messages Are Not Delivered](#)" for more information on message queues.

Symptoms of Message Queue Problems

- Disk space usage grows.
- User not receiving messages in a reasonable time.
- Message queue sizes are abnormally high.

To Monitor the Size of the Message Queues

Probably the best way to monitor the message queues is to use "[imsimta counters](#)" and `imsimta qm summarize`.

You can also monitor the number of files in the queue directories (*MessagingServer_home/data/queue/*). The number of files will be site-specific, and you'll need to build a baseline history to find out what is "too many." This can be done by recording the size of the queue files over a two week period to get an approximate average.

Checking for Held messages

If the MTA detects a message is looping, it will be **sidelined** by renaming the queue message file to **.HELD**. For more discussion of how messages can become **.HELD** and what to do about them, refer to "[Diagnosing and Cleaning up .HELD Messages](#)." To see whether there are any held messages, use the **imsimta qm summarize -held** command described in [Table 8–20](#), "[summarize Sub-Command Options](#)".

Monitoring Rate of Delivery Failure

A delivery failure is a failed attempt to deliver a message to an external site. A large increase in rate of delivery failure can be a sign of a network problem such as a dead DNS server or a remote server timing out on responding to connections.

Symptoms of Rate of Delivery

There are no outward symptoms. Lots of **Q** records will appear in to **mail.log_current**.

To Monitor the Rate of Delivery Failure

Delivery failures are recorded in the MTA logs with the logging entry code **Q**. Look at the record in the file *MessagingServer_home/data/log/mail.log_current*. Example:

```
mail.log:06-Oct-2003 00:24:03.66 501d.0b.9 ims-ms Q 5 durai.balusamy@Sun.COM
rfc822;durai.balusamy@Sun.COM durai@ims-ms-daemon
<00ce01c38bda$c7e2b240$6501a8c0@guindy>Mailbox is busy
```

Monitoring Inbound SMTP Connections

An unusual increase in the number of inbound SMTP connections from a given IP address may indicate:

- An external user is trying to relay mail.
- An external user is trying to do a service denial attack.

Symptoms of Unauthorized SMTP Connections

- **External user relaying mail** : No outward symptoms.
- **Service denial attack**: External attempt to overload the SMTP servers with message requests.

To Monitor Inbound SMTP Connections

- **External user relaying mail**: Look in *MessagingServer_home/log/mail.log_current* for records with the logging entry code **J** (rejected relays). To turn on logging of remote IP addresses run the following command: **msconfig set log_connection 1**
Note that there is a slight performance trade-off when this feature is enabled.
- **Service denial attack**: To find out who and how many users are connecting to the SMTP servers, you can run the command **netstat** and check for connections at the SMTP port (default: 25). Example:

```
Local address Remote address State
192.18.79.44.25 192.18.78.44.56035 32768 0 32768 0 CLOSE_WAIT
192.18.79.44.25 192.18.136.54.57390 8760 0 24820 0 ESTABLISHED
192.18.79.44.25 192.18.26.165.48508 33580 0 24820 0 TIME_WAIT
```

Note that you will first need to determine the appropriate number of SMTP connections and their states (**ESTABLISHED**, **CLOSE_WAIT**, etc.) for your system to determine if a particular reading is out of the ordinary. If you find many connections staying in the **SYN_RECEIVED** state this might be caused by a broken network or a denial of service attack. In addition, the lifetime of an SMTP server process is limited. This is controlled by the MTA Dispatcher configuration option **MAX_LIFE_TIME**. The default is 86,400 seconds (one day). Similarly, **MAX_LIFE_CONNS** specifies the maximum number of connections a server process can handle in its lifetime. If you find a particular SMTP server that has around for a long time you may wish to investigate.

Monitoring the Dispatcher and Job Controller Processes

The Dispatcher and Job Controller Processes must be operating for MTA to work. You should have one process of each kind.

Symptoms of Dispatcher and Job Controller Processes Down

If the Dispatcher is down or does not have enough resources, SMTP connections are refused. If the Job Controller is down, queue size will grow.

To Monitor Dispatcher and Job Controller Processes

Check to see that the processes called **dispatcher** and **job_controller** exist. See "[Check that the Job Controller and Dispatcher Are Running.](#)"

SNMP Support

Oracle Communications Messaging Server supports system monitoring through the Simple Network Management Protocol (SNMP). Using an SNMP client (sometimes called a *network manager*) such as Sun Net Manager or HP OpenView (not provided with this product), you can monitor certain parts of the Messaging Server. For more information on monitoring the Messaging Server refer to "[Monitoring Messaging Server](#)."

This information describes how to enable SNMP support for the Messaging Server. It also gives an overview of the type of information provided by SNMP. Note that it does not describe how to view this information from an SNMP client. Refer to your SNMP client documentation for details on how to use it to view SNMP-based information.

This information also describes some of the data available from the Messaging Server SNMP implementation, but complete MIB details are available from RFC 2788 (<http://www.faqs.org/rfcs/rfc2788.html>) and RFC 2789 (<http://www.faqs.org/rfcs/rfc2789.html>).

Topics:

- [SNMP Implementation](#)
- [Configuring SNMP Support for Oracle Solaris 10](#)
- [Monitoring from an SNMP Client](#)
- [SNMP Information from the Messaging Server](#)

SNMP Implementation

Messaging Server implements two standardized MIBs, the Network Services Monitoring MIB (RFC 2788) and the Mail Monitoring MIB (RFC 2789). The Network Services Monitoring MIB provides for the monitoring of network services such as POP, IMAP, HTTP, and SMTP servers. The Mail Monitoring MIB provides for the monitoring of MTAs. The Mail Monitoring MIB allows for monitoring both the active and historical state of each MTA channel. The active information focuses on currently queued messages and open network connections (for example, counts of queued messages, source IP addresses of open network connections), while the historical information provides cumulative totals (for example, total messages processed, total inbound connections).

Note: For a complete listing of Messaging Server SNMP monitoring information, refer to RFC 2788 and RFC 2789.

SNMP is supported on platforms running Oracle Solaris and Linux. Messaging Server on the Oracle Solaris 9 Operating System uses Solstice Enterprise Agents (SEA). Starting with the Oracle Solaris 10 Operating System, Messaging Server supports the open source Net-SNMP monitoring framework, relegating the Oracle Solaris 9 Solstice Enterprise Agents (SEA) technology to legacy (end of support life) status. Additionally, Net-SNMP is widely used on Linux platforms. Messaging Server uses its Net-SNMP-based SNMP subagent on Oracle Solaris 10 and later as well as Linux platforms.

With the adoption of the Net-SNMP framework, Messaging Server's SNMP subagent provides new functionality:

- Support for SNMP versions 2c and 3. This support is provided by the Net-SNMP framework. The former SNMP technology, Solstice Enterprise Agents, only provided support for SNMP version 1. Enhanced security features and access controls are the primary benefit of these two versions of SNMP.
- The subagent may be configured to run as a "standalone" SNMP agent. This provides sites with additional means of isolating their various SNMP agents running on the same system.
- Multiple "instances" of Messaging Server running on the same system may concurrently be monitored. This support is provided through either the second item in this list, or through the use of SNMP version 3 "context names". This allows for SNMP monitoring of Messaging Server in failover clusters.


Limitations of the Messaging Server SNMP support are as follows:

- Only one instance of Messaging Server per host computer can be monitored via SNMP on Oracle Solaris 9.
- The SNMP support is for monitoring only. No SNMP management is supported.
- No SNMP traps are implemented. (RFC 2788 provides similar functionality without using traps.)

SNMP Operation in the Messaging Server

The Messaging Server SNMP process is an SNMP subagent which, upon startup, registers itself with the platform's native SNMP master agent. SNMP requests from clients go to the master agent. The master agent then forwards any requests destined for the Messaging Server to the Messaging Server subagent process. The Messaging Server subagent process then processes the request and relays the response back to the client via the master agent. This process is shown in the following figure.

Figure 26-1 SNMP Information Flow



Configuring SNMP Support for Oracle Solaris 10

Note: SNMP is not supported on Solaris 11. For a workaround, contact support.

By default, SNMP monitoring is disabled within Messaging Server. This default is chosen in an attempt to minimize the number of services presented by a default Messaging Server configuration. Do not interpret this default as meaning that there is a performance penalty incurred by using SNMP monitoring. Indeed, Messaging Server's SNMP support consumes very little resources and is intended to have minimal impact upon Messaging Server. The upshot of all of this is, of course, that one time configuration steps are required before using Messaging Server's SNMP support. Additionally, the default configuration of the platform's Net-SNMP master agent, **snmpd**, typically needs to be changed to run subagents such as Messaging Server's. This change is the topic of the next section.

Net-SNMP Configuration

Messaging Server's Net-SNMP based SNMP subagent uses the AgentX protocol to communicate with the platform's SNMP master agent (RFC 2741). The Net-SNMP master agent, **snmpd**, must be configured to permit the use of the AgentX protocol. To do this, ensure that the platform's **snmpd.conf** file contains the following line:

```
master agentx
```

If that line is not present, then add it and then restart the **snmpd** daemon. Sending a SIGHUP signal to the daemon is not sufficient. Once the **snmpd** daemon has been restarted, look for the UNIX domain socket which **snmpd** creates for AgentX communications. On Oracle Solaris and Linux systems, this socket by default appears as the special file **/var/agentx/master**. However, its location and name may be changed in the **snmpd.conf** file.

The Oracle Solaris 10 **snmpd** configuration is as follows:

```
% cp /etc/sma/snmp/snmpd.conf /etc/sma/snmp/snmpd.conf.save
% cat >> /etc/sma/snmp/snmpd.conf
# Messaging Server's subagent requires the AgentX protocol
master agentx
^D
% cat >> /etc/sma/snmp/snmpd.conf
% ls -al /var/agentx/
srwxrwxrwx 1 root root 0 Aug 9 13:58 /var/agentx/master
```

Additionally, on Red Hat Enterprise Linux AS 3 systems, the default **snmpd.conf** file restricts the information which may be viewed by the "public" SNMP community. It is therefore necessary to either remove that restriction or to extend it to include the MIBs served out by Messaging Server's subagent. For initial testing, perform the later. This is accomplished by including the OID subtrees mib-2.27 and mib-2.28 in a view named "systemview" as shown in the following example. For actual deployment, each site must take their overall security policy into consideration. Note that the information provided by the SNMP subagent is "read only".

```
% cp /etc/snmp/snmpd.conf /etc/snmp/snmpd.conf.save
% cat >>/etc/snmp/snmpd.conf
# Messaging Server's subagent requires the AgentX protocol
master agentx
# Messaging Server's subagent exports mib-2.27 and .28
# Add the mib-2.27 and .28 OID subtrees to the systemview
view systemview included .1.3.6.1.2.1.27
view systemview included .1.3.6.1.2.1.28
^D
% /bin/service snmpd restart
% ls -al /var/agentx/master
srwxr-xr-x 1 root root 0 Aug 8 21:20 /var/agentx/master
```

If you are using SNMP v3 context names to distinguish between the MIBs of different instances of Messaging Server concurrently running on the same host computer, then you also need to configure at least one SNMP v3 username and password for use with your SNMP v3 queries.

Messaging Server Subagent Configuration

For basic operation of Messaging Server's SNMP subagent, you need only enable it and issue a one time manual start command. Henceforth, whenever Messaging Server is started or stopped, the subagent will likewise be started or stopped. The necessary commands to effect this configuration on both Oracle Solaris and Linux systems are as follows:

```
% ./msconfig set snmp.enable 1
% ./start-msg snmp
```

Once SNMP is running, you can test the subagent from the command line with the **snmpwalk** command. See the screen shots below for an example appropriate to Oracle Solaris and Linux. Note that the files **rfc2248.txt** and **rfc2249.txt** are copies of the Network Services and MTA MIBs. On Oracle Solaris systems, these files may also be found in the **/etc/sma/snmp/mibs/** directory under the names **NETWORK-SERVICES-MIB.txt** and **MTA-MIB.txt**. It is not necessary provide these files to the **snmpwalk** tool. However, doing so permits **snmpwalk** to print names for each of the MIB variables rather than their numeric object identifiers (OIDs).

Basic testing on Oracle Solaris:

```
% d=/opt/sun/comms/messaging64/examples/mibs /usr/sfw/bin/snmpwalk -v 1 -c public
\
-m +$D/rfc2248.txt:$D/rfc2249.txt 127.0.0.1 mib-2.27

NETWORK-SERVICES-MIB::applName.1 = STRING: /opt/SUNWmsgsr MTA on mail.example.com
...
% D=/opt/sun/comms/messaging64/examples/mibs /usr/sfw/bin/snmpwalk -v 1 -c public
\
-m +$D/rfc2248.txt:$D/rfc2249.txt 127.0.0.1 mib-2.28

MTA-MIB::mtaReceivedMessages.1 = Counter32: 1452
MTA-MIB::mtaStoredMessages.1 = Gauge32: 21
...
```

Basic testing on Linux:

```
% export d=/opt/sun/messaging/examples/mibs
% /usr/bin/snmpwalk -v 1 -c public \
-m +$D/rfc2248.txt:$D/rfc2249.txt 127.0.0.1 mib-2.27
NETWORK-SERVICES-MIB::applName.1 = STRING: /opt/sun/messaging MTA on
mail.example.com
...
% /usr/bin/snmpwalk -v 1 -c public \
-m +$D/rfc2248.txt:$D/rfc2249.txt 127.0.0.1 mib-2.28
MTA-MIB::mtaReceivedMessages.1 = Counter32: 21278
MTA-MIB::mtaStoredMessages.1 = Gauge32: 7
...
```

Running as a Standalone SNMP Agent

Before configuring Messaging Server's SNMP subagent to run as a standalone SNMP agent, you must first decide which Ethernet interface and UDP port to use to listen for SNMP requests. By default, it listens on all available Ethernet interfaces by using UDP port 161. In most cases, you should change the port number so as to not interfere with the platform's SNMP master agent, **snmpd**. In some circumstances such as HA failover, you should change the Ethernet interface from all available interfaces - **INADDR_ANY** - to a specific interface identified by its IP address. These two concepts, Ethernet interface and UDP port, are controlled by the **snmp.listenaddr** and **snmp.port** options.

Once you have made choices for the Ethernet interface and UPD port, set the value of the **snmp.standalone** option to one and restart the subagent. Once restarted, it operates as an SNMP agent independent of **snmpd** and any subagents.

For example, to run as a standalone agent listening on UDP port 9161 of the Ethernet interface with IP address 10.53.1.37, issue the commands shown below.

Configuring to run as a standalone agent:

```
% ./msconfig set snmp.port 9161
% ./msconfig set snmp.listenaddr 10.53.1.37
% ./msconfig set snmp.standalone 1
% ./stop-msg snmp
% ./start-msg snmp
% ./snmpwalk -v 1 -c public 10.53.1.37:9161 .
SNMPv2-SMI::mib-2.27.1.1.2.1 = STRING: "/opt/sun/comms/messaging64 MTA on
mail.example.com"
...
```

Monitoring Multiple Instances of Messaging Server

Two techniques for monitoring multiple instances of Messaging Server running on the same host computer are herein discussed. The first technique, running the subagent in standalone mode, is well suited to high-availability failover (HA) configurations in which the individual instances of Messaging Server may dynamically move between host computers. The second technique, the use of SNMP v3 context names, has some limited benefit in situations where multiple instances of Messaging Server are confined to a single system and it is desirable to limit the number of IP addresses polled by SNMP monitoring software (for example, when licensing of the monitoring software has a per IP address cost component). This latter technique may also be used in HA failover settings but would require polling just as many IP addresses as the standalone mode technique.

Using Standalone Agents for High-availability Failover

In a high-availability failover setting where SNMP monitoring of Messaging Server is desired, it is recommended that you run Messaging Server's SNMP subagent as a standalone agent as described in "[Running as a Standalone SNMP Agent](#)." When the subagents are run in standalone mode, each HA instance of Messaging Server should have its **snmp.listenaddr** option set to the value of that instance's failover IP address. To simplify management, each instance should use the same UDP port, but that port should be distinct from those used by the **snmpd** daemons running on each of the physical cluster hosts. Typically those daemons will be using UDP port 161 so explicitly specify a different port number with the **snmp.port** option.

When Messaging Server's SNMP support is configured as recommended here, a monitoring station can monitor each instance of Messaging Server through its failover IP address or hostname regardless of which physical cluster host the instance is running on. Moreover, you are assured that Messaging Server's standalone SNMP agents do not conflict with one another as each listens only on its own virtual Ethernet interface identified by that instance's unique failover IP address. (These virtual Ethernet interfaces are automatically created by the HA failover framework.) Owing to the careful selection of a UDP port, the agents do not conflict with the **snmpd** daemons running on systems within the cluster.

Distinguishing Multiple Instances Through SNMP v3 Context Names

While there is no downside to using Messaging Server's SNMP support in standalone mode as described in "[Running as a Standalone SNMP Agent](#)," it is recognized that some sites may prefer to use a more traditional subagent mode while still maintaining the capability of monitoring multiple instances of Messaging Server running concurrently on the same system. For instance, an SNMP monitoring system whose licensing model limits the number of IP addresses which may be polled. To achieve this goal, continue to run Messaging Server's SNMP subagent with **snmp.standalone** set to zero. Additionally, configure each instance of Messaging Server to use a distinct SNMP v3 context name by specifying a non-zero value for the **snmp.enablecontextname** option. If a context name different than the value of **base.defaultdomain** is desired, then set the desired name with the **snmp.contextname** option. Once each instance of Messaging Server's SNMP subagent is restarted, they can then be monitored with SNMP v3 queries that include the proper context names. The MIBs of two instances of Messaging Server running on the same system are distinguished by the instance's SNMP v3 context name and so no MIB object identifier (OID) conflicts will arise.

Messaging Server's Net-SNMP-based SNMP Subagent Options

The following options apply only to Messaging Server's Net-SNMP based SNMP subagent. That subagent is used on Oracle Solaris platforms running Oracle Solaris 10 and later as well as Linux platforms. The options described in this section do not apply to the legacy SNMP subagent supplied for Oracle Solaris platforms running Oracle Solaris 9 and earlier operating systems.

Table 26–1 *SNMP Subagent Options*

Option (Default)	Description
snmp.enable (0)	The Messaging Server SNMP subagent only runs when this option is given a value of 1 , in which case Messaging Server automatically stops and starts the subagent as part of its normal startup and shutdown procedures. By default this option is set to zero, which disables operation of the subagent. Before enabling the subagent, ensure that the platform's master agent has been properly configured as described in "Running as a Standalone SNMP Agent."
snmp.standalone (0)	Messaging Server's SNMP support normally runs as a SNMP subagent, receiving SNMP requests through the platform's SNMP master agent, snmpd . This operational mode is the default and is selected by giving this option a value of 0. However, as described in "Running as a Standalone SNMP Agent," the subagent may run in a "standalone" mode whereby it operates as a SNMP agent independent of snmpd . When run in standalone mode, the subagent, now an SNMP agent, listens directly for SNMP requests on the Ethernet interface and UDP port specified by, respectively, the snmp.listenaddr and snmp.port options. To run in this standalone mode, specify a value of 1 for this option. Running in standalone mode does not interfere with other SNMP master or subagents running on the system.
snmp.listenaddr (INADDR_ANY)	Hostname or IP address of the Ethernet interface to listen for SNMP requests on when running in standalone mode. By default, all available interfaces are listened on. This corresponds to specifying the value INADDR_ANY . A specific interface may be selected by specifying either the IP address or hostname associated with that interface. The interface may be either a physical interface or a virtual interface. This option is ignored when snmp.standalone is set to 0.
snmp.cachettl (30)	Time to live (TTL) in seconds for cached monitoring data. This option controls how long the subagent will report the same monitoring data before refreshing that data with new information obtained from Messaging Server. With the exception of message loop information, data is cached for no longer than 30 seconds by default. Loop information, as determined by scanning for .HELD files, is updated only once every 10 minutes. That because of the resource cost of scanning all the on-disk message queues. Note that the subagent does not continually update its monitoring data: it is only updated upon receipt of an SNMP request and the cached data has expired (that is, outlived its TTL). If the TTL is set to 30 seconds and SNMP requests are made only every five minutes, then each SNMP request causes the subagent to obtain fresh data from Messaging Server. That is, data from Messaging Server is obtained only once every five minutes. If, on the other hand, SNMP requests are made every 10 seconds, then the subagent responds to some of those requests with cached data as old as 29 seconds. Messaging Server is polled only once every 30 seconds.
snmp.servertimeout (5)	The subagent determines the operational status of each monitored service by actually opening TCP connections to each service and undergoing a protocol exchange. This timeout value, measured in seconds, controls how long the subagent waits for a response to each step in the protocol exchange. By default, a timeout value of five seconds is used.

Table 26–1 (Cont.) SNMP Subagent Options

Option (Default)	Description
<code>snmp.directoryscan</code> (1)	Use this option to control whether or not the subagent performs scans of the on-disk message queues for .HELD message files and the oldest message files. That information corresponds to the mtaGroupLoopsDetected , mtaGroupOldestMessageStored , and mtaGroupOldestMessageId MIB variables. When this option has the value 1 , then a cache of this information is maintained and updated as needed. Sites with thousands of queued messages, that are not interested in these particular MIB variables should consider setting this option's value to 0 .
<code>snmp.enablecontextname</code> (0)	The subagent has the ability to register its MIBs under an SNMP v3 <i>context name</i> . When this is done, the MIBs may only be requested by a SNMP v3 client that specifies the context name in its SNMP request. Use of context names allows multiple, independent subagents to register Network Services and MTA MIBs under the same OID tree (that is, under the same SNMP master agent). See " Monitoring Multiple Instances of Messaging Server " for further information. To enable the use of SNMP v3 context names, specify a value of 1 for this option. When that is done, the subagent defaults to using the value of the base.defaultdomain option for its context name. To use a different value for the context name, use the snmp.contextname option.
<code>snmp.contextname</code> (base.defaultdomain)	When the use of SNMP v3 context names has been enabled with snmp.enablecontextname , this option can be used to explicitly set the context name used by the subagent for its MIBs. The values supplied for this option are string values and must be appropriate for use as a SNMP v3 context name. This option is ignored when snmp.enablecontextname has the value 0 .

Monitoring from an SNMP Client

The base OIDs for RFC 2788 (<http://www.faqs.org/rfcs/rfc2788.html>) and RFC 2789 (<http://www.faqs.org/rfcs/rfc2789.html>) are:

- `mib-2.27` = 1.3.6.1.2.1.27
- `mib-2.28` = 1.3.6.1.2.1.28

Point your SNMP client at those two OIDs and access as the "public" SNMP community.

To load copies of the MIBs into your SNMP client, ASCII copies of the MIBs are located in the *MessagingServer_home/lib/config-templates* directory under the file names **rfc2788.mib** and **rfc2789.mib**. For directions on loading those MIBs into your SNMP client software, consult the SNMP client software documentation. The **SnmpAdminString** data type used in those MIBs may not be recognized by some older SNMP clients. In that case, use the equivalent files **rfc2248.mib** and **rfc2249.mib** also found in the same directory.

SNMP Information from the Messaging Server

This section summarizes the Messaging Server information provided via SNMP.

Topics in this section:

- [applTable](#)
- [assocTable](#)
- [mtaTable](#)
- [mtaGroupTable](#)
- [mtaGroupAssociationTable](#)
- [mtaGroupErrorTable](#)

For detailed information refer to the individual MIB tables in RFC 2788 (<http://www.faqs.org/rfcs/rfc2788.html>) and RFC 2789 (<http://www.faqs.org/rfcs/rfc2789.html>). Note that the RFC/MIB terminology refers to the messaging services (MTA, HTTP, and so on) as *applications* (**appl**), Messaging Server network connections as *associations* (**assoc**), and MTA channels as *MTAgroups* (**mtaGroups**).

On platforms where more than one instance of Messaging Server may be concurrently monitored, there may then be multiple sets of MTAs and servers in the **applTable**, and multiple MTAs in the other tables.

Note: The cumulative values reported in the MIBs (for example, total messages delivered, total IMAP connections, and so on) are reset to zero after a reboot.

Each site has different thresholds and significant monitoring values. A good SNMP client allows you to do trend analysis and then send alerts when sudden deviations from historical trends occur.

applTable

The **applTable** provides server information. It is a one-dimensional table with one row for the MTA and an additional row for each of the following servers, if enabled: WebMail HTTP, IMAP, POP, SMTP, and SMTP Submit. This table provides version information, uptime, current operational status (up, down, congested), number of current connections, total accumulated connections, and other related data.

Here is an example of data from **applTable** (mib-2.27.1.1).

applTable:

```
applName.1 = mailsrv-1 MTA on mailsrv-1.west.example.org (1)
applVersion.1 = 5.1
applUptime.1 = 7322 (2)
applOperStatus.1 = up (3)
applLastChange.1 = 7422 (2)
applInboundAssociations.1 = (5)
applOutboundAssociations.1 = (2)
applAccumulatedInboundAssociations.1 = 873
applAccumulatedOutboundAssociations.1 = 234
applLastInboundActivity.1 = 1054822 (2)
applLastOutboundActivity.1 = 1054222 (2)
applRejectedInboundAssociations.1 = 0 (4)
applFailedOutboundAssociations.1 = 17
applDescription.1 = Sun Java System Messaging Server 6.1
applName.2 1 = mailsrv-1 HTTP WebMail svr. mailsrv-1.example.org (1)
...
applName.3 = mailsrv-1 IMAP server on mailsrv-1.west.example.org
...
applName.4 = mailsrv-1 POP server on mailsrv-1.west.example.org
...
applName.5 = mailsrv-1 SMTP server on mailsrv-1.west.example.org
...
applName.6 = mailsrv-1 SMTP Submit server on mailsrv-1.west.example.org
...
```

Notes:

1. The application (**.appl***) suffixes (**.1**, **.2**, and so on) are the row numbers, **applIndex**. **applIndex** has the value 1 for the MTA, value 2 for the HTTP server, and so on. Thus, in this example, the first row of the table provides data on the MTA, the second on the POP server, and so on. The name after the equal sign is the name of the Messaging Server instance being monitored. In this example, the instance name is **mailsrv-1**.
2. These are SNMP TimeStamp values and are the value of **sysUpTime** at the time of the event. **sysUpTime**, in turn, is the count of hundredths of seconds since the SNMP master agent was started.
3. The operational status of the HTTP, IMAP, POP, SMTP, and SMTP Submit servers is determined by actually connecting to them by their configured TCP ports and performing a simple operation using the appropriate protocol (for example, a HEAD request and response for HTTP, a **HELO** command and response for SMTP, and so on). From this connection attempt, the status-**up** (1), **down** (2), or **congested** (4)-of each server is determined. Note that these probes appear as normal inbound connections to the servers and contribute to the value of the **applAccumulatedInboundAssociations** MIB variable for each server. For the MTA, the operational status is taken to be that of the Job Controller. If the MTA is shown to be up, then the Job Controller is up. If the MTA is shown to be down, then the Job Controller is down. This MTA operational status is independent of the status of the MTA's Service Dispatcher. The operational status for the MTA only takes on the value of up or down. Although the Job Controller does have a concept of "congested," it is not indicated in the MTA status.
4. For the HTTP, IMAP, and POP servers the **applRejectedInboundAssociations** MIB variable indicates the number of failed login attempts and not the number of rejected inbound connection attempts.

applTable Usage

Monitoring server status (**applOperStatus**) for each of the listed applications is key to monitoring each server.

If it has been a long time since the MTA last inbound activity as indicated by **applLastInboundActivity**, then something may be broken preventing connections. If **applOperStatus=2** (down), then the monitored service is down. If **applOperStatus=1** (up), then the problem may be elsewhere.

assocTable

This table provides network connection information to the MTA. It is a two-dimensional table providing information about each active network connection. Connection information is not provided for other servers. Here is an example of data from **applTable** (mib-2.27.2.1).

assocTable:

```
assocRemoteApplication.1.1 = 129.146.198.167 (1)
assocApplicationProtocol.1.1 = applTCPProtoID.25 (2)
assocApplicationType.1.1 = peerinitiator(3) (3)
assocDuration.1.1 = 400 (4)
...
```

Notes: In the **.x.y** suffix (1.1), **x** is the application index, **applIndex**, and indicates which application in the **applTable** is being reported on. In this case, the MTA. The **y** serves to enumerate each of the connections for the application being reported on.

1. The source IP address of the remote SMTP client.

2. This is an OID indicating the protocol being used over the network connection. **aplTCPProtoID** indicates the TCP protocol. The **.n** suffix indicates the TCP port in use and **.25** indicates SMTP which is the protocol spoken over TCP port 25.
3. It is not possible to know if the remote SMTP client is a user agent (UA) or another MTA. As such, the subagent always reports **peer-initiator**; **ua-initiator** is never reported.
4. This is an SNMP **TimeInterval** and has units of hundredths of seconds. In this example, the connection has been open for 4 seconds.

assocTable Usage

This table is used to diagnose active problems. For example, if you suddenly have 200,000 inbound connections, this table can let you know where they are coming from.

mtaTable

This is a one-dimensional table with one row for each MTA in the **applTable**. Each row gives totals across all channels (referred to as groups) in that MTA for select variables from the **mtaGroupTable**. Here is an example of data from **applTable** (mib-2.28.1.1).

mtaTable:

```
mtaReceivedMessages.1 = 172778
mtaStoredMessages.1 = 19
mtaTransmittedMessages.1 = 172815
mtaReceivedVolume.1 = 3817744
mtaStoredVolume.1 = 34
mtaTransmittedVolume.1 = 3791155
mtaReceivedRecipients.1 = 190055
mtaStoredRecipients.1 = 21
mtaTransmittedRecipients.1 = 3791134
mtaSuccessfulConvertedMessages.1 = 0 (1)
mtaFailedConvertedMessages.1 = 0
mtaLoopsDetected.1 = 0 (2)
```

Notes: The **.x** suffix (**.1**) provides the row number for this application in the **applTable**. In this example, **.1** indicates this data is for the first application in the **applTable**. Thus, this is data on the MTA.

1. Only takes on non-zero values for the conversion channel.
2. Counts the number of **.HELD** message files currently stored in the MTA's message queues.

mtaTable Usage

If **mtaLoopsDetected** is not zero, then there is a looping mail problem. Locate and diagnose the **.HELD** files in the MTA queue to resolve the problem.

If the system does virus scanning with a conversion channel and rejects infected messages, then **mtaSuccessfulConvertedMessages** gives a count of infected messages in addition to other conversion failures.

mtaGroupTable

This two-dimensional table provides channel information for each MTA in the **applTable**. This information includes such data as counts of stored (that is, queued) and delivered mail messages. Monitoring the count of stored messages,

mtaGroupStoredMessages, for each channel is critical. When the value becomes abnormally large, mail is backing up in your queues.

Here is an example of data from **mtaGroupTable** (mib-2.28.2.1).

mtaGroupTable:

```
mtaGroupName.1.1 = tcp_intranet 1
...
mtaGroupName.1.2 = ims-ms
...
mtaGroupName.1.3 = tcp_local
mtaGroupDescription.1.3 = mailsrv-1 MTA tcp_local channel
mtaGroupReceivedMessages.1.3 = 12154
mtaGroupRejectedMessages.1.3 = 0
mtaGroupStoredMessages.1.3 = 2
mtaGroupTransmittedMessages.1.3 = 12148
mtaGroupReceivedVolume.1.3 = 622135
mtaGroupStoredVolume.1.3 = 7
mtaGroupTransmittedVolume.1.3 = 619853
mtaGroupReceivedRecipients.1.3 = 33087
mtaGroupStoredRecipients.1.3 = 2
mtaGroupTransmittedRecipients.1.3 = 32817
mtaGroupOldestMessageStored.1.3 = 1103
mtaGroupInboundAssociations.1.3 = 5
mtaGroupOutboundAssociations.1.3 = 2
mtaGroupAccumulatedInboundAssociations.1.3 = 150262
mtaGroupAccumulatedOutboundAssociations.1.3 = 10970
mtaGroupLastInboundActivity.1.3 = 1054822
mtaGroupLastOutboundActivity.1.3 = 1054222
mtaGroupRejectedInboundAssociations.1.3 = 0
mtaGroupFailedOutboundAssociations.1.3 = 0
mtaGroupInboundRejectionReason.1.3 =
mtaGroupOutboundConnectFailureReason.1.3 =
mtaGroupScheduledRetry.1.3 = 0
mtaGroupMailProtocol.1.3 = applTCPProtoID.25
mtaGroupSuccessfulConvertedMessages.1.3 = 03 2
mtaGroupFailedConvertedMessages.1.3 = 0
mtaGroupCreationTime.1.3 = 0
mtaGroupHierarchy.1.3 = 0
mtaGroupOldestMessageId.1.3 = <01IFBV8AT8HYB4T6UA@red.ipplanet.com>
mtaGroupLoopsDetected.1.3 = 0 3
mtaGroupLastOutboundAssociationAttempt.1.3 = 1054222
```

Notes: In the **.x.y** suffix (example: **1.1**, **1.2**, **1.3**), **x** is the application index, **applIndex**, and indicates which application in the **applTable** is being reported on. In this case, the MTA. The **y** serves to enumerate each of the channels in the MTA. This enumeration index, **mtaGroupIndex**, is also used in the **mtaGroupAssociationTable** and **mtaGroupErrorTable** tables.

1. The name of the channel being reported on. In this case, the **tcp_intranet** channel.
2. Only takes on non-zero values for the conversion channel.
3. Counts the number of **.HELD** message files currently stored in this channel's message queue.

mtaGroupTable Usage

Trend analysis on **Rejected** and **Failed** might be useful in determining potential channel problems.

A sudden jump in the ratio of **mtaGroupStoredVolume** to **mtaGroupStoredMessages** could mean that a large junk mail is bouncing around the queues.

A large jump in **mtaGroupStoredMessages** could indicate unsolicited bulk email is being sent or that delivery is failing for some reason.

If the value of **mtaGroupOldestMessageStored** is greater than the value used for the undeliverable message notification times (notices channel option) this may indicate a message which cannot be processed even by bounce processing. Note that bounces are done nightly so you want to use **mtaGroupOldestMessageStored** > (maximum age + 24 hours) as the test.

If **mtaGroupLoopsDetected** is greater than 0, a mail loop has been detected.

mtaGroupAssociationTable

This is a three-dimensional table whose entries are indices into the **assocTable**. For each MTA in the **applTable**, there is a two-dimensional sub-table. This two-dimensional sub-table has a row for each channel in the corresponding MTA. For each channel, there is an entry for each active network connection which that channel has currently underway. The value of the entry is the index into the **assocTable** (as indexed by the entry's value and the **applIndex** index of the MTA being looked at). This indicated entry in the **assocTable** is a network connection held by the channel.

In simple terms, the **mtaGroupAssociationTable** table correlates the network connections shown in the **assocTable** with the responsible channels in the **mtaGroupTable**.

Here is an example of data from **mtaGroupAssociationTable** (mib-2.28.3.1).

mtaGroupAssociationTable:

```
mtaGroupAssociationIndex.1.3.1 = 1 1
mtaGroupAssociationIndex.1.3.2 = 2
mtaGroupAssociationIndex.1.3.3 = 3
mtaGroupAssociationIndex.1.3.4 = 4
mtaGroupAssociationIndex.1.3.5 = 5
mtaGroupAssociationIndex.1.3.6 = 6
mtaGroupAssociationIndex.1.3.7 = 7
```

Notes: In the **.x.y.z** suffix, **x** is the application index, **applIndex**, and indicates which application in the **applTable** is being reported on. In this case, the MTA. The **y** indicates which channel of the **mtaGroupTable** is being reported on. In this example, 3 indicates the **tcp_local channel**. The **z** serves to enumerate the associations open to or from the channel.

- The value here is an index into the **assocTable**. Specifically, **x** and this value become, respectively, the values of the **applIndex** and **assocIndex** indices into the **assocTable**. Or, put differently, this is saying that (ignoring the **applIndex**) the first row of the **assocTable** describes a network connection controlled by the **tcp_local** channel.

mtaGroupErrorTable

This is another three-dimensional table which gives the counts of temporary and permanent errors encountered by each channel of each MTA while attempting delivery of messages. Entries with index values of 4000000 are temporary errors while those with indices of 5000000 are permanent errors. Temporary errors result in the message being re-queued for later delivery attempts. Permanent errors result in either the message being rejected or otherwise returned as undeliverable.

Here is an example of data from **mtaGroupErrorTable** (mib-2.28.5.1).

mtaGroupErrorTable:

```
mtaGroupInboundErrorCount.1.1.4000000 1 = 0
mtaGroupInboundErrorCount.1.1.5000000 = 0
mtaGroupInternalErrorCount.1.1.4000000 = 0
mtaGroupInternalErrorCount.1.1.5000000 = 0
mtaGroupOutboundErrorCount.1.1.4000000 = 0
mtaGroupOutboundErrorCount.1.1.5000000 = 0
mtaGroupInboundErrorCount.1.2.4000000 1 = 0
...
mtaGroupInboundErrorCount.1.3.4000000 1 = 0
...
```

Notes:

- In the **.x.y.z** suffix, **x** is the application index, **applIndex**, and indicates which application in the **applTable** is being reported on. In this case, the MTA. The **y** indicates which channel of the **mtaGroupTable** is being reported on. In this example, 1 specifies the **tcp_intranet** channel, 2 the **ims-ms** channel, and 3 the **tcp_local** channel. Finally, the **z** is either 4000000 or 5000000 and indicates, respectively, counts of temporary and permanent errors encountered while attempting message deliveries for that channel.

mtaGroupErrorTable Usage

A large jump in error count may likely indicate an abnormal delivery problem. For instance, a large jump for a **tcp_** channel may indicate a DNS or network problem. A large jump for the **ims-ms** channel may indicate a delivery problem to the message store (for example, a partition is full, **stored** problem, and so on).

Short Message Service (SMS)

This information describes how to implement the Short Message Service (SMS) in Unified Configuration on the Oracle Communications Messaging Server.

Topics:


- Introduction
- SMS Channel Theory of Operation
- SMS Channel Configuration
- SMS Gateway Server Theory of Operation
- SMS Gateway Server Configuration
- SMS Gateway Server Storage Requirements
- SMS Configuration Examples

Introduction

Messaging Server implements email-to-mobile and mobile-to-email messaging using a Short Message Service (SMS). SMS can be configured to be either one-way (email-to-mobile only) or two-way (both email-to-mobile and mobile-to-email). To enable one-way service only, you must add and configure the SMS channel. To enable two-way service, you must add and configure the SMS channel, and in addition, configure the SMS Gateway Server.

For both one- and two-way SMS, the generated SMS messages are submitted to a Short Message Service Center (SMSC) using the Short Message Peer to Peer (SMPP) protocol. Specifically, the SMSC must provide a V3.4 or later SMPP server that supports TCP/IP.

[Figure 27-1](#) illustrates the logical flow of messages for both one-way and two-way SMS.

Figure 27–1 Logical Flow For One-Way and Two-Way SMS

One-Way SMS

To enable one-way service, the Messaging Server implements an SMPP client (the MTA SMS channel) that communicates with remote SMSCs. The SMS channel converts enqueued email messages to SMS messages as described in "The Email to SMS Conversion Process" of multipart MIME messages as well as character set translation issues.

Operating in this capacity, the SMS channel functions as an (SMPP) External Short Message Entity (ESME).

Two-Way SMS

Two-Way SMS enables the mail server not only to send email to remote devices, but allows for receiving replies from the remote devices and for remote device email origination.

Enabling two-way SMS service requires both the MTA SMS channel (SMPP client), as explained in the previous topic, and the SMS Gateway Server. Messaging Server installs an SMS Gateway Server as part of its general installation process, which you must then configure. The SMS Gateway Server performs two functions:

- **SMPP relay** The SMS Gateway Server acts as a transparent SMPP client between the MTA SMS channel and SMSCs. However, in addition, while acting as a relay, the SMS Gateway Server generates unique SMS source addresses for relayed messages, and saves the message IDs returned by the remote SMSCs for later correlation with SMS notification messages.
- **SMPP server** The SMS Gateway Server acts as an SMPP server to receive mobile originated SMS messages, replies to prior email messages, and SMS notifications. The SMS Gateway Server extracts destination email addresses from the SMS messages using profiles that define the conversion process. Profiles also describe how to handle notification messages returned by remote SMSCs in response to previously sent email-to-mobile messages.

Note: Messaging Server does not support the two-way SMS on the Windows platform.

Requirements

This manual assumes that you have read Logica CMG's SMPP specification, and the SMPP documentation for your SMSC.

In order to implement SMS, the SMSC must support SMPP V3.4, or later, over TCP/IP and there must be TCP/IP connectivity between the host running Messaging Server and the SMSC.

For storage planning information for the SMS Gateway Server, see "[SMS Gateway Server Storage Requirements](#)."

SMS Channel Theory of Operation

The SMS channel is a multi-threaded channel which converts queued email messages to SMS messages and then submits them for delivery to an SMSC.

The following channel operation topics are covered in this section:

- [Directing Email to the Channel](#)
- [The Email to SMS Conversion Process](#)
- [The SMS Message Submission Process](#)
- [Site-defined Address Validity Checks and Translations](#)
- [Site-defined Text Conversions](#)

Directing Email to the Channel

When the SMS channel is configured as per "[SMS Channel Configuration](#)" purposes of discussion, let us assume that the host name **sms.example.org** is a host name associated with the channel. In that case, email is directed to the channel with an address of the form:

local-part@sms.example.org

in which **local-part** is either the SMS destination address (for example, a wireless phone number, pager ID, etc.) or an attribute-value pair list in the format:

/attribute1=value1/attribute2=value2/.../@sms.example.org

The recognized attribute names and their usages are given in [Table 27-1](#). These attributes allow for per-recipient control over some channel options.

Table 27-1 SMS Attributes

Attribute	Attribute Value and Usage
ID	SMS destination address (for example, wireless phone number, pager ID, etc.) to direct the SMS message to. This attribute and associated value must be present.
FROM	SMS source address. Ignored when option USE_HEADER_FROM=0 .
FROM_NPI	Use the specified NPI value. Ignored when option USE_HEADER_FROM=0 .

Table 27-1 (Cont.) SMS Attributes

Attribute	Attribute Value and Usage
FROM_TON	Use the specified TON value. Ignored when option <code>USE_HEADER_FROM=0</code> .
MAXLEN	The maximum, total bytes (that is, eight bit bytes) to place into the generated SMS message or messages for this recipient. The lower value of either <code>MAXLEN</code> and the value specified by the " <code>MAX_MESSAGE_SIZE</code> " channel option is used.
MAXPAGES	The maximum number of SMS messages to split the email message into for this recipient. The lower value of either <code>MAXPAGES</code> and the value specified by the " <code>MAX_PAGES_PER_MESSAGE</code> " channel option is used.
NPI	Specify a Numeric Plan Indicator (NPI) value for the destination SMS address specified with the <code>ID</code> attribute. See the description of the " <code>DEFAULT_DESTINATION_NPI</code> " channel option for information on the accepted values for this attribute. When this attribute is used, its value overrides the value given by the <code>DEFAULT_DESTINATION_NPI</code> channel option.
PAGELEN	Maximum number of bytes to place into a single SMS message for this recipient. The minimum of this value and that specified with the " <code>MAX_PAGE_SIZE</code> " channel option is used.
TO	Synonym for <code>ID</code> .
TO_NPI	Synonym for <code>NPI</code> .
TO_TON	Synonym for <code>TON</code> .
TON	Specify a Type of Number (TON) value for the destination SMS address given with the <code>ID</code> attribute. See the description of the " <code>DEFAULT_DESTINATION_TON</code> " channel option for information on the accepted values for this attribute. When this attribute is used, its value overrides the value given by the <code>DEFAULT_DESTINATION_TON</code> channel option.

Some example addresses:

```
123456@sms.example.org
/id=123456/@sms.example.org
/id=123456/maxlen=100/@sms.example.org
/id=123456/maxpages=1/@sms.example.org
```

For information on performing translations, validity checks, and other operations on the SMS destination address portion of the email address, see "[Site-defined Address Validity Checks and Translations](#)."


The Email to SMS Conversion Process

In order for email to be sent to a remote site, email must be converted to SMS messages that can be understood by the remote SMSCs. This section describes the process of converting an email message queued to the SMS channel to one or more SMS messages. As described below, options allow control over the maximum number of SMS messages generated, the maximum total length of those SMS messages, and the maximum size of any single SMS message. Only text parts (that is, MIME text content types) from the email message are used and the maximum number of parts converted may also be controlled.

Character sets used in the email message's header lines and text parts are all converted to Unicode and then converted to an appropriate SMS character set.

When there is no **SMS_TEXT** mapping table (see "Site-defined Text Conversions") an email message queued to the SMS channel receives the processing illustrated in Figure 27-2.

Figure 27-2 Channel Email Processing



5. For each text part, while the **MAX_MESSAGE_PARTS** limit has not been reached, the text part is decoded to Unicode and appended to the buffer, along with the value of the **LINE_STOP** channel option.
6. The resulting output buffer is then converted from Unicode to either the SMSC's default character set or UCS2 (UTF-16). The SMSC's default character set is specified with the "**SMSC_DEFAULT_CHARSET**" option.
7. After being converted, it is then truncated to not exceed "**MAX_MESSAGE_SIZE**" bytes.
8. The converted string from "[The Email to SMS Conversion Process](#)" is then broken into one or more SMS messages, no single SMS message longer than **MAX_PAGE_SIZE** bytes. At most, "**MAX_PAGES_PER_MESSAGE**" SMS messages will be generated.

Note: As an email message may have multiple recipients, Step 6 through Step 8 may need to be done for each recipient address which makes use of the **MAXLEN**, **MAXPAGES**, or **PAGELEN** attributes described in "[Directing Email to the Channel](#)."

Sample Email Message Processing

For example, with the channel's default settings, the email message:

```
From: John Doe
To: 1234567@sms.example.org
Subject: Today's meeting
Date: Fri, 26 March 2001 08:17
```

The staff meeting is at 14:30 today in the big conference room.

Would be converted to the SMS message:

jdoe@example.org (Today's meeting) The staff meeting is at 14:30 today in the big conference room.

A different set of option settings, that follows:

```
CONTENT_PREFIX=Msg:
FROM_FORMAT=From:{$pa}
SUBJECT_FORMAT=Subj:{$s}
```

would instead produce:

From:John Doe Subj:Today's meeting Msg:The staff meeting is at 14:30 today in the big conference room.

The SMS Message Submission Process

Once an email message has been converted to one or more SMS messages, with possibly different sets for each recipient, the SMS messages are then submitted to the destination SMSC. The submissions are effected using SMPP V3.4 over TCP/IP. The hostname (**SMPP_SERVER**) of the SMPP server is taken to be the official host name associated with the SMS channel; the TCP port (**SMPP_PORT**) to use is specified with the **port** channel option.

When there are messages to process, the channel is started. The channel binds to the SMPP server as a transmitter, presenting the credentials specified with the **ESME_** channel options described in "[SMPP Options](#)." [Table 27-2](#) lists the fields set in a

BIND_TRANSMITTER PDU (Protocol Data Unit), and gives their values:

Table 27–2 Fields in Generated in a BIND_TRANSMITTER PDU

Field	Channel Options	Value
system_id	ESME_SYSTEM_ID	Default value is an empty string.
password	ESME_PASSWORD	Default value is an empty string.
system_type	ESME_SYSTEM_TYPE	Default value is an empty string.
interface_version	n/a	0x34 indicating SMPP V3.4.
addr_ton	ESME_ADDRESS_TON	Default value is 0x00 indicating an unknown TON.
addr_npi	ESME_ADDRESS_NPI	Default value is 0x00 indicating an unknown NPI.
addr_range	ESME_IP_ADDRESS	Default value is an empty string.

Note that the channel is multithreaded. Depending on how much mail there is to send, the channel may have multiple dequeue threads running. (There can even be multiple channel processes running.) Each thread does a **BIND_TRANSMITTER** and then on that TCP/IP connection, sends all of the SMS messages it has to send, and then sends an **UNBIND**, and then closes the connection. No attempt is made to hold a connection open for a period of idle time for potential reuse. If the remote SMPP server sends back a throttle error, then an **UNBIND** is issued, the TCP/IP connection is closed, and a new connection and **BIND** established. It behaves similarly if the remote SMPP server sends an **UNBIND** before it is finished sending its SMS messages.

The SMS messages are then submitted using SMPP **SUBMIT_SM** PDUs. If a permanent error is returned (for example, **ESME_RINVDSTADR**), then the email message is returned as undeliverable. If a temporary error is returned, then the email message is re-enqueued for a later delivery attempt. To clarify, a permanent error is one for which the condition is likely to exist indefinitely and for which repeated delivery attempts will have no positive effect, such as invalid SMS destination addresses. Whereas, a temporary error is one for which the condition is likely to not exist in the near future, such as a server down or server congested condition.

If the **USE_HEADER_FROM** option has the value 1, then the source address for the submitted SMS message is set. The value used is derived from the originating email message and is chosen to be the most likely (email) address to which any replies should be directed. Accordingly, the source address taken from one of the following seven sources, shown in decreasing order of preference:

1. Resent-reply-to:
2. Resent-from:
3. Reply-to:
4. From:
5. Resent-sender:
6. Sender:
7. Envelope From:

Note that the **Resent-reply-to:** and **Reply-to:** header lines are only considered if the "**USE_HEADER_REPLY_TO**" option has the value 1. The default value is the value 0. As such, only items 4, 6, and 7 are considered by the default configuration. Finally, since the source address in an SMS message is limited to 20 bytes, the source address chosen will be truncated if it exceeds that limit.

Table 27–3 shows the mandatory fields set in a **SUBMIT_SM** PDU:

Table 27–3 Mandatory Fields in Generated SUBMIT_SM PDUs

Field	Value
service_type	"DEFAULT_SERVICE_TYPE" channel option; default value is an empty string.
source_addr_ton	"DEFAULT_SOURCE_TON" channel option; if USE_HEADER_FROM=1 , then this field is usually forced to the value 0x05 indicating an alphanumeric TON; otherwise, the default value is 0x01 indicating an international TON.
source_addr_npi	"DEFAULT_SOURCE_NPI" channel option; default value is 0x00.
source_addr	"DEFAULT_SOURCE_ADDRESS" channel option if USE_HEADER_FROM=0 ; otherwise, an alphanumeric string representing the originator of the email message.
dest_addr_ton	TON addressing attribute or "DEFAULT_DESTINATION_TON" channel option; default value is 0x01 indicating an international TON.
dest_addr_npi	NPI addressing attribute or "DEFAULT_SOURCE_NPI" channel option; default value is 0x00 indicating an unknown NPI.
dest_addr	Destination SMS address derived from the local part of the email envelope To: address; see " Directing Email to the Channel ."
esm_class	For one-way SMS, set to 0x03, indicating store and forward mode, default SMSC message type, and do not set reply path. For a two-way MSM message, set to 0x83.
protocol_id	0x00; unused for CDMA and TDMA; for GSM, 0x00 indicates no Internet, but SME-to-SME protocol.
priority_flag	0x00 for GSM & CDMA and 0x01 for TDMA, all indicating normal priority; See the description of the " DEFAULT_PRIORITY " channel option.
schedule_delivery_time	Empty string indicating immediate delivery.
validity_period	"DEFAULT_VALIDITY_PERIOD" channel option; default value is an empty string indicating that the SMSC's default should be used.
registered_delivery	0x00 indicating no registered delivery.
replace_if_present_flag	0x00 indicating that any previous SMS messages should not be replaced.
data_coding	0x00 for the SMSC's default character set; 0x08 for the UCS2 character set.
sm_default_msg_id	0x00 indicating not to use a pre-defined message.
sm_length	Length and content of the SMS message; see " The Email to SMS Conversion Process ."
short_message	Length and content of the SMS message; see " The Email to SMS Conversion Process ."

Table 27–4 shows the optional fields in a **SUBMIT_SM** PDU:

Table 27-4 Optional Fields in Generated *SUBMIT_SM* PDUs.

Field	Value
privacy	See the description of the "DEFAULT_PRIVACY" channel option; default is to not provide this field unless the email message has a Sensitivity: header line
sar_refnum	See the description of the "USE_SAR" channel option; default is to not provide these fields
sar_total	See sar_refnum above.
sar_seqnum	See sar_refnum above.

The channel remains bound to the SMPP server until either it has no more SMS messages to submit (the message queue is empty), or "**MAX_PAGES_PER_BIND**" has been exceeded. In the latter case, a new connection is made and bind operation performed if there remain further SMS messages to send.

Note that the SMS channel is multithreaded. Each processing thread in the channel maintains its own TCP connection with the SMPP server. For example, if there are three processing threads each with SMS messages to submit, then the channel will have three open TCP connections to the SMPP server. Each connection will bind to the SMPP server as a transmitter. Moreover, any given processing thread will only have one outstanding SMS submission at a time. That is, a given thread will submit an SMS message, then wait for the submission response (that is, **SUBMIT_SM_RESP** PDU) before submitting another SMS message.

Site-defined Address Validity Checks and Translations

Sites may wish to apply validity checks or translations to SMS destination addresses encoded in the recipient email addresses described in "[Directing Email to the Channel](#)."

- Strip non-numeric characters (for example, translating 800.555.1212 to 8005551212)
- Prepend a prefix (for example, translating 8005551212 to +18005551212)
- Validate for correctness (for example, 123 is too short)

The first two tasks can be done specifically with the "**DESTINATION_ADDRESS_NUMERIC**" and "**DESTINATION_ADDRESS_PREFIX**" channel options. In general, all three of these tasks, and others can be implemented using mapping tables: either mapping table callouts in the rewrite rules or by means of a **FORWARD** mapping table. Using a mapping table callout in the rewrite rules will afford the most flexibility, including the ability to reject the address with a site-defined error response. The remainder of this section will focus on just such an approach - using a mapping table callout from the rewrite rules.

Let us suppose that destination addresses need to be numeric only, be 10 or 11 digits long, and be prefixed with the string "+1". This can be accomplished with the following rewrite rules:

```
sms.example.org ${X-REWRITE-SMS-ADDRESS,$U}@sms.example.org
sms.example.org $?Invalid SMS address
```

The first rewrite rule above calls out to the site-define mapping table named **X-REWRITE-SMS-ADDRESS**. That mapping table is passed the local part of the email address for inspection. If the mapping process decides that the local part is acceptable, then the address is accepted and rewritten to the SMS channel. If the mapping process does not accept the local part, then the next rewrite rule is applied.

Since it is a \$? rewrite rule, the address is rejected with the error text "Invalid SMS address".

The **X-REWRITE-SMS-ADDRESS** mapping table is shown below. It performs the necessary validation steps for local parts in either attribute-value pair list format or just a raw SMS destination address.

```
X-VALIDATE-SMS-ADDRESS
```

```
! Iteratively strip any non-numeric characters
$_*[$ -/:-~]* $0$2$R
! Accept the address if it is of the form lnnnnnnnnnn or nnnnnnnnnn
! In accepting it, ensure that we output +lnnnnnnnnn
1%%%%%%%% +1$0$1$2$3$4$5$6$7$8$9$Y
%%%%%%%% +1$0$1$2$3$4$5$6$7$8$9$Y
! We didn't accept it and consequently it's invalid
* $N
```

```
X-REWRITE-SMS-ADDRESS
```

```
*/id=$_*/ * $C$0/id=$|X-VALIDATE-SMS-ADDRESS;$1|/$2$Y$E
*/id=$_*/ * $N
* $C$|X-VALIDATE-SMS-ADDRESS;$0|$Y$E
* $N
```

With the above set up, be sure that "DESTINATION_ADDRESS_NUMERIC" option has the value 0 (the default). Otherwise, the "+" will be stripped from the SMS destination address.

Site-defined Text Conversions

Sites may customize Steps 1 - 6 described in "The Email to SMS Conversion Process" using a mapping table.

The name of the mapping table should be *SMS_Channel_TEXT* where *SMS_Channel* is the name of the SMS channel; for example, **SMS_TEXT** if the channel is named **sms** or **SMS_MWAY_TEXT** if the channel is named **sms_mway**.

Two types of entries may be made in this mapping table. However, before explaining the format of those entries, let it be made clear that an understanding of how to use mappings is essential in order to understand how to construct and use these entries. An example mapping table is given after the description of these two types of entries.

Now, the two types of entries are:

- [Message Header Entries](#)
- [Message Body Entries](#)

Message Header Entries

These entries specify which message header lines should be included in an SMS message and how they should be abbreviated or otherwise converted. Only if a header line is successfully mapped to a string of non-zero length by one of these entries will it be included in the SMS message being generated. Each entry has the format

H | *patternreplacement-text*

If a message header line matches the pattern then it will be replaced with the replacement text *replacement-text* using the mapping's pattern matching and string substitution facilities. The final result of mapping the header line will then be included in the SMS message provided that the metacharacter \$Y was specified in the replacement text. If a header line does not match any pattern string, if it maps to a

string of length zero, or if the \$Y metacharacter is not specified in the replacement text, then the header line will be omitted from the SMS message. The two entries:

```
H|From:* F:$0$Y
H|Subject:* S:$0$Y
```

cause the **From:** and **Subject:** header lines to be included in SMS messages with **From:** and **Subject:** abbreviated as **F:** and **S:**. The entries:

```
H|Date:* H|D:$0$R$Y
H|D:*,*%19%*:*:* H|D:$0$ $5:$6$R$Y
```

cause the **Date:** header line to be accepted and mapped such that, for instance, the header line

Date: Wed, 16 Dec 1992 16:13:27 -0700 (PDT)

will be converted to

D: Wed 16:13

Very complicated, iterative mappings may be built. Sites wishing to set up custom filters will first need to understand how mappings work. The **H|** in the right-hand-side of the entry may be omitted, if desired. The **H|** is allowed in that side so as to cut down on the number of table entries required by sets of iterative mappings.

Message Body Entries

Body mappings are not supported.

Example SMS Mapping Table

An example **SMS_TEXT** mapping table is shown in "[Example SMS_TEXT Mapping Table](#)." The numbers inside parentheses at the end of each line correspond to the item numbers in the section titled "[Explanatory Text](#)" that follows this table.

Example SMS_TEXT Mapping Table

SMS_TEXT

```
H|From:* H|F:$0$R$Y (1)
H|Subject:* H|S:$0$R$Y (1)
H|F:*&lt;*>* H|F:$1$R$Y (1)
H|F:*(*)* H|F:$0$2$R$Y (2)
H|F:**** H|F:$0$2$R$Y (3)
H|F:*@* H|F:$0$R$Y (4)
H|%:$ * H|$0:$1$R$Y (5)
H|%:*$ H|$0:$1$R$Y (5)
H|%:*$ $ * H|$0:$1$ $2$R$Y (6)
B|*--* B|$0-$1$R (7)
B|*..* B|$0.$1$R (7)
B|*!* B|$0!$1$R (7)
B|*??* B|$0?$1$R (7)
B|*$ $ * B|$0$ $1$R (6)
B|$ * B|$0$R (5)
B|*$ B|$0$R (5)
```

Explanatory Text

The entries in the example **SMS_TEXT** mapping table above are explained below:

In the example above, the metacharacter \$R is used to implement and control iterative application of the mappings. By iterating on these mappings, powerful filtering is achieved. For instance, the simple mappings to remove a single leading or trailing space (6) or reduce two spaces to a single space (7) become, when taken as a whole, a filter which strips all leading and trailing spaces and reduces all consecutive multiple spaces to a single space. Such filtering helps reduce the size of each SMS message.

1. These two entries cause **From:** and **Subject:** header lines to be included in an SMS message. **From:** and **Subject:** are abbreviated as, respectively, **F:** and **S:**. Some of the other entries may have further effects on **From:** and **Subject:** header lines.

This entry will reduce a **From:** header line containing a <...> pattern to only the text within the angle brackets. For example:

F: "John C. Doe" <jdoe@example.org> (Hello) will be replaced with:**F: jdoe@example.org**

2. This entry will remove, inclusively, everything inside of a (...) pattern in a **From:** header line. For example:

F: "John C. Doe" <jdoe@example.org> (Hello) will be replaced with:**F: "John C. Doe" <jdoe@example.org>**

3. This entry will remove, inclusively, everything inside of a "..." pattern in a **From:** header line. For example:

F: "John C. Doe" <jdoe@example.org> (Hello) will be replaced with:**F: <jdoe@example.org> (Hello)**

4. This entry will remove, inclusively, everything to the right of an at-sign, @, in a **From:** header line. For example:

F: "John C. Doe" <jdoe@example.org> (Hello) will be replaced with:**F: "John C. Doe" <jdoe@**

5. These four entries remove leading and trailing spaces from lines in the message header and body.
6. These two entries reduce two spaces to a single space in lines of the message header and body.
7. These four entries reduce double dashes, periods, exclamation and question marks to single occurrences of the matching character. Again, this helps save bytes in an SMS message.

The order of the entries is very important. For instance, with the given ordering, the body of the message **From:** header line:

From: "John C. Doe" (Hello)

will be reduced to:

jdoe

The steps taken to arrive at this are as follows:

1. We begin with the From: header line:

From: "John C. Doe" (Hello)

The pattern in the first mapping entry matches this and produces the result:

F: "John C. Doe" (Hello)

The \$R metacharacter in the result string causes the result string to be remapped.

2. The mapping is applied to the result string of the last step. This produces:

F: jdoe@example.org

The \$R in the mapping causes the entire set of mappings to be re-applied to the result of this step.

3. Next, the mapping is applied producing:

F: jdoe The \$R in the mapping causes the entire set of mappings to be re-applied to the result of this step.

4. Next, the mapping is applied producing:

F:jdoe The \$R in the mapping causes the entire set of mappings to be re-applied to the result of this step.

5. Since no other entries match, the final result string:

F:jdoe is incorporated into the SMS message.

Note: The **imsimta test -mapping** utility may be used to test a mapping table. For instance,

```
imsimta test -mapping -noimage_file -mapping_file=test.txt
Enter table name: SMS_TEXT
Input string: H|From: "John C. Doe" (Hello)
Output string: H|F:jdoe
Output flags: [0,1,2,89]
Input string: ^D
```

For further details, see the "[imsimta test](#)" utility.

SMS Channel Configuration

This section gives directions on how to set up the SMS channel for both one-way (email-to-mobile) and two-way (email-to-mobile and mobile-to-email) functionality. The SMS channel is set up the same for both one-way and two-way functionality, with the exceptions noted in the topic "[Configuring the SMS Channel for Two-Way SMS](#)."

This section includes the following topics:

- [Adding an SMS Channel](#)
- [Setting SMS Channel Options](#)
- [Available Options](#)
- [Adding Additional SMS Channels](#)
- [Adjusting the Frequency of Delivery Retries](#)
- [Sample One-Way Configuration \(MobileWay\)](#)
- [Configuring the SMS Channel for Two-Way SMS](#)

Adding an SMS Channel

Two steps are required to add an SMS channel to a Messaging Server configuration:

1. [Adding the Channel Definition and Rewrite Rules](#).
2. [Setting SMS Channel Options](#).

While there are no channel options which must be set in all situations, it is likely that one or more of the following options may need to be set: "[ESME_PASSWORD](#),"

"ESME_SYSTEM_ID," "MAX_PAGE_SIZE," "DEFAULT_SOURCE_TON," and "DEFAULT_DESTINATION_TON." And, as described, the SMPP server's hostname or IP address and TCP port must be set with channel options.

You may configure more than one SMS channel, giving different characteristics to different SMS channels. See ["Adding Additional SMS Channels"](#) for further information on the use of multiple SMS channels.

Note for the instructions that follow: if you change channel definitions or rewrite rules, you must recompile.

Note also that the time before a channel change takes effect can differ depending on what the change is. Many channel option changes take effect in all channels started since the change was made, which may seem almost instantaneous since the Job Controller is often starting new channels. Some of the changes don't take effect until you recompile and restart the SMTP server. These options are processed as a message is enqueued to the channel and not when the channel itself runs.

Adding the Channel Definition and Rewrite Rules

To add the channel definition and rewrite rules, do the following:

To Add Channel Definition and Rewrite Rules

1. Before adding an SMS channel to the MTA's configuration, you need to pick a name for the channel. The name of the channel may be either **sms** or **sms_x** where **x** is any case-insensitive string whose length is between one and thirty-six bytes. For example, **sms_mway**.
2. To add the channel definition, run **msconfig edit channels**. At the bottom of the channel definitions, add a blank line followed by the two lines:

```
_channel-name_ port _p_ threaddepth _t_ \ backoff "pt2m" "pt5m" "pt10m"
"pt30m" notices 1
_smpp-host-name_
```

where *channel-name* is the name you chose for the channel, *p* is the TCP port the SMPP server listens on, *t* is the maximum simultaneous number of SMPP server connections per delivery process, and *smpp-host-name* is the host name of the system running the SMPP server. For example, you might specify a channel definition as follows:

```
sms_mway port 55555 threaddepth 20 \
backoff "pt2m" "pt5m" "pt10m" "pt30m" notices 1
smpp.example.org
```

For instructions on how to calculate **threaddepth**, see ["Controlling the Number of Simultaneous Connections."](#)

See ["Adjusting the Frequency of Delivery Retries"](#) for a discussion of the **backoff** and **notices** channel options.

If you wish to specify an IP address rather than a host name, for **smpp-host-name**, specify a domain literal. For example, if the IP address is 127.0.0.1, then specify [127.0.0.1] for **smpp-host-name**. Alternatively, consider using the "SMPP_SERVER" channel option.

Note: The use of the **master** channel option is ignored if present.

- Once the channel definition has been added, run **msconfig edit rewrite** and add a rewrite rule in this format:

```
smpp-host-name$u@smpp-host-name
```

For example, **smpp.example.org \$u@smpp.example.org**

- Save the changes.
- Recompile the configuration with the **imsimta cnbuild** command.
- Restart the SMTP server with the **imsimta restart dispatcher** command.
- With the above configuration, email messages are directed to the channel by addressing them to **id@smpp-host-name** (for example, **123456@smpp.example.org**).

See "[The Email to SMS Conversion Process](#)" for further information on addressing.

- Optionally, if you wish to hide the SMPP server's host name from users or associate other host names with the same channel, then add additional rewrite rules. For instance, to associate **host-name-1** and **host-name-2** with the channel, add the following to rewrite rules:

```
host-name-1 $U%host-name-1@smpp-host-name
host-name-2 $U%host-name-2@smpp-host-name
```

For example, if the SMPP server's host name is **smpp.example.org** but you want users to address email to **id@sms.example.com**, then add the rewrite rule: **sms.example.com \$U%sms.example.com@smpp.example.org**

Note that the "**SMPP_SERVER**" and "**SMPP_PORT**" channel options will override the channel's official host name and **port** channel option settings. When the **SMPP_PORT** option is used, it is not necessary to also use the **port** option. The advantage of using these two options is that they can be put into effect and subsequently changed without needing to recompile the configuration. An additional use of the **SMPP_SERVER** option is described in "[Adding Additional SMS Channels](#)."

Controlling the Number of Simultaneous Connections

The **threaddepth** channel option controls the number of messages to assign to each delivery thread within a delivery process. To calculate the total number of concurrent connections allowed, multiply the values of the two following options: **SMPP_MAX_CONNECTIONS** and **job_limit** (**SMPP_MAX_CONNECTIONS * job_limit**). The "**SMPP_MAX_CONNECTIONS**" option controls the maximum number of delivery threads in a delivery process. And, the **job_limit** option, for the Job Controller processing pool in which the channel is run, controls the maximum number of simultaneous delivery processes.

To limit the total number of concurrent connections, you must adjust appropriately either or both of these options. For instance, if the remote SMPP server allows only a single connection, then both **SMPP_MAX_CONNECTIONS** and **job_limit** must be set to **1**. When adjusting the values, it's preferable to allow **job_limit** to exceed **1**.

Setting SMS Channel Options

In general, a channel options contain site-specific options required for the operation of the channel. Channel options are not required for SMS. If you determine that some are necessary for your installation, set them using **msconfig**. For example:

```
msconfig
```

```

msconfig> set channel:sms_mway.options.profile GSM
msconfig# set channel:sms_mway.options.smsc_default_charset iso-8859-1
msconfig# set channel:sms_mway.options.use_ucs2 1
msconfig# write

```

For a list of available SMS channel options and a description of each, see "[Available Options](#)."

Available Options

The SMS channel contains a number of options which divide into six broad categories:

- *Email to SMS conversion*: Options which control the email to SMS conversion process.
- *SMS Gateway Server Option*: Gateway profile option.
- *SMS fields*: Options which control SMS-specific fields in generated SMS messages.
- *SMPP protocol*: Options associated with the use of the SMPP protocol over TCP/IP.
- *Localization*: Options which allow for localization of text fields inserted into SMS messages.
- *Miscellaneous*: Debug and logging options.

These options are summarized in the table below, and described more fully in the sections which follow.

Table C-5 SMS Channel Options

Table 27-5 SMS Channel Options

Option	Description	Default
Email to SMS Conversion	-	-
<code>GATEWAY_NOTIFICATIONS</code>	Specify whether or not to convert email notification messages to SMS messages.	0
<code>MAX_MESSAGE_PARTS</code>	Max. number of message parts to extract from an email message.	2
<code>MAX_MESSAGE_SIZE</code>	Maximum number of bytes to extract from an email message.	960
<code>MAX_PAGE_SIZE</code>	Maximum number of bytes to put into a single SMS message.	160
<code>MAX_PAGES_PER_MESSAGE</code>	Max. number of SMS messages to break an email message into.	6
<code>ROUTE_TO</code>	Route SMS messages to the specified IP host name.	NA
<code>SMSC_DEFAULT_CHARSET</code>	The default character set used by the SMSC.	US-ASCII
<code>USE_HEADER_FROM</code>	Set the SMS source address.	0
<code>USE_HEADER_PRIORITY</code>	Control the use of priority information from the email message's header.	1
<code>USE_HEADER_REPLY_TO</code>	Control the use of Reply-to : header lines when generating SMS source addresses.	0
<code>USE_HEADER_SENSITIVITY</code>	Control the use of privacy information from the email message's header.	1
<code>USE_UCS2</code>	Use the UCS2 character set in SMS messages when applicable.	1

Table 27-5 (Cont.) SMS Channel Options

Option	Description	Default
SMS Gateway Server Option	-	-
GATEWAY_PROFILE	Match the gateway profile name configured in the SMS Gateway Server's configuration file, <code>sms_gateway.cnf</code> .	NA
SMS Fields Options	-	-
DEFAULT_DESTINATION_NPI	Default NPI for SMS destination addresses.	0x00
DEFAULT_DESTINATION_TON	Default TON for SMS destination addresses.	0x01
DEFAULT_PRIORITY	Default priority setting for SMS messages.	0=GSM, CDMA1=TDMA
DEFAULT_PRIVACY	Default privacy value flag for SMS messages.	-1
DEFAULT_SERVICE_TYPE	SMS application service associated with submitted SMS messages.	NA
DEFAULT_SOURCE_ADDRESS	Default SMS source address.	0
DEFAULT_SOURCE_NPI	Default NPI for SMS source addresses.	0x00
DEFAULT_SOURCE_TON	Default TON for SMS source addresses.	0x01
DEFAULT_VALIDITY_PERIOD	Default validity period for SMS messages.	NA
DESTINATION_ADDRESS_NUMERIC	Reduce the destination SMS address to only the characters 0 - 9.	0
DESTINATION_ADDRESS_PREFIX	Text string to prefix destination SMS addresses with.	N/A
PROFILE	SMS profile to use.	GSM
USE_SAR	Sequence multiple SMS messages using the SMS <code>sar_</code> fields.	0
SMPP Protocol Options	-	-
ESME_ADDRESS_NPI	ESME NPI to specify when binding to the SMPP server.	0x00
ESME_ADDRESS_TON	ESME TON to specify when binding to the SMPP server.	0x00
ESME_IP_ADDRESS	IP address of the host running Messaging Server.	NA
ESME_PASSWORD	Password to present when binding to the SMPP server.	NA
ESME_SYSTEM_ID	System identification to present to the SMSC when binding.	NA
ESME_SYSTEM_TYPE	System type to present to the SMSC when binding.	NA
MAX_PAGES_PER_BIND	Maximum number of SMS messages to submit during a single session with an SMPP server.	1024
REVERSE_ORDER	Transmission sequence of multi-part SMS messages.	0
SMPP_MAX_CONNECTIONS	Maximum number of simultaneous SMPP server connections.	20
SMPP_PORT	For one-way SMS, TCP port the SMPP server listens on. For two-way SMS, same TCP port used for the <code>LISTEN_PORT</code> for the SMPP relay.	NA

Table 27-5 (Cont.) SMS Channel Options

Option	Description	Default
SMPP_SERVER	For one-way SMS, host name of the SMPP server to connect to. For two-way SMS, set to point to the host name or IP address of the SMS Gateway server. If using the SMPP relay's LISTEN_INTERFACE_ADDRESS option, then be sure to use the host name or IP address associated with the specified network interface address.	NA
TIMEOUT	Timeout for completion of reads and writes with the SMPP server.	30
Localization Options	-	-
CONTENT_PREFIX	Text to introduce the content of the email message.	Msg:
DSN_DELAYED_FORMAT	Formatting string for delivery delay notifications.	an empty string
DSN_FAILED_FORMAT	Formatting string for delivery failure notifications.	see description
DSN_RELAYED_FORMAT	Formatting string for relay notifications.	see description
DSN_SUCCESS_FORMAT	Formatting string to successful delivery notifications.	see description
FROM_FORMAT	Text to display indicating the originator of the email message.	\$a
FROM_NONE	Text to display when there is no originator.	NA
LANGUAGE	(i-default) Language group to select text fields from.	i-default
LINE_STOP	Text to place at the end of each line extracted from the email message.	space character
NO_MESSAGE	Text to indicate that the message had no content.	[no message]
SUBJECT_FORMAT	Text to display indicating the subject of the email message.	\$s
SUBJECT_NONE	Text to display when there is no subject for the email message.	NA
Miscellaneous Options	-	-
DEBUG	Enable verbose debug output.	6
LISTEN_CONNECTION_MAX	Maximum number of concurrent, inbound TCP connections to allow across all SMPP relay and server instantiations.	10,000
LOG_PAGE_COUNT	Controls the value recorded in the mail.log file's message size field to be page count instead of blocks.	0

Email to SMS Conversion Options

The following options control the conversion of email messages to SMS messages. The value range for the options are in parenthesis. In general, a given email message may be converted into one or more SMS messages. See "The Email to SMS Conversion Process."

GATEWAY_NOTIFICATIONS

(0 or 1) Specifies whether or not to convert email notifications to SMS notifications. Email notification messages must conform to RFCs 1892, 1893, 1894. The default value is 0.

When GATEWAY_NOTIFICATIONS=0, such notifications are discarded and are not converted to SMS notifications.

To enable the notifications to be converted to SMS notifications, set **GATEWAY_NOTIFICATIONS=1**. When the option set to *1*, the localization options (**DSN_*_FORMAT**) control which notification types (success, failure, delay, relayed) are converted into SMS messages and sent through the gateway. (If the notification type has a value of an empty string, then that type notification is not converted into SMS messages.)

MAX_MESSAGE_PARTS

(*integer*) When converting a multi-part email message to an SMS message, only the first **MAX_MESSAGE_PARTS** number of text parts will be converted. The remaining parts are discarded. By default, **MAX_MESSAGE_PARTS** is 2. To allow an unlimited number of message parts, specify a value of -1. When a value of 0 is specified, then no message content will be placed into the SMS message. This has the effect of using only header lines from the email message (for example, **Subject:**) to generate the SMS message.

Note that an email message containing both text and an attachment will typically consist of two parts. Note further that only plain text message parts are converted. All other MIME content types are discarded.

MAX_MESSAGE_SIZE

(*integer, >= 10*) With this option, an upper limit may be placed on the total number of bytes placed into the SMS messages generated from an email message. Specifically, a maximum of **MAX_MESSAGE_SIZE** bytes will be used for the one or more generated SMS messages. Any additional bytes are discarded.

By default, an upper limit of 960 bytes is imposed. This corresponds to **MAX_MESSAGE_SIZE=960**. To allow any number of bytes, specify a value of zero.

The count of bytes used is made after converting the email message from Unicode to either the SMSC's default character set or UCS2. This means, in the case of UCS2, that a **MAX_MESSAGE_SIZE** of 960 bytes will yield, at most, 480 characters since each UCS2 character is at least two bytes long.

Note that the **MAX_MESSAGE_SIZE** and "**MAX_PAGES_PER_MESSAGE**" options both serve the same purpose: to limit the overall size of the resulting SMS messages. Indeed, "**MAX_PAGE_SIZE**"=960 and "**MAX_PAGE_SIZE**"=160 implies **MAX_PAGES_PER_MESSAGE**=6. So why are there two different options? So as to allow control of the overall size or number of pages without having to consider the maximal size of a single SMS message, **MAX_PAGE_SIZE**. While this may not be important in the channel options themselves, it is important when directing email to the channel described in "[Directing Email to the Channel](#)" and addressing attributes described in [Table 27-1, "SMS Attributes"](#).

Finally, note that the smaller of the two limits of **MAX_MESSAGE_SIZE** and **MAX_PAGE_SIZE * MAX_PAGES_PER_MESSAGE** is used.

MAX_PAGE_SIZE

(*integer, >= 10*) The maximum number of bytes to allow in a single SMS message is controlled with the **MAX_PAGE_SIZE** option. By default, a value of 160 bytes is used. This corresponds to **MAX_PAGE_SIZE=160**.

MAX_PAGES_PER_MESSAGE

(*integer, 1 - 255*) The maximum number of SMS messages to generate for a given email message is controlled with this option. In effect, this option truncates the email message, only converting to SMS messages that part of the email message which fits

into **MAX_PAGES_PER_MESSAGE** SMS messages. See the description of the "**MAX_PAGE_SIZE**" option for further discussion.

By default, **MAX_PAGES_PER_MESSAGE** is set to the larger of 1 or "**MAX_MESSAGE_SIZE**" divided by "**MAX_PAGE_SIZE**."

ROUTE_TO

(string, IP host name, 1-64 bytes) All SMS messages targeted to the profile will be rerouted to the specified IP host name using an email address of the form:

```
SMS-destination-address@route-to
```

where **SMS-destination-address** is the SMS message's destination address and the **route-to** is the IP host name specified with this option. The entire content of the SMS message is sent as the content of the resulting email message. The **PARSE_RE_*** options are ignored.

Note: Use of **PARSE_RE_*** and **ROUTE_TO** options are mutually exclusive. Use of both in the same gateway profile is a configuration error.

SMSC_DEFAULT_CHARSET

(string) With this option, the SMSC's default character set may be specified. Use the character set names given in the file

```
installation-directory/config/charsets.txt
```

When this option is not specified, then US-ASCII is assumed. Note that the mnemonic names used in **charsets.txt** are defined in **charnames.txt** in the same directory.

When processing an email message, the header lines and text message parts are first decoded and then converted to Unicode. Next, the data is then converted to either the SMSC's default character set or UCS2, depending on the value of the "**USE_UCS2**" option and whether or not the SMS message contains at least one glyph not found in the default SMSC character set. Note that the UCS2 character set is a 16-bit encoding of Unicode and is often referred to as UTF-16.

USE_HEADER_FROM

(integer, 0-2) Set this option to allow the **From:** address to be passed to the SMSC. The value indicates where the **From:** address is taken from and what format it will have. Table 27-6 shows the allowable values and their meaning.

Table 27-6 *USE_HEADER_FROM* Values

Value	Description
0	SMS source address never set from the From: address. Use attribute-value pair found
1	SMS source address set to from-local@from-domain , where the From: address is: @from-route:from-local@from-domain
2	SMS source address set to <i>from-local</i> , where the From: address is: @from-route:from-local@from-domain

USE_HEADER_PRIORITY

(0 or 1) This option controls handling of RFC 822 **Priority:** header lines. By default, information from the **Priority:** header line is used to set the resulting SMS message's

priority flag, overriding the default SMS priority specified with the "[DEFAULT_PRIORITY](#)" option. This case corresponds to `USE_HEADER_PRIORITY=1`. To disable use of the RFC 822 **Priority:** header line, specify `USE_HEADER_PRIORITY=0`.

See the description of the "[DEFAULT_PRIORITY](#)" option for further information on the handling the SMS priority flag.

USE_HEADER_REPLY_TO

(0 or 1) When `USE_HEADER_FROM=1`, this option controls whether or not a **Reply-to:** or **Resent-reply-to:** header line is considered for use as the SMS source address. By default, **Reply-to:** and **Resent-reply-to:** header lines are ignored. This corresponds to an option value of 0. To enable consideration of these header lines, use an option value of 1.

Note that RFC 2822 has deprecated the use of **Reply-to:** and **Resent-reply-to:** header lines.

USE_HEADER_SENSITIVITY

(0 or 1) The `USE_HEADER_SENSITIVITY` option controls handling of RFC 822 **Sensitivity:** header lines. By default, information from the **Sensitivity:** header line is used to set the resulting SMS message's privacy flag, overriding the default SMS privacy specified with the "[DEFAULT_PRIVACY](#)" option. This case, which is the default, corresponds to `USE_HEADER_SENSITIVITY=1`. To disable use of RFC 822 **Sensitivity:** header lines, specify `USE_HEADER_SENSITIVITY=0`.

See the description of the "[DEFAULT_PRIVACY](#)" option for further information on the handling the SMS privacy flag.

USE_UCS2

(0 or 1) When appropriate, the channel will use the UCS2 character set in the SMS messages it generates. This is the default behavior and corresponds to `USE_UCS2=1`. To disable the use of the UCS2 character set, specify `USE_UCS2=0`. See the description of the "[SMSC_DEFAULT_CHARSET](#)" option for further information on character set issues.

Table 27-7 Valid Values for USE_UCS2

USE_UCS2	Result
1 (default)	The SMSC default character set will be used whenever possible. When the originating email message contains glyphs not in the SMSC default character set, then the UCS2 character set will be used.
0	The SMSC default character set will always be used. Glyphs not available in that character set will be represented by mnemonics (for example, "AE" for AE-ligature).

SMS Gateway Server Option

The following option describes the SMS Gateway Server.

GATEWAY_PROFILE

The name of the gateway profile in the SMS Gateway Server configuration file, `sms_gateway.cnf`.

SMS Options

The following options allow for specification of SMS fields in generated SMS messages.

DEFAULT_DESTINATION_NPI

(*integer, 0 - 255*) By default, destination addresses will be assigned an NPI (Numeric Plan Indicator) value of zero. With this option, an alternate integer value in the range 0 to 255 may be assigned. Typical NPI values include those found in [Table 27–8](#) that follows:

Table 27–8 *Numeric Plan Indicator Values*

Value	Description
0	Unknown
1	ISDN (E.163, E.164)
3	Data (X.121)
4	Telex (F.69)
6	Land Mobile (E.212)
8	National
9	Private
10	ERMES
14	IP address (Internet)
18	WAP client ID
>= 19	Undefined

Values for this option may be specified in one of three ways:

- A decimal value (for example, 10).
- A hexadecimal value prefixed by "0x" (for example, 0x0a).
- One of the following case-insensitive text strings (the associated decimal value is shown in parentheses): data (3), default (0), e.163 (1), e.164 (1), e.212 (6), ermes (10), f.69 (4), Internet (14), ip (14), isdn (1), land-mobile (6), national (8), private (9), telex (4), unknown (0), wap (18), x.121 (3).

DEFAULT_DESTINATION_TON

(*integer, 0 - 255*) By default, destination addresses will be assigned a TON (Type of Number) designator value of zero. With this option, an alternate integer value in the range 0 to 255 may be assigned. Typical TON values include those found in [Table 27–9](#) that follows:

Table 27–9 *Typical TON Values*

Value	Description
0	Unknown
1	International
2	National
3	Network specific
4	Subscriber number

Table 27–9 (Cont.) Typical TON Values

Value	Description
5	Alphanumeric
6	Abbreviated
>=7	Undefined

Values for this option may be specified in one of three ways:

- A decimal value (for example, 10)
- A hexadecimal value prefixed by "0x" (for example, 0x0a)
- One of the following case-insensitive text strings (the associated decimal value is shown in parentheses): abbreviated (6), alphanumeric (5), default (0), international (1), national (2), network-specific (3), subscriber (4), unknown (0).

DEFAULT_PRIORITY

(integer, 0 - 255) SMS messages have a mandatory priority field. The interpretation of SMS priority values is shown in Table 27–10 that follows:

Table 27–10 SMS Priority Values Interpreted for Each SMS Profile Type

Value	GSM	TDMA	CDMA
0	Non-priority	Bulk	Normal
1	Priority	Normal	Interactive
2	Priority	Urgent	Urgent
3	Priority	Very urgent	Emergency

With this option, the default priority to assign to SMS messages may be specified. When not specified, a default priority of 0 is used for **PROFILE=GSM** and **CDMA**, and a priority of 1 for "**PROFILE**"=TDMA.

Note that if "**USE_HEADER_PRIORITY**"=1 and an email message has an RFC 822 **Priority:** header line, then the priority specified in that header line will instead be used to set the priority of the resulting SMS message. Specifically, if **USE_HEADER_PRIORITY**=0, then the SMS priority flag is always set in accord with the **DEFAULT_PRIORITY** option and the RFC 822 **Priority:** header line is always ignored. If **USE_HEADER_PRIORITY**=1, then the originating email message's RFC 822 **Priority:** header line is used to set the SMS message's priority flag. If that header line is not present, then the SMS priority flag is set using the **DEFAULT_PRIORITY** option.

The mapping used to translate RFC 822 **Priority:** header line values to SMS priority flags is shown in table that follows:

Table 27–11 Mapping for Translating Priority Header to SMS Priority Flags

RFC 822	SMS Priority Flag	-	-
Priority: value	GSM	TDMA	CDMA
Third	Non-priority (0)	Bulk (0)	Normal (0)
Second	Non-priority (0)	Bulk (0)	Normal (0)
Non-urgent	Non-priority (0)	Bulk (0)	Normal (0)

Table 27–11 (Cont.) Mapping for Translating Priority Header to SMS Priority Flags

RFC 822	SMS Priority Flag	-	-
Normal	Non-priority (0)	Normal (1)	Normal (0)
Urgent	Priority (1)	Urgent (2)	Urgent (2)

DEFAULT_PRIVACY

(integer, -1, 0 - 255) Whether or not to set the privacy flag in an SMS message, and what value to use is controlled with the **DEFAULT_PRIVACY** and "**USE_HEADER_SENSITIVITY**" options. By default, a value of -1 is used for **DEFAULT_PRIVACY**. [Table 27–12](#) that follows shows the result of setting the **DEFAULT_PRIVACY** and "**USE_HEADER_SENSITIVITY**" options to various values.

Table 27–12 Result of Values for DEFAULT_PRIVACY and USE_HEADER_SENSITIVITY

DEFAULT_PRIVACY	USE_HEADER_SENSITIVITY	Result
-1	0	The SMS privacy flag is never set in SMS messages.
n >= 0	0	The SMS privacy flag is always set to the value n. RFC 822 Sensitivity : header lines are always ignored.
-1 (default)	1 (default)	The SMS message's privacy flag is only set when the originating email message has an RFC 822 Sensitivity : header line. In that case, the SMS privacy flag is set to correspond to the Sensitivity : header line's value. This is the default.
n >= 0	1	The SMS message's privacy flag is set to correspond to the originating email message's RFC 822 Sensitivity : header line. If the email message does not have a Sensitivity : header line, then the value of the SMS privacy flag is set to n.

The SMS interpretation of privacy values is shown in [Table 27–13](#) that follows:

Table 27–13 SMS Interpretation of Privacy Values

Value	Description
0	Unrestricted
1	Restricted
2	Confidential
3	Secret
>= 4	Undefined

The mapping used to translate RFC 822 **Sensitivity**: header line values to SMS privacy values is shown in [Table 27–14](#) that follows:

Table 27–14 Mapping Translation of Sensitivity Headers to SMS Privacy Values

RFC 822	SMS Privacy Value
Personal	1 (Restricted)
Private	2 (Confidential)
Company confidential	3 (Secret)

DEFAULT_SERVICE_TYPE

(string, 0 - 5 bytes) Service type to associate with SMS messages generated by the channel. By default, no service type is specified (that is, a zero length string). Some common service types are: **CMT** (cellular messaging), **CPT** (cellular paging), **VMN** (voice mail notification), **VMA** (voice mail alerting), **WAP** (wireless application protocol), and **USSD** (unstructured supplementary data services).

DEFAULT_SOURCE_ADDRESS

(string, 0 - 20 bytes) Source address to use for SMS messages generated from email messages. Note that the value specified with this option is overridden by the email message's originator address when **USE_HEADER_FROM=1**. By default, the value is disabled, that is, has a value of 0.

DEFAULT_SOURCE_NPI

(integer, 0 - 255) By default, source addresses will be assigned an NPI value of zero. With this option, an alternate integer value in the range 0 to 255 may be assigned. See the description of the "**DEFAULT_SOURCE_NPI**" option for a table of typical NPI values.

DEFAULT_SOURCE_TON

(integer, 0 - 255) By default, source addresses will be assigned a TON designator value of zero. With this option, an alternate integer value in the range 0 to 255 may be assigned. See the description of the "**DEFAULT_DESTINATION_TON**" option for a table of typical TON values.

DEFAULT_VALIDITY_PERIOD

(string, 0 - 252 bytes) By default, SMS messages are not given a relative validity period; instead, they use the SMSC's default value. Use this option to specify a different relative validity period. Values may be specified in units of seconds, minutes, hours, or days. [Table 27-15](#) that follows specifies the format and description of the various values for this option:

Table 27-15 **DEFAULT_VALIDITY_PERIOD** Format and Values

Format	Description
<i>nnn</i>	Implicit units of seconds; for example, 604800
<i>nnns</i>	Units of seconds; for example, 604800s
<i>nnnm</i>	Units of minutes; for example, 10080m
<i>nnnh</i>	Units of hours; for example, 168h
<i>nnnd</i>	Units of days; for example, 7d

A specification of 0, 0s, 0m, 0h, or 0d may be used to select the SMSC's default validity period. That is, when a specification of 0, 0s, 0m, 0h, or 0d is used, an empty string is specified for the validity period in generated SMS messages.

Note that this option does not accept values in UTC format.

DESTINATION_ADDRESS_NUMERIC

(0 or 1) Use this option to strip all non-numeric characters from the SMS destination address extracted from the email envelope **To:** address. For instance, if the envelope **To:** address is:

"(800) 555-1212"@sms.example.org

then it will be reduced to:

8005551212@sms.example.org

To enable this stripping, specify a value of 1 for this option. By default, this stripping is disabled which corresponds to an option value of 0. Note that when enabled, the stripping is done before any destination address prefix is added via the "DESTINATION_ADDRESS_PREFIX" option.

DESTINATION_ADDRESS_PREFIX

(string) In some instances, it may be necessary to ensure that all SMS destination addresses are prefixed with a fixed text string; for example, "+". This option may be used to specify just such a prefix. The prefix will then be added to any SMS destination address which lacks the specified prefix. To prevent being stripped by the "DESTINATION_ADDRESS_NUMERIC" option, this option is applied after the DESTINATION_ADDRESS_NUMERIC option.

PROFILE

(string) Specify the SMS profiling to be used with the SMSC. Possible values are **GSM**, **TDMA**, and **CDMA**. When not specified, GSM is assumed. This option is only used to select defaults for other channel options such as "DEFAULT_PRIORITY" and "DEFAULT_PRIVACY."

USE_SAR

(0 or 1) Sufficiently large email messages may need to be broken into multiple SMS messages. When this occurs, the individual SMS messages can optionally have sequencing information added using the SMS sar_ fields. This produces a "segmented" SMS message which can be re-assembled into a single SMS message by the receiving terminal. Specify **USE_SAR=1** to indicate that this sequencing information is to be added when applicable. The default is to not add sequencing information and corresponds to **USE_SAR=0**.

When **USE_SAR=1** is specified, the "REVERSE_ORDER" option is ignored.

SMPP Options

The following options allow for specification of SMPP protocol options. The options with names beginning with the string "ESME_" serve to identify the MTA when it acts as an External Short Message Entity (ESME); that is, when the MTA binds to an SMPP server in order to submit SMS messages to the server's associated SMSC.

ESME_ADDRESS_NPI

(integer, 0 - 255) By default, bind operations will specify an ESME NPI value of zero indicating an unknown NPI. With this option, an alternate integer value in the range 0 to 255 may be assigned. See the description of the "DEFAULT_DESTINATION_NPI" option for a table of typical NPI values.

ESME_ADDRESS_TON

(integer, 0 - 255) By default, bind operations will specify an ESME TON value of 0. With this option, an alternate integer value in the range 0 to 255 may be assigned. See the description of the "DEFAULT_DESTINATION_TON" option for a table of typical TON values.

ESME_IP_ADDRESS

(*string, 0 - 15 bytes*) When binding to the SMPP server, the *BIND* PDU indicates that the client's (that is, ESME's) address range is an IP address. This is done by specifying a TON of 0x00 and an NPI of 0x0d. The value of the address range field is then set to be the IP address of the host running the SMS channel. Specify the IP address in dotted decimal format; for example, 127.0.0.1.

ESME_PASSWORD

(*string, 0 - 8 bytes*) When binding to the SMPP server, a password may be required. If so, then specify that password with this option. By default, a zero-length password string is presented.

ESME_SYSTEM_ID

(*string, 0 - 15 bytes*) When binding to the SMPP server, a system ID for the MTA may be supplied. By default, no system ID is specified (that is, a zero-length string is used). To specify a system ID, use this option.

ESME_SYSTEM_TYPE

(*string, 0 - 12 bytes*) When binding to the SMPP server, a system type for the MTA may be supplied. By default, no system type is specified (that is, a zero-length string is used).

MAX_PAGES_PER_BIND

(*integer, >= 0*) Some SMPP servers may limit the maximum number of SMS messages submitted during a single, bound session. In recognition of this, this option allows specification of the maximum number of SMS messages to submit during a single session. Once that limit is reached, the channel will unbind, close the TCP/IP connection, re-connect, and then rebind.

By default, a value of 1024 is used for **MAX_PAGES_PER_BIND**. Note that the channel will also detect **ESME_RTHROTTLED** errors and adjust **MAX_PAGES_PER_BIND** during a single run of the channel accordingly.

REVERSE_ORDER

(*0 or 1*) When an email message generates more than one SMS message, those SMS messages can be submitted to the SMSC in sequential order (**REVERSE_ORDER=0**), or reverse sequential order (**REVERSE_ORDER=1**). Reverse sequential order is useful for situations where the receiving terminal displays the last received message first. In such a case, the last received message will be the first part of the email message rather than the last. By default, **REVERSE_ORDER=1** is used.

Note that this option is ignored when "**USE_SAR**"=1 is specified.

SMPP_MAX_CONNECTIONS

(*integer, 1 - 50*) This option controls the maximum number of simultaneous SMPP connections per process. As each connection has an associated thread, this option also places a limit on the maximum number of "worker" threads per process. By default, **SMPP_MAX_CONNECTIONS=20**.

SMPP_PORT

(*integer, 1 - 65535*) The TCP port which the SMPP server listens on may be specified with either this option or the *port* channel option. This port number must be specified through either of these two mechanisms. If it is specified with both mechanisms, then

the setting made with the **SMPP_PORT** option takes precedence. Note that there is no default value for this option.

For two-way SMS, make sure its the same port as the **LISTEN_PORT** for the SMPP relay.

SMPP_SERVER

(*string, 1 - 252 bytes*) For one-way SMS, by default, the IP host name of the SMPP server to connect to is the official host name associated with the channel; that is, the host name shown on the second line of the channel's definition in MTA's configuration. This option may be used to specify a different host name or IP address which will override that specified in the channel definition. When specifying an IP address, use dotted decimal notation; for example, 127.0.0.1.

For two-way SMS, set to point to the host name or IP address of the SMS Gateway Server. If using the SMPP relay's **LISTEN_INTERFACE_ADDRESS** option, then be sure to use the host name or IP address associated with the specified network interface address.

TIMEOUT

(*integer, >= 2*) By default, a timeout of 30 seconds is used when waiting for data writes to the SMPP server to complete or for data to be received from the SMPP server. Use the **TIMEOUT** option to specify, in units of seconds, a different timeout value. The specified value should be at least 1second.

Localization Options

In constructing SMS messages, the SMS channel has a number of fixed text strings it puts into those messages. These strings, for example, introduce the email's **From:** address and **Subject:** header line. With the channel options described in this section, versions of these strings may be specified for different languages and a default language for the channel then specified. "[Example Language Specification Options](#)" shows the language part of the option file:

Example Language Specification Options

```
msconfig
msconfig> set channel:mway_sms.options.language <default-language>
msconfig# set channel:mway_sms.options.from_prefix From:
msconfig# set channel:mway_sms.options.subject_prefix Subj:
msconfig# set channel:mway_sms.options.content_prefix Msg:
msconfig# set channel:mway_sms.options.line_stop
set channel:mway_sms.options.no_message [no message]
msconfig# set channel:mway_sms.options.reply_prefix re:
msconfig# write
msconfig> msconfig> set channel:mway_sms.options.language en
msconfig# set channel:mway_sms.options.from_prefix From:
msconfig# set channel:mway_sms.options.subject_prefix Subj:
msconfig# set channel:mway_sms.options.content_prefix Msg:
msconfig# set channel:mway_sms.options.line_stop
msconfig# set channel:mway_sms.options.no_message [no message]
msconfig# set channel:mway_sms.options.reply_prefix Re:
msconfig# write
```

```
LANGUAGE=_default-language_
```

```
[language=i-default]
FROM_PREFIX=From:
SUBJECT_PREFIX=Subj:
```

```
CONTENT_PREFIX=Msg:
LINE_STOP= NO_MESSAGE=[no message]
REPLY_PREFIX=Re:

[language=en]
FROM_PREFIX=From:
SUBJECT_PREFIX=Subj:
CONTENT_PREFIX=Msg:
LINE_STOP=
NO_MESSAGE=[no message]
REPLY_PREFIX=Re:
...
```

Within each [**language=x**] block, the localization options relevant to that language may be specified. If a particular option is not specified within the block, then the global value for that option is used. A localization option specified outside of a [**language=x**] block sets the global value for that option.

For the options listed below, the string values must be specified using either the US-ASCII or UTF-8 character sets. Note that the US-ASCII character set is a special case of the UTF-8 character set.

CONTENT_PREFIX

(string, 0 - 252 bytes) Text string to place in the SMS message before the content of the email message itself. Default global value is the US-ASCII string "Msg:".

DSN_DELAYED_FORMAT

(string, 0-256 characters) Formatting string for delivery delay notifications. By default, an empty string is used for this option, thereby inhibiting the conversion to SMS of delay notifications. Note that "GATEWAY_NOTIFICATIONS" must be set to **1** for this option to be in effect. This option is ignored when **GATEWAY_NOTIFICATIONS=0**.

DSN_FAILED_FORMAT

(string, 0-256 characters) Formatting string for permanent delivery failure notifications. The default value of this option is the string:

```
Unable to deliver your message to $a; no further delivery attempts will be
made.
```

To inhibit conversion of failure notifications, specify an empty string for this option. Note that "GATEWAY_NOTIFICATIONS" must be set to **1** for this option to be in effect. This option is ignored when **GATEWAY_NOTIFICATIONS=0**.

DSN_RELAYED_FORMAT

(string, 0-256 characters) Formatting string for relay notifications. The default value is the string:

```
Your message to $a has been relayed to a messaging system which may not
provide a final delivery confirmation
```

To inhibit conversion of relay notifications, specify an empty string for this option. Note that "GATEWAY_NOTIFICATIONS" must be set to **1** for this option to be in effect. This option is ignored when **GATEWAY_NOTIFICATIONS=0**.

DSN_SUCCESS_FORMAT

(string, 0-256 characters) Formatting string for successful delivery notifications. The default value is the string:

Your message to \$a has been delivered

To inhibit conversion of successful delivery notifications, specify an empty string for this option. Note that "GATEWAY_NOTIFICATIONS" must be set to **1** for this option to be in effect. This option is ignored when **GATEWAY_NOTIFICATIONS=0**.

FROM_FORMAT

(*string, 0 - 252 bytes*) Formatting template to format the originator information to insert into the SMS message. The default global value is the US-ASCII string "\$a" which substitutes in the originator's email address. See "Formatting Templates."

FROM_NONE

(*string, 0 - 252 bytes*) Text string to place in the SMS message when there is no originator address to display. The default global value is an empty string.

Note that normally, this option will never be used as sites will typically reject email messages which lack any originator address.

LANGUAGE

(*string, 0 - 40 bytes*) The default language group to select text strings from. If not specified, then the language will be derived from the host's default locale specification. If the host's locale specification is not available or corresponds to "C", then i-default will be used. (i-default corresponds to "English text intended for an international audience.")

LINE_STOP

(*string, 0 - 252 bytes*) Text string to place in the SMS message between lines extracted from the email message. The default global value is the US-ASCII space character, " ".

NO_MESSAGE

(*string, 0 - 252 bytes*) Text string to place in the SMS message to indicate that the email message had no content. The default global value is the US-ASCII string "[no message]".

SUBJECT_FORMAT

(*string, 0 - 252 bytes*) Formatting template to format the content of the **Subject:** header line for display in the SMS message. The global default value for this option is the US-ASCII string "\$s". See "Formatting Templates" for further details.

See the **SUBJECT_NONE** option for a description of the handling when there is no **Subject:** header line or the content of that header line is an empty string.

SUBJECT_NONE

(*string, 0 - 252 bytes*) Text string to display when the originating email message either has no **Subject:** header line, or the **Subject:** header line's value is an empty string. The default global value for this option is the empty string.

DEBUG

(*integer, bitmask*) Enable debug output. The default value is 6 which selects warning and error messages. Any non-zero value enables debug output for the channel itself, the same as specifying **master_debug** on the channel definition. Table 27-16 defines the bit values of the **DEBUG** bitmask.

Table 27–16 *DEBUG Bitmask*

Bit	Value	Description
0-31	-1	Extremely verbose output
0	1	Informational messages
1	2	Warning messages
3	4	Error messages
3	8	Subroutine call tracing
4	16	Hash table diagnostics
5	32	I/O diagnostics, receive
6	64	I/O diagnostics, transmit
7	128	SMS to email conversion diagnostics (mobile originate and SMS notification)
8	256	PDU diagnostics, header data
9	512	PDU diagnostics, body data
10	1024	PDU diagnostics, type-length-value data
11	2048	Option processing; sends all option settings to the log file.

Formatting Templates

The formatting templates specified with the "FROM_FORMAT," "SUBJECT_FORMAT," and all the DSN_* channel options are UTF-8 strings which may contain a combination of literal text and substitution sequences. Assuming the sample email address of

Jane Doe <user@example>

The recognized substitution sequences are shown in Table 27–17 that follows:

Table 27–17 *Substitution Sequences*

Sequence	Description
\$a	Replace with the local and domain part of the originator's email address (for example, "user@example")
\$d	Replace with the domain part of the originator's email address (for example, "domain")
\$p	Replace with the phrase part, if any, of the originator's email address (for example, "Jane Doe")
\$s	Replace with the content of the Subject: header line
\$u	Replace with the local part of the originator's email address (for example, "user")
\x	Replace with the literal character "x"

For example, the formatting template

From: \$a

produces the text string

From: user@example

The construct,

`#{xy:alternate text}`

may be used to substitute in the text associated with the sequence `x`. If that text is the empty string, the text associated with the sequence `y` is instead used. And, if that text is the empty string, to then substitute in the alternate text. For example, consider the formatting template

From: `#{pa:unknown sender}`

For the originator email address

John Doe <`jd@example.org`>

which has a phrase part, the template produces:

From: John Doe

However, for the address

`jd@example.org`

which lacks a phrase, it produces

From: `jd@example.org`

And for an empty originator address, it produces

From: unknown sender

Adding Additional SMS Channels

You may configure the MTA to have more than one SMS channel. There are two typical reasons to do this:

1. To communicate with different SMPP servers. This is quite straightforward: just add an additional SMS channel to the configuration, being sure to (a) give it a different channel name, and (b) associate different host names with it. For example,

```
sms_mway port 55555 threaddepth 20
smpp.example.org
```

```
sms_ace port 777 threaddepth 20
sms.ace.net
```

Note that no new rewrite rule is needed. If there is no directly matching rewrite rule, Messaging Server looks for a channel with the associated host name. For example, if the server is presented with `user@host.domain`, it would look for a channel of the name "host.domain". If it finds such a channel, it routes the message there. Otherwise, it starts looking for a rewrite rule for the ".domain" and if none is there, then for the dot (".") rule. For more information on rewrite rules, see ["Configuring Rewrite Rules."](#)

2. To communicate with the same SMPP server but using different channel options. To communicate with the same SMPP server, using different channel options, specify the same SMPP server in the "SMPP_SERVER" channel option for each channel definition. Using this mechanism is necessary since two different channels cannot have the same official host name (that is, the host name listed in the second line of the channel definition). To allow them to communicate with the same SMPP server, define two separate channels, with each specifying the same SMPP server in their "SMPP_SERVER" in their channel options. For example, you could have the following channel definitions,

```
sms_mway_1 port 55555 threaddepth 20
```

```
SMS-DAEMON-1  
  
sms_mway_2 port 55555 threaddepth 20  
SMS-DAEMON-2
```

and rewrite rules,

```
sms-1.example.org $u%sms-1.example.org@SMS-DAEMON-1  
sms-2.example.org $U%sms-2.example.org@SMS-DAEMON-2
```

Then, to have them both use the same SMPP server, each of these two channels would specify `"SMPP_SERVER"=smpp.example.org` in their channel options.

Adjusting the Frequency of Delivery Retries

When an SMS message cannot be delivered owing to temporary errors (for example, the SMPP server is not reachable), the email message is left in the delivery queue and retried again later. Unless configured otherwise, the Job Controller will not re-attempt delivery for an hour. For SMS messaging, that is likely too long to wait. As such, it is recommended that the backoff channel option be used with the SMS channel to specify a more aggressive schedule for delivery attempts. For example,

```
sms_mway port 55555 threaddepth 20 \  
backoff "pt2m" "pt5m" "pt10m" "pt30m" notices 1  
smpp.example.org
```

With the above settings, a redelivery attempt will be made at two minutes after the first attempt. If that then fails, then five minutes after the second attempt. Then ten minutes later and finally every thirty minutes. The `notices 1` channel option causes the message to be returned as undeliverable if it cannot be delivered after a day.

Sample One-Way Configuration (MobileWay)

The MTA SMS channel may be used with any SMPP V3.4 compatible SMPP server. For purposes of illustrating an example configuration, this section explains how to configure the SMS channel for use with a MobileWay SMPP server. MobileWay (<http://www.mobilway.com>) is a leading provider of global data and SMS connectivity. By routing your SMS traffic through MobileWay, you can reach SMS subscribers on most of the major SMS networks throughout the world.

When requesting an SMPP account with MobileWay, you may be asked to answer the following questions:

- IP address of your SMPP client: Supply the IP address of your Messaging Server system as seen by other domains on the Internet.
- Default validity period: This is the SMS validity period which MobileWay will use should a validity period not be specified in the SMS messages you submit. SMS messages which cannot be delivered before this validity period expires will be discarded. Supply a reasonable value (for example, 2 days, 7 days, etc.).
- Window size: This is the maximum number of SMS messages your SMPP client will submit before it will stop and wait for responses from the SMPP server before submitting any further SMS messages. You must supply a value of 1 message.
- Timezone: Specify the timezone in which your Messaging Server system operates. The timezone should be specified as an offset from GMT.
- Timeout: Not relevant to one-way SMS messaging.

- IP address and TCP port for outbound requests: Not relevant for one-way SMS messaging.

After supplying MobileWay with the answers to the above questions, they will provide you with an SMPP account and information necessary to communicate with their SMPP servers. This information includes

```
Account Address: a.b.c.d:p
Account Login: system-id
Account Passwd: secret
```

The Account Address field is the IP address, **a.b.c.d**, and TCP port number, **P**, of the MobileWay SMPP server you will be connecting to. Use these values for the "**SMPP_SERVER**" and "**SMPP_PORT**" channel options. The Account Login and Passwd are, respectively, the values to use for the "**ESME_SYSTEM_ID**" and "**ESME_PASSWORD**" channel options. Using this information, your channel's options should be set as follows:

```
msconfig
msconfig> set channel:mway_sms.options.smpp_server a.b.c.d
msconfig# set channel:mway_sms.options.smpp_port p
msconfig# set channel:mway_sms.options.esme_system_id system-id
msconfig# set channel:mway_sms.options.esme_password secret
msconfig# write
```

Now, to interoperate with MobileWay you need to make two additional option settings

```
msconfig
msconfig> set channel:mway_sms.options.esme_address_ton 0x01
msconfig# set channel:mway_sms.options.default_destination_ton 0x01
msconfig# write
```

The rewrite rule can appear as

```
sms.your-domain $u@sms.your-domain
```

And, the channel definition can appear as

```
sms_mobileway
sms.your-domain
```

Once the channel options, rewrite rules, and channel definitions are in place, a test message may be sent. MobileWay requires International addressing of the form

```
+<country-code><subscriber-number>
```

For instance, to send a test message to the North American subscriber with the subscriber number (800) 555-1212, you would address your email message to

```
+18005551212@sms.your-domain
```

Debugging

To debug the channel, specify the **master_debug** channel option in the channel's definition. For example,

```
sms_mway port 55555 threaddepth 20 \
backoff "pt2m" "pt5m" "pt10m" "pt30m" notices 1 master_debug
```

With the **master_debug** channel option, basic diagnostic information about the channel's operation will be output to the channel's log file. For verbose diagnostic

information about the SMPP transactions undertaken by the channel, also set the channel debug option to 1 by running:

```
msconfig set channel:mway_sms.options.debug -1
```

Configuring the SMS Channel for Two-Way SMS

For general directions on configuring the SMS channel, refer to earlier topics starting with "SMS Channel Configuration." Configure the SMS channel as though it will be talking directly to the remote SMSC, with the exceptions listed in Table 27–18 that follows:

Table 27–18 Two-Way Configuration Exceptions

Exception	Explanation
master channel option	Remove the master channel option, if present. It is no longer needed for SMS channel configuration.
SMPP_SERVER	Set to point to the host name or IP address of the SMS Gateway Server. If using the SMPP relay's LISTEN_INTERFACE_ADDRESS option (see "Configuration Options"), then be sure to use the host name or IP address associated with the specified network interface address.
SMPP_PORT	Same TCP port as used for the LISTEN_PORT setting used to instantiate the SMPP relay (see "An SMPP Relay").
DEFAULT_SOURCE_ADDRESS	Pick a value and then configure the remote SMSC to route this address back to the Gateway SMPP server. In the SMS channel's options, specify the chosen value with this option.
GATEWAY_PROFILE	Set to match the gateway profile name. See "A Gateway Profile."
USE_HEADER_FROM	Set to 0.

All other channel configurations should be done as described in the SMS Channel documentation.

As mentioned in "Setting Up Bidirectional SMS Routing," the remote SMSC needs to be configured to route the SMS address, defined in the **DEFAULT_SOURCE_ADDRESS** channel option, to the Gateway's SMPP server using the TCP port number specified with the **LISTEN_PORT** option. (For how to specify the **LISTEN_PORT**, see "An SMPP Server.")

Note that multiple SMS channels may use the same SMPP relay. Similarly, there need be only one SMPP server or gateway profile to handle SMS replies and notifications for multiple SMS channels. The ability to configure multiple relays, servers, and gateway profiles exists to effect different usage characteristics through configuration options.

SMS Gateway Server Theory of Operation

The SMS Gateway Server facilitates two-way SMS through mechanisms that allow mobile originated SMS messages to be matched to the correct email address. The following SMS Gateway Server topics are covered in this section:

- [Function of the SMS Gateway Server](#)
- [Behavior of the SMPP Relay and Server](#)
- [Remote SMPP to Gateway SMPP Communication](#)

- [SMS Reply and Notification Handling](#)

Function of the SMS Gateway Server

The SMS Gateway Server simultaneously functions as both an SMPP relay and server. It may be configured to have multiple "instantiations" of each function. For instance, it may be configured to have three different SMPP relays, each listening on different TCP ports or network interfaces and relaying to different remote SMPP servers. Similarly, it may be configured to have four different SMPP servers, each listening on different combinations of TCP ports and network interfaces.

The SMS Gateway Server may be configured with zero or more gateway profiles for sending SMS messages to email. Each gateway profile describes which destination SMS addresses match the profile, how to extract the destination email addresses from SMS messages, and various characteristics of the SMS to email conversion process. Each SMS message presented to the SMS Gateway Server through either its SMPP relay or server are compared to each profile. If a match is found, then the message is routed to email.

Finally, the gateway profiles also describe how to handle notification messages returned by remote SMSCs in response to previous email-to-mobile messages.

Behavior of the SMPP Relay and Server

When acting as an SMPP relay, the SMS Gateway Server attempts to be as transparent as possible, relaying all requests from local SMPP clients on to a remote SMPP server and then relaying back the remote server's responses. There are two exceptions:

- When a local SMPP client submits a message whose SMS destination address matches one of the configured gateway profiles, the submitted SMS message is sent directly back to email; the SMS message is not relayed to a remote SMPP server.
- When a local or remote SMPP client submits a message whose SMS destination address matches a unique SMS source address previously generated by the SMPP relay, the SMS message is a reply to a previously relayed message. This reply is directed back to the originator of the original message.

Note that typically the SMS Gateway Server will be configured such that the unique SMS source addresses which it generates match one of the gateway profiles.

Note: The SMS Gateway Server's SMPP relay is only intended for use with qualified, SMPP clients, that is, the Messaging Server's SMS channel. It is not intended for use with arbitrary SMPP clients.

When acting as an SMPP server, the SMS Gateway Server directs SMS messages to email for three circumstances:

- The SMS messages are mobile originated and match a gateway profile.
- The SMS messages are mobile originated and the SMS destination address matches a previously generated unique SMS source address.
- The SMS messages are SMS notifications which correspond to email-to-mobile messages previously relayed by the SMS Gateway Server's SMPP relay.

All other SMS messages are rejected by the SMPP server.

Remote SMPP to Gateway SMPP Communication

Remote SMPP clients communicate to the Gateway SMPP server with Protocol Data Units (PDUs). Remote SMPP clients emit request PDUs to which the Gateway SMPP server responds. The Gateway SMPP server operates synchronously. It completes the response to a request PDU before it processes the next request PDU from the connected remote SMPP client.

Table 27–19 lists the request PDUs the Gateway SMPP server handles, and specifies the Gateway SMPP server's response.

Table 27–19 SMPP Server Protocol Data Units

Request	SMPP Server Response
BIND_TRANSMITTER BIND_RECEIVER UNBIND	Responded to with the appropriate response PDU. Authentication credentials are ignored.
OUTBIND	Gateway SMPP server sends back a BIND_RECEIVER PDU. Authentication credentials presented are ignored.
SUBMIT_SMDATA_SM	Attempts to match the destination SMS address with either a unique SMS source address or the SELECT_RE setting of a Gateway profile. If neither is matched, the PDU is rejected with an ESME_RINVDSTADR error.
DELIVER_SM	Attempts to find either the destination SMS address or the receipted message ID in the historical record. If neither is matched, returns the error ESME_RINVMSGID .
BIND_RECEIVER	Not supported. Returns a GENERIC_NAK PDU with an ESME_RINVCMDID error.
SUBMIT_MULTI	Not supported. Returns a GENERIC_NAK PDU with an ESME_RINVCMDID error.
REPLACE_SM	Not supported. Returns a GENERIC_NAK PDU with an ESME_RINVCMDID error.
CANCEL_SM	Not supported. Returns a GENERIC_NAK PDU with an ESME_RINVCMDID error.
QUERY_SM	Not supported. Returns a GENERIC_NAK PDU with an ESME_RINVCMDID error.
QUERY_LAST_MSGS	Not supported. Returns a GENERIC_NAK PDU with an ESME_RINVCMDID error.
QUERY_MSG_DETAILS	Not supported. Returns a GENERIC_NAK PDU with an ESME_RINVCMDID error.
ENQUIRE_LINK	Returns ENQUIRE_LINK_RESP PDU.
ALERT_NOTIFICATION	Accepted but ignored.

SMS Reply and Notification Handling

The SMS Gateway Server maintains a historical record of each SMS message relayed through its SMPP relays. The need to use historical data arises from the fact that when submitting an email message to SMS it is generally not possible to convert the email address of the message's originator to an SMS source address. Since any SMS replies and notifications will be directed to this SMS source address, a problem arises. This is resolved by using automatically generated, unique SMS source addresses in relayed messages. The remote SMSCs are then configured to route these SMS source addresses back to the Gateway SMPP server.

The historical data is represented as an in-memory hash table of message IDs and generated, unique SMS source addresses. This data along with the associated email origination data are also stored on disk. The disk based storage is a series of files, each file representing **HASH_FILE_ROLLOVER_PERIOD** seconds of transactions (the default is 30 minutes). Each file is retained for **RECORD_LIFETIME** seconds (the default is 3 days).

Each record has three components:

- Email origination data (such as, envelope From: and To: addresses). This data is supplied by the MTA SMS channel when it submits a message.
- The unique SMS source address generated by the SMPP relay and inserted into the relayed SMS message.
- The resulting receipted message ID returned by the remote SMSC's SMPP server when it accepts a submission.

Routing Process for SMS Replies

The Gateway SMPP relays and servers use historical records to handle SMS replies, notifications and mobile originated messages. When an SMS message is presented to the SMPP relay or server, the following routing process is followed:

1. The SMS destination address is compared against the historical record to see if there is a matching, unique SMS source address that the SMPP relay previously generated. If a match is found, see Step 6. .
2. If there is no match, but the message is an SMS notification (SMPP **DELIVER_SM** PDU), then the receipted message ID, if present, is compared against the historical record. If a match is found, go to Step 8. The SMS Gateway Server actually allows these to be presented to either the SMPP relay or SMPP server.
3. If there is no match, then the destination SMS address is compared against the **SELECT_RE** option expressions for each configured gateway profile. If a match is found, then go to Step 9.
4. If there is no match and the SMS message was presented to the Gateway SMPP relay, then the message is relayed to the remote SMPP server.
5. If there is no match and the SMS message was presented to the Gateway SMPP server, then the message is determined to be an invalid message and an error response is returned in the SMPP response PDU. For email to SMS, a Non Delivery Notification (NDN) is eventually generated.
6. If a matching unique SMS source address was found, then the SMS message is further inspected to see if it is a reply or a notification message. To be a notification message it must be a **SUBMIT_SM** PDU with a receipted message ID. Otherwise, it is considered to be a reply.
7. If it is a reply then the SMS message is converted to an email message using the origination email information from the historical record.
8. If it is a notification, then the SMS message is converted to an email Delivery Status Notification (DSN) as per RFC 1892-1894. Note that the ESMTP **NOTIFY** flags (RFC 1891) of the original email message will be honored (For example, if the SMS message is a "success" DSN but the original email message requested only "failure" notifications, then the SMS notification will be discarded).
9. If the destination SMS address matches a **SELECT_RE** option in a configured gateway profile, then the SMS message is treated as a mobile originated message and converted back to email message as per the **PARSE_RE_n** rules for that

gateway profile. If the conversion fails, then the SMS message is invalid and an error response is returned.

SMS Gateway Server Configuration

This section gives directions on how to set up the SMS Gateway Server for both email-to-mobile and mobile-to-email functionality. This section includes the following topics:

- [Setting Up Bidirectional SMS Routing](#)
- [Enabling and Disabling the SMS Gateway Server](#)
- [Starting and Stopping the SMS Gateway Server](#)
- [SMS Gateway Server Configuration File](#)
- [Configuring Email-To-Mobile on the Gateway Server](#)
- [Configuring Mobile-to-Email Operation](#)
- [Configuration Options](#)
- [Global Options](#)
- [SMPP Relay Options](#)
- [SMPP Server Options](#)
- [Gateway Profile Options](#)
- [Configuration Example for Two-Way SMS](#)

Setting Up Bidirectional SMS Routing

The recommended way to set up bidirectional email and SMS routing between the MTA and SMSC is a three step process:

- [Set the SMS Address Prefix](#)-- Choose an SMS address prefix. Any prefix may be used, so long as it is ten characters or less.
- [Set the Gateway Profile](#)-- Reserve the prefix for use with the SMS Gateway Server (by setting the gateway profile).
- [Configure the SMSC](#)-- Configure the SMSC to route SMS destination addresses to the SMS Gateway SMPP server that start with the prefix. Mobile originated email will have only the prefix. Replies and notifications will have the prefix followed by exactly ten decimal digits.

Set the SMS Address Prefix

The source SMS addresses generated by the MTA SMS channel should be set to match the selected SMS address prefix. Do this by setting the following:

- MTA SMS channel options:
USE_HEADER_FROM=0
DEFAULT_SOURCE_ADDRESS=*prefix*

The first setting causes the channel to not attempt to set the SMS source address from information contained in the email message. The second setting causes the SMS source address to be set (to the selected prefix) when it is not set from any other source.

- Recognize the prefix as an SMS destination address to accept and route to email. Do this by specifying the **SELECT_RE** gateway profile option as follows:

```
SELECT_RE=prefix
```

Set the Gateway Profile

The SMS Gateway Server's gateway profile should then be set to make all relayed SMS source addresses unique. This is the default setting but may be explicitly set by specifying the gateway profile option **MAKE_SOURCE_ADDRESSES_UNIQUE=1**. This will result in relayed SMS source addresses of the form:

```
prefixnnnnnnnnnn
```

where **nnnnnnnnnn** will be a unique, ten digit decimal number.

Configure the SMSC

Finally, the SMSC should be configured to route all SMS destination addresses matching the prefix (either just the prefix, or the prefix plus a ten digit number) to the SMS Gateway Server's SMPP server. The regular expression for such a routing will be similar to:

```
prefix([0-9]{10,10}){0,1}
```

where **prefix** is the value of **DEFAULT_SOURCE_ADDRESS**, **[0-9]** specifies the allowed values for the ten digit number, **{10, 10}** specifies that there will be a minimum of ten digits and a maximum of ten digits, and **{0, 1}** specifies that there can be zero or one of the 10-digit numbers.

Enabling and Disabling the SMS Gateway Server

- To enable the SMS Gateway Server, the option **sms_gateway.enable** must be set to the value **1**. Use the following **msconfig** command to set it:

```
msconfig set sms_gateway.enable 1
```

- To disable the gateway server, set **sms_gateway.enable** to the value **0**, using the following command:

```
msconfig sms_gateway.enable 0
```

Starting and Stopping the SMS Gateway Server

After the SMS Gateway Server is enabled, it may be started and stopped with the following commands:

```
start-msg sms
```

and

```
stop-msg sms
```

SMS Gateway Server Configuration File

In order to operate, the SMS Gateway Server requires a configuration file. The configuration file is a Unicode text file encoded using UTF-8. The file can be an ASCII text file. The name of the file must be:

```
installation-directory/config/sms_gateway.cnf
```

Each option setting in the file has the following format:

```
option-name=option-value
```

Options that are part of an option group appear in the following format:

```
[group-type=group-name]
option-name-1=option-value-1
option-name-2=option-value-2
...
option-name-n=option-value-n
```

Configuring Email-To-Mobile on the Gateway Server

To implement the email-to-mobile part of two-way SMS, you must configure the following:

- [A Gateway Profile](#)
- [An SMPP Relay](#)
- [An SMPP Server](#)

A Gateway Profile

To configure an email-to-mobile gateway profile, follow these steps:

To Configure an Email-to-mobile Gateway Profile

1. Add a gateway profile to the SMS Gateway Server configuration file. To add an option group, use the following format:

```
[GATEWAY_PROFILE=profile_name]
option-name-1=option-value-1
option-name-2=option-value-2a
...
option-name-n=option-value-n
```

The length of the gateway profile name, **profile_name** in the preceding format, must not exceed 11 bytes. The name must be the same as the name for the **GATEWAY_PROFILE** channel option in the SMS channel option file. The name is case insensitive. For a list of the valid channel options, see "[Available Options](#)."

2. Set the gateway profile options (for example, **SMSC_DEFAULT_CHARSET**) to match characteristics of the remote SMSC.
3. Set the other gateway profile options to match the SMS channel's email characteristics.

A complete description of gateway profile options, see "[Gateway Profile Options](#)."

4. Set the **CHANNEL** option.

Set its value to the name of the MTA SMS channel. When a notification is sent to email through the gateway, the resulting email message will be enqueued to the MTA using this channel name.

An SMPP Relay

To configure an SMPP Relay, complete the following steps:

To Configure an SMPP Relay

1. Add an SMPP relay instantiation (option group) to the SMS Gateway Server's configuration file. To add an option group, use the following format:

```
[SMPP_RELAY=relay_name]
option-name-1=option-value-1
option-name-2=option-value-2
...
option-name-n=option-value-n
```

Any name may be used for the relay's name. All that matters is that the name not be used for any other SMPP relay instantiation within the same configuration file.

2. Set the `LISTEN_PORT` option.

The value used for the SMS channel's `SMPP_PORT` option must match that used for the relay's `LISTEN_PORT` option. For the `LISTEN_PORT`, select a TCP port number which is not used by any other SMPP relay or server instantiation nor by any other server running on the same computer.

3. Set the `SERVER_HOST` option.

The relay's `SERVER_HOST` option should give the host name for the remote SMSC's SMPP server. An IP address may be used in place of a host name.

4. Set the `SERVER_PORT` option.

The relay's `SERVER_PORT` option should give the TCP port for the remote SMSC's SMPP server. For a complete description of all SMPP relay options, see "[SMPP Relay Options](#)."

An SMPP Server

To configure an SMPP server, complete the following steps:

To Configure an SMPP Server

1. Add an SMPP server instantiation (option group) to the SMS Gateway Server's configuration file. To add an option group, use the following format:

```
[SMPP_SERVER=server_name]
option-name-1=option-value-1
option-name-2=option-value-2...
option-name-n=option-value-n
```

Any name may be used for the server's name. All that matters is that the name not be used for any other SMPP server instantiation within the same configuration file.

2. Set `LISTEN_PORT` option. Select a TCP port number which is not being used by any other server or relay instantiation. Additionally, the port number must not be being used by any other server on the same computer. The remote SMSC needs to be configured to route notifications via SMPP to the SMS Gateway Server system using this TCP port. For a complete description of all SMPP server options, see "[SMPP Server Options](#)."

Configuring Mobile-to-Email Operation

To configure mobile-to-email functionality, two configuration steps must be performed:

- [Configure a Mobile-to-Email Gateway Profile](#)
- [Configure a Mobile-Email SMPP Server](#)

Note that multiple gateway profiles may use the same SMPP server instantiation. Indeed, the same SMPP server instantiation may be used for both email-to-mobile and mobile-to-email applications.

Configure a Mobile-to-Email Gateway Profile

For mobile origination, a gateway profile provides two key pieces of information: how to identify SMS messages intended for that profile and how to convert those messages to email messages. Note that this profile can be the same one used for email-to-mobile with the addition of the *SELECT_RE* option.

To configure the gateway profile, follow these steps:

To Configure the Gateway Profile

1. Add a gateway profile (option group) to the SMS Gateway Server's configuration file.

To add an option group, use the following format:

```
[GATEWAY_PROFILE=profile_name]
option-name-1=option-value-1
option-name-2=option-value-2
...
option-name-n=option-value-n
```

Any name of 11 characters or less may be used for the profile's name. All that matters is that it is not already used for another gateway profile within the same configuration file.

2. Set the **SELECT_RE** option must be specified for each gateway profile.

The value of this option is an ASCII regular expression with which to compare SMS destination addresses. If an SMS destination address matches the regular expression, then the SMS message is sent through the gateway to email using the characteristics described by the matching profile. It is important to note that it is possible to configure multiple gateway profiles which have overlapping sets of SMS addresses (for example, a profile which matches the address 000 and another which matches any other three-digit address). However, so doing should be avoided as an SMS message will be passed off to only one gateway profile: the first one which matches. And, moreover, the order in which they are compared is undefined.

3. Set the **CHANNEL** option.

Its value should be the name of the MTA's SMS channel. For a complete description of all mobile origination options, see "[Gateway Profile Options](#)."

Configure a Mobile-Email SMPP Server

Adding an SMPP server is the same as for the email-to-mobile SMPP server (see "[An SMPP Server](#)").

The remote SMSC needs to be configured to route SMS traffic to the gateway SMPP server. To do this, the SMS destination address used by the SMSC to route mobile-to-email traffic should be the value set for the gateway profile option **SELECT_RE**.

For example, if the SMS address 000 is to be used for mobile-to-email traffic, then the SMSC needs to be configured to route traffic for the SMS destination address 000 to the gateway SMPP server. The gateway profile should use the option setting **SELECT_RE=000**.

Configuration Options

The SMS Gateway Server configuration file options are detailed in this section. The tables that follow list all the available configuration options with a brief description of each. There is a table each for global options, SMPP relay options, SMPP server options, and SMS Gateway Server profile options.

In the subsections that follow, complete descriptions are given for all the available configuration options. The subsections are:

- [Global Options](#)
 - Global options must be placed at the top of the configuration file, before any option groups. The remaining options must appear within option groups.
- [SMPP Relay Options](#)
- [SMPP Server Options](#)
- [Gateway Profile Options](#)

Global Options

The SMS Gateway Server presently has three categories of global options:

- [Thread Tuning Options](#)
- [Historical Data Tuning](#)
- [Miscellaneous](#)

All global options must be specified at the top of the configuration file, before any option groups are specified. [Table 27–20](#) lists all global configuration options.

Table 27–20 Global Options

Option	Default	Description
DEBUG	6	Selects the types of diagnostic output generated.
HISTORY_FILE_DIRECTORY	no default value	Absolute directory path for files of historical data.
HISTORY_FILE_MODE	0770	Permissions for files of historical data.
HISTORY_FILE_ROLLOVER_PERIOD	30 mins	Maximum length of time to write to the same file of historical data.
LISTEN_CONNECTION_MAX	10,000	Maximum number of concurrent inbound connections across all SMPP relay and server instantiations.
RECORD_LIFETIME	3 days	Lifetime of a record in the historical data archive.
THREAD_COUNT_INITIAL	10 threads	Initial number of worker threads.
THREAD_COUNT_MAXIMUM	50 threads	Maximum number of worker threads.
THREAD_STACK_SIZE	64 Kb	Stack size for each worker thread.

Thread Tuning Options

Each inbound TCP connection represents an SMPP session. The processing for a session is handled by a worker thread from a pool of threads. When the session processing needs to wait for completion of an I/O request, the worker thread parks the session and is given other work to perform. When the I/O request completes, the session is resumed by an available worker thread from the pool.

The following options allow for tuning of this pool of worker thread processes: "THREAD_COUNT_INITIAL," "THREAD_COUNT_MAXIMUM," "THREAD_STACK_SIZE."

THREAD_COUNT_INITIAL

(*integer, > 0*)_ Number of threads to initially create for the pool of worker threads. This count does not include the dedicated threads used to manage the in-memory historical data (2 threads) nor the dedicated threads used to listen for incoming TCP connections (one thread per TCP port/interface address pair the SMS Gateway Server listens on). The default value is for **THREAD_COUNT_INITIAL** is 10 threads.

THREAD_COUNT_MAXIMUM

(*integer, >= THREAD_COUNT_INITIAL*) Maximum number of threads to allow for the pool of worker threads. The default value is 50 threads.

THREAD_STACK_SIZE

(*integer, > 0*)_ Stack size in bytes for each worker thread in the pool of worker threads. The default value is 65,536 bytes (64 Kb).

Historical Data Tuning

When an SMS message is relayed, the message ID generated by the receiving, remote SMPP server is saved in an in-memory hash table. Along with this message ID, information about the original email message is also saved. Should that message ID subsequently be referenced by an SMS notification, this information may be retrieved. The retrieved information can then be used to send the SMS notification to the appropriate email recipient.

The in-memory hash table is backed to disk by a dedicated thread. The resulting disk files are referred to as "history files". These history files serve two purposes: to save, in nonvolatile form, the data necessary to restore the in-memory hash table after a restart of the SMS Gateway Server, and to conserve virtual memory by storing potentially lengthy data on disk. Each history file is only written to for **HASH_FILE_ROLLOVER_PERIOD** seconds after which time it is closed and a new history file created. When a history file exceeds an age of **RECORD_LIFETIME** seconds, it is deleted from disk.

The following options allow for tuning historical files: "HISTORY_FILE_DIRECTORY," "HISTORY_FILE_MODE," "HISTORY_FILE_ROLLOVER_PERIOD," "RECORD_LIFETIME."

HISTORY_FILE_DIRECTORY

(*string, absolute directory path*) Absolute path to the directory to which to write the history files. The directory path will be created if it does not exist. The default value for this option is:

MessagingServer_home/data/sms_gateway_cache/

The directory used should be on a reasonably fast disk system and have more than sufficient free space for the anticipated storage; see "SMS Gateway Server Storage Requirements" to change this option to a more appropriate value.

HISTORY_FILE_MODE

(*integer, octal value*) File permissions to associated with the history files. By default, a value of 0770 (octal) is used.

HISTORY_FILE_ROLLOVER_PERIOD

(*integer, seconds*) The current history file is closed and a new one created every **HASH_FILE_ROLLOVER_PERIOD** seconds. By default, a value of 1800 seconds (30 minutes) is used.

RECORD_LIFETIME

(*integer, seconds > 0*) Lifetime in seconds of a historical record. Records older than this lifetime will be purged from memory; history files older than this lifetime will be deleted from disk. By default, a value of 259,200 seconds (3 days) is used. Records stored in memory are purged in sweeps by a thread dedicated to managing the in-memory data. These sweeps occur every **HASH_FILE_ROLLOVER_PERIOD** seconds. Files on disk are purged when it becomes necessary to open a new history file.

Miscellaneous

This section describes the miscellaneous options.

DEBUG

(*integer, bitmask*) Enable debug output. The default value is 6 which selects warning and error messages.

Table 27–21 defines the bit values of the DEBUG bitmask.

Table 27–21 *DEBUG Bitmask*

Bit	Value	Descriptions
0-31	-1	Extremely verbose output.
0	1	Informational messages.
1	2	Warning messages.
3	4	Error messages.
3	8	Subroutine call tracing.
4	16	Hash table diagnostics.
5	32	I/O diagnostics, receive.
6	64	I/O diagnostics, transmit.
7	128	SMS to email conversion diagnostics (mobile originate and SMS notification).
8	256	PDU diagnostics, header data.
9	512	PDU diagnostics, body data.
10	1024	PDU diagnostics, type-length-value data.
11	2048	Option processing; sends all option settings to the log file.

LISTEN_CONNECTION_MAX

(*integer, >= 0*) The maximum number of concurrent, inbound TCP connections to allow across all SMPP relay and server instantiations. A value of 0 (zero) indicates that there is no global limit on the number of connections. There may, however, be per relay or server limits imposed by a given relay or server instantiation. Default: 10,000

SMPP Relay Options

The SMS Gateway Server can have multiple instantiations of its SMPP relay, each with different characteristics chief of which will be the TCP port and interface listened on. Put differently, for each network interface and TCP port pair the SMPP relay listens on, distinct characteristics may be ascribed. These characteristics are specified using the options described in this section.

Each instantiation should be placed within an option group of the form:

```
[SMPP_RELAY=relay-name]
option-name-1=option-value-1
option-name-2=option-value-2
...
option-name-n=option-value-n
```

The string **relay-name** merely serves to differentiate this instantiation from other instantiations.

Table 27–22 lists the SMPP relay configuration options.

Table 27–22 SMPP Relay Options

Options	Default	Description
<code>LISTEN_BACKLOG</code>	255	Connection backlog for inbound SMPP client connections.
<code>LISTEN_CONNECTION_MAX</code>	no default value	Maximum number of concurrent inbound connections.
<code>LISTEN_INTERFACE_ADDRESS</code>	no default value	Network interface for inbound SMPP client connections.
<code>LISTEN_PORT</code>	no default value	TCP port for inbound SMPP client connections.
<code>LISTEN_RECEIVE_TIMEOUT</code>	600 s	Read timeout for inbound connections from SMPP clients.
<code>LISTEN_TRANSMIT_TIMEOUT</code>	120 s	Write timeout for inbound connections from SMPP clients.
<code>MAKE_SOURCE_ADDRESSES_UNIQUE</code>	1	Make relayed SMS source addresses unique and able to be replied to.
<code>SERVER_HOST</code>	no default value	Host name or IP address of the SMPP server to relay to.
<code>SERVER_PORT</code>	no default value	TCP port of the SMPP server to relay to.
<code>SERVER_RECEIVE_TIMEOUT</code>	600 s	Read timeout for outbound SMPP server connections.
<code>SERVER_TRANSMIT_TIMEOUT</code>	120 s	Write timeout for outbound SMPP server connections.

LISTEN_BACKLOG

(*integer, in [0,255]*) Connection backlog allowed by the TCP stack for inbound SMPP client connections. The default value is 255.

LISTEN_CONNECTION_MAX

(*integer, >= 0*) The maximum number of concurrent, inbound TCP connections to allow for this SMPP relay instantiation. Note that this value will be ignored if it exceeds the global `LISTEN_CONNECTION_MAX` setting.

LISTEN_INTERFACE_ADDRESS

(*string, "INADDR_ANY" or dotted decimal IP address*) The IP address of the network interface to listen to for inbound SMPP client connections. May be either the string "INADDR_ANY" (all available interfaces) or an IP address in dotted decimal form. (For example, 193.168.100.1) The default value is "INADDR_ANY". Clustered HA configurations will need to set this value to correspond to the HA logical IP address.

LISTEN_PORT

(*integer, TCP port number*) TCP port to bind to for accepting inbound SMPP client connections. Specification of this option is mandatory; there is no default value for this option. Note also that there is no Internet Assigned Numbers Authority (IANA) assignment for this service.

LISTEN_RECEIVE_TIMEOUT

(*integer, seconds > 0*) Timeout to allow when waiting to read data from an SMPP client. The default value is 600 seconds (10 minutes).

LISTEN_TRANSMIT_TIMEOUT

(*integer, seconds > 0*) Timeout to allow when sending data to an SMPP client. The default value is 120 seconds (2 minutes).

MAKE_SOURCE_ADDRESSES_UNIQUE

(*0 or 1*) By default, the SMPP relay will append to each SMS source address a unique, ten digit string. The resulting SMS source address is then saved along with the other historical data. The result is a unique SMS address which may then be replied to by SMS users. The SMPP server will detect this address when used as an SMS destination address and will then send the SMS message to the correct email originator.

To disable this generating of unique SMS source addresses (for one-way SMS), specify a value of 0 (zero) for this option.

SERVER_HOST

(*string, TCP hostname or dotted decimal IP address*) SMPP server to relay SMPP client traffic to. Either a hostname or IP address may be specified. Specification of this option is mandatory; there is no default value for this option.

SERVER_PORT

(*integer, TCP port number*) TCP port for the remote SMPP server to which to relay. Specification of this option is mandatory; there is no default value for this option. There is no IANA assignment for this service; do not confuse with the IANA assignment for SNPP.

SERVER_RECEIVE_TIMEOUT

(*integer, seconds > 0*) Timeout to allow when waiting to read data from the SMPP server. The default value is 600 seconds (10 minutes).

SERVER_TRANSMIT_TIMEOUT

(*integer, seconds > 0*) Timeout to allow when sending data to the SMPP server. The default value is 120 seconds (2 minutes).

SMPP Server Options

The SMS Gateway Server can have multiple instantiations of its SMPP server, each with different characteristics chief of which will be the TCP port and interface listened on. Put differently, for each network interface and TCP port pair the SMPP server listens on, distinct characteristics may be ascribed. These characteristics are specified using the options described in this section.

Each instantiation should be placed within an option group of the form:

```
[SMPP_SERVER=server-name]
option-value-1=option-value-1
option-value-2=option-value-2
...
option-name-n=option-value-n
```

The string **server-name** merely serves to differentiate the instantiation from other instantiations.

Table 27–23 lists the SMPP server configuration options.

Table 27–23 SMPP Server Options

Options	Default	Description
<code>LISTEN_BACKLOG</code>	255	Connection backlog for inbound SMPP server connections
<code>LISTEN_CONNECTION_MAX</code>	no default value	Maximum number of concurrent inbound connections.
<code>LISTEN_INTERFACE_ADDRESS</code>	no default value	Network interface for inbound SMPP server connections.
<code>LISTEN_PORT</code>	no default value	TCP port for inbound SMPP server connections.
<code>LISTEN_RECEIVE_TIMEOUT</code>	600 s	Read timeout for inbound SMPP server connections.
<code>LISTEN_TRANSMIT_TIMEOUT</code>	120 s	Write timeout for inbound SMPP server connections.

LISTEN_BACKLOG

(*integer in[0,255]*) Connection backlog allowed by the TCP stack for inbound SMPP client connections. The default value is 255.

LISTEN_CONNECTION_MAX

(*integer >= 0*) The maximum number of concurrent, inbound TCP connections to allow for this SMPP server instantiation. Note that this value will be ignored if it exceeds the global `LISTEN_CONNECTION_MAX` setting.

LISTEN_INTERFACE_ADDRESS

(*string, "INADDR_ANY" or dotted decimal IP address*) The IP address of the network interface to listen to for inbound SMPP client connections on. May be either the string "INADDR_ANY" (all available interfaces) or an IP address in dotted decimal form. (For example, 193.168.100.1.) The default value is "INADDR_ANY".

LISTEN_PORT

(*integer, TCP port number*) TCP port to bind to for accepting inbound SMPP client connections. Specification of this option is mandatory; there is no default value for this option. Note that there is no IANA assignment for this service.

LISTEN_RECEIVE_TIMEOUT

(*integer, seconds > 0*) Timeout to allow when waiting to read data from an SMPP client. The default value is 600 seconds (10 minutes).

LISTEN_TRANSMIT_TIMEOUT

(*integer, seconds > 0*) Timeout to allow when sending data to an SMPP client. The default value is 120 seconds (2 minutes).

Gateway Profile Options

There may be zero or more gateway profiles. In the SMS Gateway Sever's configuration file, each gateway profile is declared within an option group in the following format:

```
[GATEWAY_PROFILE=profile-name]
option-name-1=option-value-1
option-name-2=option-value-2
...
option-name-n=option-value-n
```

The string **profile-name** merely serves to differentiate the profile from other origination profiles.

Table 27–24 lists the SMS Gateway Server profile options.

Table 27–24 SMS Gateway Server Profile Options

Options	Default	Description
CHANNEL	sms	Channel to enqueue message as.
EMAIL_BODY_CHARSET	US-ASCII	Character set for email message bodies.
EMAIL_HEADER_CHARSET	US-ASCII	Character set for email message headers.
FROM_DOMAIN	no default value	Domain name for routing email back to SMS.
PARSE_RE_0, PARSE_RE_1, ..., PARSE_RE_9	no default value	Regular expressions for parsing SMS message text.
PROFILE	GSM	SMS profile to operate under: GSM, TDMA, or CDMA.
SELECT_RE	no default value	Regular expression for selecting the plugin.
SMSC_DEFAULT_CHARSET	US-ASCII	SMSC's default character set.
USE_SMS_PRIORITY	0	Gateway SMS priority flags to email.
USE_SMS_PRIVACY	0	Gateway SMS privacy indicators to email.

CHANNEL

(*string, 1-40 characters*) Name of the MTA channel used to enqueue email messages. If not specified, then "sms" is assumed. The specified channel must be defined in the MTA's configuration.

EMAIL_BODY_CHARSET

(*string, character set name*) The character set to translate SMS text to prior to insertion into an email message's body. If necessary, the translated text will be MIME encoded. The default value is US-ASCII. If the SMS message contains glyphs not available in the charset, they will be converted to mnemonic characters, which may or may not be meaningful to the recipient.

A list of the character sets known to the MTA may be found in the following file:

installation-directory/config/charsets.txt

EMAIL_HEADER_CHARSET

(*string, character set name*) The character set to translate SMS text to prior to insertion into an RFC 822 **Subject: header** line. If necessary, the translated string will be MIME encoded. The default value is US-ASCII. If the SMS message contains glyphs not available in the charset, they will be converted to mnemonic characters, which may or may not be meaningful to the recipient

FROM_DOMAIN

(*string, IP host name, 1-64 characters*) Domain name to append to SMS source addresses when constructing envelope *From:* addresses for email messages. The name specified should be the correct name for routing email back to SMS. (For example, the host name associated with the MTA SMS channel.) If not specified, then the official host name of the channel specified with the **CHANNEL** option will be used.

PARSE_RE_0, PARSE_RE_1, ..., PARSE_RE_9

(*string, UTF-8 regular expression*) For mobile origination of email, the gateway profile needs to extract a destination email address from the text of the SMS message. This is done by means of one or more POSIX-compliant regular expressions (REs). The text of the SMS message will be evaluated by each regular expression until either a match producing a destination email address is found or the list of regular expressions exhausted.

Note: Use of **PARSE_RE_*** and **ROUTE_TO** options are mutually exclusive. Use of both in the same gateway profile is a configuration error.

Each regular expression must be POSIX compliant and encoded in the UTF-8 character set. The regular expressions must output as string 0 the destination address. They may optionally output text to use in a **Subject:** header line as string 1, and text to use in the message body as string 2. Any text not "consumed" by the regular expression will also be used in the message body, following any text output as string 2.

The regular expressions will be tried in the order **PARSE_RE_0**, **PARSE_RE_1**, ..., up to **PARSE_RE_9**. If no regular expressions are specified, then the following default regular expression is used:

```
[ \t]*([^\( ]*)[ \t]*(?:\(((\[^\)]*)\))?[ \t]*(.*)
```

This default regular expression breaks into the following components:

```
[ \t]*
```

Ignore leading white space characters (*SPACE* and *TAB*).

```
([^\( ]*)
```

Destination email address. This is the first reported string.

```
[ \t]*
```

Ignore white space characters.

```
(?:\(((\[^\)]*)\))?
```

Optional subject text enclosed in parentheses. This is the second reported string. The leading `?` causes the outer parentheses to not report a string. They are being used merely for grouping their contents together into a single RE for the trailing `?`. The trailing `?` causes this RE component to match only zero or one time and is equivalent to the expression:

```
{0,1}
```

```
[ \t]*
```

Ignore white space characters.

```
(.*)
```

Remaining text to message body. This is the third reported string.

For example, with the above regular expression, the sample SMS message:

dan@example.com(Testing)This is a test

yields the email message:

```
To: dan@example.com
Subject: Testing
```

This is a test

As a second example, the SMS message:

sue@example.com This is another test

would yield:

To: sue@example.com

This is another test

Note that the SMS message, prior to evaluation with these regular expressions, will be translated to the UTF-16 encoding of Unicode. The translated text is then evaluated with the regular expressions which were previously converted from UTF-8 to UTF-16. The results of the evaluation are then translated to US-ASCII for the destination email address, `EMAIL_HEADER_CHARSET` for the **Subject:** text, if any, and `EMAIL_BODY_CHARSET` for the message body, if any.

PROFILE

(*string, "GSM", "TDMA", or "CDMA"*) SMS profile to assume. Presently this information is only used to map SMS priority flags to RFC 822 *Priority:* header lines. Consequently, this option has no effect when `USE_SMS_PRIORITY=0` which is the default setting for that option.

SELECT_RE

(*string, US-ASCII regular expression*) A US-ASCII POSIX-compliant regular expression to compare against each SMS message's SMS destination address. If an SMS message's destination address matches this RE, then the SMS message will be sent through the gateway to email in accord with this gateway profile.

Note that since an SMS message's destination address is specified in the US-ASCII character set, this regular expression must also be expressed in US-ASCII.

SMSC_DEFAULT_CHARSET

(*string, character set name*) The name of the default character set used by the remote SMSC. The two common choices for this option are US-ASCII and UTF-16-BE (USC2). If not specified, US-ASCII is assumed.

USE_SMS_PRIORITY

(*integer, 0 or 1*) By default (with **USE_SMS_PRIORITY=0**), priority flags in SMS messages are ignored and not sent with the email messages. To have the priority flags passed with the email, specify **USE_SMS_PRIORITY=1**. When passed with the email, the mapping from SMS to email is as shown in [Table 27–25](#):

Table 27–25 Priority Flag Mapping from SMS to Email

SMS Profile	SMS Priority Flag	Email Priority: Header Line
GSM	0 (Non-priority)1, 2, 3 (Priority)	No header line (implies Normal) Urgent .
TDMA	0 (Bulk)1 (Normal)2 (Urgent)3 (Very Urgent)	Nonurgent No header line (implies Normal) UrgentUrgent .
CDMA	0 (Normal)1 (Interactive)2 (Urgent)3 (Emergency)	No header line (implies Normal) UrgentUrgentUrgent .

Note that the email **Priority:** header line values are **Nonurgent**, **Normal**, and **Urgent**.

USE_SMS_PRIVACY

(*integer, 0 or 1*) By default (with **USE_SMS_PRIVACY=0**), SMS privacy indications are ignored and not sent with the email messages. To have this information passed with the email, specify **USE_SMS_PRIVACY=1**. When passed along with email, the mapping from SMS to email is as shown in [Table 27–26](#):

Table 27–26 Privacy Flags Mapping from SMS to Email

SMS Privacy Flag	Email Sensitivity: Header Line
0 (Not restricted)	No header line.
1 (Restricted)	Personal .
2 (Confidential)	Private .
3 (Secret)	Company-confidential .

Note that the values of the email *Sensitivity:* header line are *Personal*, *Private*, and *Company-confidential*.

Configuration Example for Two-Way SMS

This section shows a configuration example for a two-way SMS. It provides assumptions on behavior and further assumptions and assignments. The section also describes an SMS channel configuration, SMS channel options, an SMS gateway server configuration, and additional sms_option file settings.

Assumptions on Behavior

For the sake of this example, let us assume that the following behavior is desired:

- Email messages addressed to
sms-id@sms.domain.com

are to be sent to the SMS address

sms-id

and given a unique SMS source address in the range *000nnnnnnnnnnnn*.

- Mobile SMS messages addressed to the SMS address *000* are to be sent through the gateway to email with the email address extracted from the start of the SMS message text.

For example, if the SMS message text is:

jdoe@domain.com Interested in a movie?

then the message "Interested in a movie?" is to be sent to *jdoe@domain.com*.

- SMS notifications sent to **000nnnnnnnnnn** are to be sent through the gateway to email and directed to the originator of the message being received.

In order to bring about this behavior, the following assumptions and assignments are made

Further Assumptions and Assignments

- The MTA's SMS channel uses the domain name *sms.domain.com*.
- The SMS Gateway Server runs on the host *gateway.domain.com* and uses:
 - TCP port 503 for its SMPP relay
 - TCP port 504 for its SMPP server
- The remote SMSC's SMPP server runs on the host **smpp.domain.com** and listens on TCP port 377.
- The remote SMSC's default character set is UCS2 (aka, UTF-16).

SMS Channel Configuration

To effect the above behavior, the following SMS channel configuration may be used (add these lines to the bottom of the channel blocks after running **msconfig edit channels**):

```
(blank line)
sms
sms.domain.com
```

SMS Channel Options

The channel's options would then be set as follows:

```
msconfig
msconfig> set channel:mway_sms.options.smpp_server gateway.domain.com
msconfig# set channel:mway_sms.options.smpp_port 503
msconfig# set channel:mway_sms.options.use_header_from 0
msconfig# set channel:mway_sms.options.default_source_address 000
msconfig# set channel:mway_sms.options.gateway_profile sms1
msconfig# set channel:mway_sms.options.smsc_default_charset UCS2
```

SMS Gateway Server Configuration

Finally, the Gateway Server configuration file, **sms_gateway.cnf**, should look like the following:

```
HISTORY_FILE_DIRECTORY=/sms_gateway_cache/
```

```
[SMPP_RELAY=relay1]
LISTEN_PORT=503SERVER_HOST=smpp.domain.com
SERVER_PORT=377
```

```
[SMPP_SERVER=server1]
LISTEN_PORT=504
```

```
[GATEWAY_PROFILE=sms1]
SELECT_RE=000 ([0-9]{10,10}){0,1}
SMSC_DEFAULT_CHARSET=UCS2
```

Testing this Configuration

If you do not have an SMSC to test on, you may want to perform some loopback tests. With some additional sms channel option settings, some simple loop back tests may be performed for the above configuration.

Additional sms_option File Settings

The additional settings for the sms channel options are:

```
msconfig
msconfig> set channel:mway_sms.options.FROM_FORMAT
msconfig# set channel:mway_sms.options.SUBJECT_FORMAT
msconfig# set channel:mway_sms.options.CONTENT_PREFIX
```

Without these settings, an email containing:

```
user@domain.com (Sample subject) Sample text
```

would get converted into the SMS message:

```
From:user@domain.com Subject:Sample Subject Msg:Sample text
```

That, in turn, would not be in the format expected by the mobile-to-email code, which wants to see:

```
user@domain.com (Sample subject) Sample text
```

Hence the need (for loopback testing) to specify empty strings for the *FROM_FORMAT*, *SUBJECT_FORMAT*, and *CONTENT_PREFIX* options.

Performing the Loopback test

Send test email messages addressed to *000@sms.domain.com*, such as:

```
user@domain.com (Test message) This is a test message which should loop back
```

The result is that this email message should be routed back to the email recipient *user@domain.com*. Be sure to have added *sms.domain.com* to your DNS or host tables for the test.

SMS Gateway Server Storage Requirements

To determine the amount of resources you will need for the SMS Gateway Server, use the numbers you generate from the requirements in [Table 27–27](#) along with your expected number of relayed messages per second and the **RECORD_LIFETIME** setting.

[Table 27–27](#) covers the requirements for the historical records, the SMPP relay, and SMPP server.

Table 27–27 SMS Gateway Server Storage Requirements

Component	Requirements
In-memory historical record	<p>Each relayed message requires $33+m+s$ bytes of virtual memory, where m is the length of the message's SMS message ID ($1 \leq m \leq 64$) and s is the length of the message's SMS source address ($1 \leq s \leq 20$).</p> <p>When <code>MAKE_SOURCE_ADDRESS_UNIQUE=0</code>, then only $16+m$ bytes are used. For 64-bit operating systems, $49+m+s$ bytes of virtual memory are consumed per record [$24+m$ when <code>MAKE_SOURCE_ADDRESS_UNIQUE=0</code>].</p> <p>Note also, that the heap allocator may actually allocate larger size pieces of virtual memory for each record. The maximum number of records is 43 billion:</p> $(2^{32}-1)$ <p>For fewer than 16.8 million records:</p> (2^{24}) <p>the hash table consumes approximately 16 Mb. For fewer than 67.1 million records:</p> (2^{26}) <p>the hash table consumes approximately 64 Mb; for more than 67.1 million records, the hash table consumes approximately 256 Mb.</p> <p>Double the memory consumptions for 64 bit operating systems.</p> <p>These consumptions are in addition to the memory consumption required for each record itself.</p>
On-disk historical record	<p>Each relayed message requires on average the following number of bytes:</p> $81+m+2s+3a+S+2i$ <p>where:</p> <ul style="list-style-type: none"> ■ m is the average length of SMS message IDs, and $1 \leq m \leq 64$ ■ s is the average length of SMS source addresses, and $1 \leq s \leq 20$ ■ a is the average length of email addresses, and $3 \leq a \leq 129$ ■ S is the average length of Subject: header lines, and $0 \leq S \leq 80$ $78+m+3a+S+2i$
SMPP relay	Each relayed SMPP session consumes two TCP sockets: one with the local SMPP client and another with the remote SMPP server. Approximately 1 Kb of virtual memory is consumed per connection on 32 bit operating systems; 2 Kb on 64 bit operating systems.
SMPP server	Each incoming connection consumes a TCP socket. Approximately 1 Kb of virtual memory is consumed per connection on 32 bit operating systems; 2 Kb on 64 bit operating systems.

For instance, if on average 50 messages per second are expected to be relayed, SMS source addresses are 13 bytes long, SMS message IDs have a typical length of 12 bytes, email addresses 24 bytes, **Subject:** lines 40 bytes, email message and envelope IDs 40 bytes each, and historical data is retained for 7 days, then:

- There will be 30.24 million historical records to store, each on average 58 bytes in memory and 311 bytes long on disk;
- The in-memory consumption of the historical records will be about 1.70 Gb (1.63 Gb + 64 Mb); and
- The on-disk storage will be approximately 8.76 Gb.

While a sufficiency of disk may be supplied to handle any on disk requirements, the virtual memory requirement on a 32-bit machine will be a hard limit of approximately 2 Gb. To reduce the amount of virtual memory or disk storage required, use the `RECORD_LIFETIME` option to reduce the length of time records are retained.

SMS Configuration Examples

The following examples outline a couple of ways to configure one-way and two-way SMS:

- [Configuring Messaging Server for One-Way SMS](#)
- [Configuring Messaging Server for One-Way SMS](#)

Configuring Messaging Server for One-Way SMS

This example describes how to configure Oracle Communications Messaging Server for one-way SMS by using Unified Configuration recipes. You could also perform the same configuration manually by using the `msconfig` command, however, using recipes is faster and makes the task repeatable. You run one recipe on the system where you are adding the SMS channelsmpp. For more information on Unified Configuration recipes, see the discussion on recipe language in the *Messaging Server Reference* and "Using Recipes."

Use the "Example Recipe to Configure One-Way SMS with Messaging Server" to configure Messaging Server for one-way SMS.

1. Create the source filter file (`process_sms.filter`) with the Messaging Server unified configuration.

- a. To create a filter definition file for the source filter, enter a definition as in the following example:

```
require ["body", "imap4flags"];
if body :raw :contains "Action: failed"
{
  addflag "sms";
  addflag "smserror";
}
```

- b. Save the filter definition file.

2. Make a copy of the example recipe file and save it as `recipe.rcp`.
3. To run the recipe, type the following command:

```
cd MessagingServer_home/bin
msconfig run <recipe_name>
```

4. If running a compiled configuration, recompile the MTA configuration to capture the changes you have made:

```
cd MessagingServer_home/bin
imsimta cnbuild
```

5. Restart Messaging Server.

```
stop-msg
start-msg
```

Example Recipe to Configure One-Way SMS with Messaging Server

```
set_option("mta.enable_sieve_body", "1");

if_exists_channel("sms") {
delete_channel("sms");
}
if_exists_channel("process_sms") {
delete_channel("process_sms");
}

add_channel("sms", ["notificationchannel", "process_sms", "backoff", "PT10M PT20M
PT30M", "notices", "1", "official_host_name", "sms-handle", "options.smpp_server",
"127.0.0.1", "options.smpp_port", "8500", "options.default_source_address", "000",
"options.smsc_default_charset", "UTF-16-BE", "options.esme_password", "password",
"options.esme_system_id", "smppclient"]);

add_channel("process_sms", ["sourcefilter",
"file:/opt/sun/comms/messaging64/config/process_sms.filter", "official_host_name",
"process_sms-daemon"]);

#Setting rewrite rules
append_rewrites(["sms", "$U@sms-handle"]);
```

Configuring Messaging Server for Two-Way SMS

This example describes how to configure Oracle Communications Messaging Server for two-way SMS by using Unified Configuration recipes. You could also perform the same configuration manually by using the `msconfig` command, however, using recipes is faster and makes the task repeatable. You run one recipe on the system where you are adding the SMS channelsmpp. For more information on Unified Configuration recipes, see the discussion on recipe language in the *Messaging Server Reference* and "Using Recipes."

Use the following "Example Recipe to Configure Two-Way SMS with Messaging Server" to configure Messaging Server for two-way SMS.

1. Create the source filter file (`process_sms.filter`) with the Messaging Server Unified Configuration.

- a. To create a filter definition file for the source filter, enter a definition as in the following example. This is the source filter for the `process_sms` channel:

```
require ["body", "imap4flags"];
if body :raw :contains "Action: failed"
{
  addflag "sms";
  addflag "smserror";
}
```

- b. Save the process definition file.

2. Make a copy of the example recipe file and save it as `recipe.rcp`.
3. To run the recipe, type the following command:

```
cd MessagingServer_home/bin
msconfig run <recipe_name>
```

4. If running a compiled configuration, recompile the Messaging Server configuration to capture the changes you have made:

```
imsimta cnbuild
```

5. Restart Messaging Server.

```
cd MessagingServer_home/bin
stop-msg
start-msg
```

Example Recipe to Configure Two-Way SMS with Messaging Server

```

set_option("mta.enable_sieve_body", "1");
set_option("sms_gateway.enable", "1");

if exists_channel("sms") {
delete_channel("sms");
}
if exists_channel("process_sms") {
delete_channel("process_sms");
}

add_channel("sms", ["notificationchannel", "process_sms", "backoff", "PT10M
PT20M PT30M", "notices", "1", "sourcefilter",
"file://opt/sun/comms/messaging64/config/sms_channel.filter", "official_host_
name", "sms-handle", "options.smpp_server", "127.0.0.1", "options.smpp_port",
"8500", "options.gateway_profile", "GATEWAY", "options.default_source_
address", "000", "options.smsc_default_charset", "UTF-16-BE", "options.use_
header_from", "0", "options.esme_password", "password", "options.esme_system_
id", "smppclient"]);

add_channel("process_sms", ["sourcefilter",
"file://opt/sun/comms/messaging64/config/process_sms.filter", "official_host_
name", "process_sms-daemon"]);

Setting rewrite rules
append_rewrites(["sms", "$U@sms-handle"]);

if exists_group("sms_gateway.smpp_server:SMPPSERVER") {
delete_group("sms_gateway.smpp_server:SMPPSERVER");
}

add_group("sms_gateway.smpp_server:SMPPSERVER", ["tcp_ports", "4080", "server_
host", "127.0.0.1", "server_port", "950", "ESME_PASSWORD", "password", "ESME_
SYSTEM_ID", "smppclient"]);

if exists_group("sms_gateway.smpp_relay:SMPPRELAY") {
delete_group("sms_gateway.smpp_relay:SMPPRELAY");
delete_group("sms_gateway.smpp_relay:SMPP_RELAY");
}

add_group("sms_gateway.smpp_relay:SMPPRELAY", ["server_port", "950", "tcp_
ports", "8500", "server_host", "127.0.0.1"]);

if exists_group("sms_gateway.gateway_profile:GATEWAY") {
delete_group("sms_gateway.gateway_profile:GATEWAY");
}

add_group("sms_gateway.gateway_profile:GATEWAY", ["select_re", "([0-9]*)",
"mta_channel", "sms", "smsc_default_charset", "UTF-16-BE", "email_body_
charset", "UTF-8", "email_header_charset", "UTF-8", "text_to_subject", "1"]);

```

6. Set the JMQ options to create non-delivery failure notifications of any SMS messages. For more information, see ["Java Message Queue Notification Messages and Properties."](#) To set the following options, see ["Configure and Enable the Messaging Server jmqnotify Plugin"](#):

- NewMsg.enable -v 1
- jmqHost -v "127.0.0.1"
- jmqPort -v "7777"

-
- jmqUser -v "user"
 - jmqpwd -v "xxx"
 - DestinationType -v "queue"
 - jmqQueue -v "ucsms1"
 - Priority -v 3
 - ttl -v 1000
 - Persistent -v 1
 - maxheadersize -v 1024
 - noneinbox.enable -v 1
 - msgflags.enable -v 0
 - readmsg.enable -v 0

Using the iSchedule Channel to Handle iMIP Messages

Oracle Communications Messaging Server is capable of posting a calendar event received in an iMIP (iCalendar Message-Based Interoperability Protocol) message to Calendar Server by using the iSchedule protocol. This capability enables "internal" users to automatically process calendar invitations from "external" users. To enable this interoperability between calendaring systems, you configure a Messaging Server "iSchedule" channel to process the iMIP messages. For additional information, see the discussion on enabling the iSchedule channel to handle iMIP messages in the *Calendar Server System Administrators Guide*.

Topics:

- [Inviting Users on Internal and External Calendar Systems Background](#)
- [Message Server iMIP Configuration Overview](#)
- [Configuring the iSchedule Channel for iMIP Messages in Unified Configuration](#)
- [Configuring the iSchedule Channel in Legacy Configuration](#)
- [Troubleshooting the iSchedule Configuration](#)

Inviting Users on Internal and External Calendar Systems Background

Calendar Server meetings often have multiple invitees. These invitees can be both *internal* users, who reside on the same Calendar Server deployment, or *external* users, who reside either on a different Calendar Server deployment administered by a separate group, or on an outside calendaring system, such as Exchange, Google Calendar, and so on. For "internal" invitees, Calendar Server automatically adds the meeting request to their calendars (referred to as *implicit scheduling*) and also sends them email notification about the meeting request. All "external" invitees are sent an iMIP (iCalendar Message-Based Interoperability Protocol) email with the meeting request as an attachment. External invitees must manually process these messages to add the invite to their calendars.

Manually Accepting External Invitations

Meeting invitations from external organizers are sent to the user's mailbox. Mail clients, such as Outlook or Thunderbird, enable users to process these invitations and add the invitation to their calendar. How the invitation is added to the user's calendar depends on the specific mail client, but the invitation is not added until the user has manually read the email and accepted the meeting request.

Automatically Accepting External Invitations

You can configure your Calendar Server deployment to automatically process invitations coming from external calendar systems. To users, handling an external invite then appears just like an internal invite.

This capability involves an intermediary in the form of Messaging Server. You configure the Messaging Server MTA to process the calendar invite email (which is an iMIP message), extract the pertinent calendar information, then use the iSchedule protocol to add the invite to the attendee's calendar database. As a consequence, external event invitations automatically appear in the user's calendar without the need for a manual intervention, even when using a "non-calendar" aware client.

Once you have configured your deployment accordingly, users have a choice on how to process invitations. Users can either accept the "external" meeting invite directly from their calendar client (either desktop or mobile iOS CalDAV clients) or they can still accept it from their email client. That is, CalDAV clients now receive iMIP messages in their scheduling-inbox, and are able to process them just like regular CalDAV-based invitations and replies. Because the invitation is already in the user's calendar, invitation replies and cancel are also merged automatically. Thus, based upon the user accepting or rejecting the request, the calendar client merely has to update the attendee status in the invitation. That status change also enables Calendar Server to send a response to the organizer indicating the disposition of the meeting request. As the response is sent directly by Calendar Server, it does not matter how the user accepted the invitation (whether from a calendar client on a mobile device, desktop, or from an email client). Finally, because of the addition of meta data to the email message, the Convergence (web-based) client is able to display a scheduling-specific form to users that enables them to accept, decline, or indicate a "maybe" to meeting invitations directly from their email without having to switch to the calendar client.

Message Server iMIP Configuration Overview

A Messaging Server MTA channel (an "iSchedule" channel) handles automatic processing of external calendar invites by:

1. Intercepting incoming emails containing an iMIP message An iMIP message has an iCalendar attachment of type 'text/calendar' with a **method=<action>** option in the 'Content-Type:' header.
2. Injecting the corresponding iTIP message into the regular calendar server workflow
3. Adding meta-data (email X- headers) to the iMIP email before delivering it to its recipients
4. Posting the invitation request to a calendar iSchedule URL

The Calendar Server then consumes the invitation from the iSchedule URL just as it would have done if an external calendar server had posted an invitation to one of its users. On the Calendar Server side, an iSchedule database, which is a separate table from the Calendar Server database, acts as the global inbox and outbox for external invites.

Administering the iMIP configuration involves:

- Enabling or disabling iMIP messaging processing
- Configuring the iSchedule service URL
- Configuring the criteria for messages to be selected for processing

You should configure the iSchedule channel on the message store systems if you are not using LMTP. If you are using LMTP, configure the iSchedule on the MTAs.

For more information on the iCalendar Message-Based Interoperability Protocol, see RFC 6047 (<http://www.faqs.org/rfcs/rfc6047.html>). For more information on iCalendar Transport-independent Interoperability Protocol (iTIP), see RFC 5546 (<http://www.faqs.org/rfcs/rfc5546.html>).

Configuring the iSchedule Channel for iMIP Messages in Unified Configuration

You can use a Messaging Server Unified Configuration recipe to automate the configuration process or you can manually perform the necessary configuration. After completing the configuration, you also need to verify the Calendar Server configuration, as described in "Verifying the Calendar Server Configuration."

You do not need to perform any additional Convergence configuration for Convergence to automatically process invitations coming from external calendar systems. If you have configured Messaging Server and Calendar Server correctly, Convergence users see a UI form that they use to reply to external invites.

Topics in this section:

- [Using the iSchedule Recipe to Automate Configuring the iSchedule Channel in Unified Configuration](#)
- [Manually Configuring the iSchedule Channel in Unified Configuration](#)
- [Verifying the Calendar Server Configuration](#)
- [Modifying iSchedule Channel Options](#)

Using the iSchedule Recipe to Automate Configuring the iSchedule Channel in Unified Configuration

Unified Configuration provides a recipe language and some stock recipes to automate certain configuration tasks (see "Using Recipes"). To set up the iSchedule channel, you can use a recipe called `iSchedule.rcp`, which automatically sets up the channel definition, job controller configuration, channel options, Sieve rule, and CONVERSION mapping.

To use the `iSchedule.rcp` recipe:

1. Run the `msconfig` command with the recipe name.

```
/opt/sun/comms/messaging64/bin/msconfig run iSchedule.rcp
```

2. Respond to the prompts, for example:

```
HTTP URL for iSchedule server: http://host1.example.com:8080/dav/ischedule/
Destination channel for messages to check (<RET> if no more): ims-ms
```

Use the iSchedule URL and destination channels based on your deployment.

Note: Be sure to add the trailing forward slash (/) in the iSchedule URL, otherwise you will receive the error message "HTTP Error 401 Unauthorized."

3. If you are using a compiled configuration, recompile the configuration.

```
/opt/sun/comms/messaging64/bin/imsimta cbuild
```

4. Restart Messaging Server.

```
/opt/sun/comms/messaging64/bin/stop-msg
/opt/sun/comms/messaging64/bin/start-msg
```

5. Verify the Calendar Server configuration. See "[Verifying the Calendar Server Configuration](#)."

Manually Configuring the iSchedule Channel in Unified Configuration

The high-level steps to manually configure the iSchedule channel by using the **msconfig** command involve:

- Adding the channel
- Configuring the conversion mapping
- Specifying messages to be processed by iSchedule

To manually configure the iSchedule Channel, job controller master command, **include_conversiontag** MTA option, and conversion mapping:

1. Use the **msconfig** command in interactive mode to configure the iSchedule channel, the job controller master command for the channel, the **include_conversiontag** MTA option if you want to have a **TAG=** clause included in your conversion mapping probes, and the conversion mapping.

```
/opt/sun/comms/messaging64/bin/msconfig
msconfig> set channel:ischedule.official_host_name ischedule-daemon
msconfig# set channel:ischedule.options.handle-imip 1
msconfig# set channel:ischedule.options.ischedule-url
http://<host>:<port>/dav/ischedule/
msconfig# set instance.job_controller.channel_class:ischedule.master_command
IMTA_BIN:ischedule
msconfig# set mapping:conversion.rule
"IN-CHAN=ischedule;OUT-CHAN=*;TAG=*;CONVERT" NO
msconfig# set mapping:conversion.rule
"IN-CHAN=*;OUT-CHAN=*;TAG=ISCHEDULE;CONVERT" "YES,CHANNEL=ischedule"
msconfig# set include_conversiontag 2
msconfig# write
```

Use the host name and alternately the port for the Calendar Server configured for the iSchedule database.

Note: Be sure to add the trailing forward slash (/) in the **ischedule-url**, otherwise you will receive the error message "HTTP Error 401 Unauthorized."

2. Edit the filters block to specify messages to be processed by iSchedule.

```
msconfig> edit filter
```

The filter block appears in the editor that is the default configured for your login.

3. Add the following lines to create a filter that selects all the messages that have "text/calendar" MIME as an attachment:

```
require ["mime", "environment"];
if allof(environment :is "vnd.sun.destination-channel" ["ims-ms"],
```

```
header :mime :anychild :contenttype :is "content-type" "text/calendar",
NOT header :contains "X-Oracle-CS-iSchedule-Ignore" "Yes") {
addconversiontag "ISCHEDULE";
}
```

iMIP messages generated by Calendar Server contain a "X-Oracle-CS-iSchedule-Ignore: Yes" header to indicate that the event was already added to the user's calendar. So, the Sieve rule should ignore those iMIP messages by not tagging them with an ISCHEDULE conversion tag. Failing to do so results in an iSchedule post of the event that is already present in the user's calendar.

4. Write the configuration and exit the **msconfig** interactive mode.

```
msconfig# write
msconfig> exit
#
```

5. If you are using a compiled configuration, recompile the configuration.

```
/opt/sun/comms/messaging64/bin/imsimta cbuild
```

6. Restart Messaging Server.

```
/opt/sun/comms/messaging64/bin/stop-msg
/opt/sun/comms/messaging64/bin/start-msg
```

7. Verify the Calendar Server configuration. See "[Verifying the Calendar Server Configuration](#)."

Verifying the Calendar Server Configuration

To ensure that your Calendar Server configuration is setup properly, use the following steps to verify SMTP settings and iMIP email notifications. iMIP email notifications need to also be configured for internal users, that is, users on the same Calendar Server host. If necessary, restart the GlassFish Server container on which Calendar Server is deployed.

1. Check the SMTP configuration for the following settings.

```
cd CalendarServer_home/sbin
davadmin config list|grep smtp
notification.dav.smtphost=host2.example.com
notification.dav.smtpuser=user
notification.dav.smtppassword=*****
notification.dav.smtpport=25
notification.dav.smtpstarttls=true
notification.dav.smtpusessl=false
notification.dav.smtpdebug=false
notification.dav.smtpauth=false
```

2. Check the email notifications configuration.

```
davadmin config modify -o notification.dav.smtpstarttls -v false
Enter Admin password:
davadmin config list|grep imip
notification.dav.enableimipemailnotif=false
davadmin config modify -o notification.dav.enableimipemailnotif -v true
Enter Admin password:
```

3. Check the whitelist configuration for the iSchedule port. The **service.dav.ischedulewhitelist** configuration option prevents denial of service

attacks on the iSchedule port. See the discussion on enabling the iSchedule channel to handle iMIP messages in the *Calendar Server System Administrators Guide* for more information.

4. If necessary, restart GlassFish Server. For example:

```
/opt/SUNWappserver/bin/asadmin stop-domain domain1
/opt/SUNWappserver/bin/asadmin start-domain domain1
```

Modifying iSchedule Channel Options

After you have configured the iSchedule channel, you might need to change iSchedule channel options as described in this section.

Topics in this section:

- [To Enable or Disable iMIP Message Processing](#)
- [To Modify the iSchedule Service URL](#)

To Enable or Disable iMIP Message Processing

- Use the **msconfig** command to enable or disable iMIP message processing by setting the **handle-imip** option to 1 or 0 respectively. For example, the following command disables iMIP message process:

```
/opt/sun/comms/messaging64/bin/msconfig
msconfig> set instance.channel:ischedule.options.handle-imip 0
msconfig# write
msconfig> exit
```

To Modify the iSchedule Service URL

- Use the **msconfig** command to modify the iSchedule URL by editing the **ischedule-url** option. For example:

```
/opt/sun/comms/messaging64/bin/msconfig
msconfig> set instance.channel:ischedule.options.ischedule-url
http://<host>:<port>/dav/ischedule/
msconfig# write
msconfig> exit
```

Configuring the iSchedule Channel in Legacy Configuration

This section describes how to configure the iSchedule channel for Messaging Server in legacy configuration (that is, Messaging Server 7 Update 5 and greater but you either did not convert an existing deployment to Unified Configuration, or choose to install a fresh deployment using Unified Configuration.)

1. Create the iSchedule channel.
 - a. Add following lines to the *MessagingServer_home/config/imta.cnf* file:

```
ischedule
ischedule-daemon
```
 - b. Add the following lines to the *MessagingServer_home/config/job_controller.cnf* file:

```
[CHANNEL=ischedule]
```

```
master_command=IMTA_BIN:ischedule
```

2. Configure CONVERSION mapping by adding the following lines to the *MessagingServer_home/config/mappings* file:

```
CONVERSION
```

```
IN-CHAN=ischedule;OUT-CHAN=*;TAG=*;CONVERT NO
IN-CHAN=*;OUT-CHAN=*;TAG=ISCHEDULE;CONVERT YES,CHANNEL=ischedule
```

3. Enable or disable iMIP message processing. To enable or disable iMIP message processing, create a channel options file, *MessagingServer_home/config/ischedule_option*, and add the following line:

```
handle-imip=1 (to enable)
handle-imip=0 (to disable)
```

4. Configure the iSchedule service URL. In the channel options file, specify the iSchedule Service URL as follows:

```
ischedule-url=http://<host>:<port>/dav/ischedule/
```

5. Configure the **include_conversiontag** MTA option if you want to have a {TAG=} clause included in your conversion mapping probes by adding the following line to the *MessagingServer_home/config/option.dat* file:

```
INCLUDE_CONVERSIONTAG=2
```

6. Specify messages to be processed by iSchedule. Run the following Sieve script to select all the messages that have text/calendar MIME as an attachment. This script should be placed in the location of your system-wide scripts.

```
require ["mime", "environment"];
if allof(environment :is "vnd.sun.destination-channel" ["ims-ms"],
header :mime :anychild :contenttype :is "content-type" "text/calendar",
NOT header :contains "X-Oracle-CS-iSchedule-Ignore" "Yes") {
addconversiontag "ISCHEDULE";
}
```

7. If you are using a compiled configuration, recompile the configuration.

```
/opt/sun/comms/messaging64/bin/imsimta cbuild
```

8. Restart Messaging Server.

```
/opt/sun/comms/messaging64/bin/stop-msg
/opt/sun/comms/messaging64/bin/start-msg
```

9. Verify the Calendar Server configuration. See "[Verifying the Calendar Server Configuration.](#)"

Troubleshooting the iSchedule Configuration

Use the following information to troubleshoot your iSchedule configuration:

- All messages processed through the iSchedule channel have a "Received:" header containing **ischedule-daemon.host.domain**. This is true whether handling iMIP is enabled or not. If iMIP handling is enabled, the iMIP messages have an extra header, "X-Oracle-CS-iSchedule-Status;," which contains the HTTP status code sent by the iSchedule service in response to the posted iSchedule message.

- The iSchedule channel log files are located in *MessagingServer_home/log/ischedule_master.log*.*
- Use the *MessagingServer_home/bin/imsimta qm counters* command to list the number of messages processed by the iSchedule channel.
- One common misconfiguration is to specify a wrong destination channel in the Sieve rule. If you do not see the "X-Oracle-CS-iSchedule-Status:" header in iMIP e-mails, or do not see the iSchedule counters increase when you use the **imsimta qm counters** command, check if the destination channel that you specified in the Sieve rule matches the destination channel that you have configured for that user.

Handling sendmail Clients

If users (or system utilities, for example, **cron**) send messages through **sendmail** clients, you can configure Oracle Communications Messaging Server to work with those clients over protocol. Users can continue to use the UNIX **sendmail** client.

To create compatibility between **sendmail** clients and Messaging Server, you can create and modify a **sendmail** configuration file.

Each time a new **sendmail** patch is applied to your system, you will need to modify the **submit.cf** file as described in ["To Create the sendmail Configuration File on Oracle Solaris 9 Platforms."](#)

On Oracle Solaris 9 platforms, **sendmail** is no longer a **setuid** program. Instead, it is a **setgid** program.

To Create the sendmail Configuration File on Oracle Solaris 8 Platforms

1. Find the file **main-v7sun.mc** file in directory **/usr/lib/mail/cf** and create a copy of this file.

In the example in this section, a copy called **sunone-msg.mc** is created.

2. In the **sunone-msg.mc** file, add the following lines before the **MAILER** macros:

```
FEATURE(`nullclient', `smtp:rhino.west.example.com')dnl
MASQUERADE_AS(`west.example.com')dnl
define(`confDOMAIN_NAME', `west.example.com')dnl
```

rhino.west.example.com is the localhost name and **west.example.com** is the default email domain as described in creating the initial Messaging Server runtime configuration in the *Messaging Server Installation and Configuration Guide*. In an HA environment, use the logical host name. See the *Messaging Server Installation and Configuration Guide* for information about logical hostnames for high availability.

3. Compile the **sunone-msg.mc** file:

```
/usr/ccs/bin/make sunone-msg.cf
```

The **sunone-msg.mc** will output **sunone-msg.cf**.

4. Make a backup copy of the existing **sendmail.cf** file located in the **/etc/mail** directory.
 - a. Copy and rename **/usr/lib/mail/cf/sunone-msg.cf** to **sendmail.cf** file.
 - b. Move the new **sendmail.cf** file to the **/etc/mail** directory.

To Create the sendmail Configuration File on Oracle Solaris 9 Platforms

1. Find the file **submit.mc** file in directory `/usr/lib/mail/cf` and create a copy of this file.

In the example in this section, a copy called **sunone-submit.mc** is created.

2. Change the following line in the file **sunone-submit.mc**:

```
FEATURE(`msp')dn
```

to

```
FEATURE(`msp', `rhino.west.example.com')dn1
```

rhino.west.example.com is the localhost name and **west.example.com** is the default email domain as described in creating the initial Messaging Server runtime configuration in the *Messaging Server Installation and Configuration Guide*. In an HA environment, use the logical host name. See the *Messaging Server Installation and Configuration Guide* for information about logical hostnames for high availability.

3. Compile the **sunone-submit.mc** file:

```
/usr/ccs/bin/make sunone-submit.cf
```

The **sunone-submit.mc** will output **sunone-submit.cf**.

4. Make a backup copy of the existing **submit.cf** file in the `/etc/mail` directory.
 - a. Copy and rename `/usr/lib/mail/cf/sunone-submit.cf` file to `submit.cf` file.
 - b. Move the new **submit.cf** file to the `/etc/mail` directory.

Handling Forged Email by Using the Sender Policy Framework

Spam producers and email scammers often forge email by using false domain names and email addresses or by using legitimate domain names and email addresses to fool users into thinking that a message is from someone or some company they know. For example, a spammer could send email from an address such as **president@whitehouse.gov** and the user could be fooled into thinking the mail was actually from this address. Forging email might fool users into opening the unsolicited message, or worse, provide information to a false authority. Also, spammers prefer to send their email from legitimate domains that are not on an RBL list.

Sender Policy Framework (SPF) is a technology that can detect and reject forged email during the SMTP dialogue. Specifically, SPF is a protocol that allows a domain to explicitly authorize the hosts that may use its domain name. In addition, a receiving host may be configured to check this authorization. SPF can thus significantly reduce the instances of forged email.

Note: When using Unified Configuration, you use the **msconfig** command to configure options instead of editing the legacy configuration files.

Topics:

- [Theory of Operations](#)
- [SPF Limitations](#)
- [SPF Pre-Deployment Considerations](#)
- [Setting up the Technology](#)
- [Reference Information](#)
- [Testing SPF by Using spfquery](#)
- [Handling Forwarded Mail in SPF by Using the Sender Rewriting Scheme \(SRS\)](#)

Theory of Operations

When a message comes into Oracle Communications Messaging Server, the MTA does an SPF query to determine if the address actually came from the domain on the address. An SPF query consults the DNS for TXT records belonging to the domain of the message (domain). Domain is either the domain name specified as the argument for HELO or EHLO (if the **spfhello** channel option is used) or the domain name in the

originator's address given in the MAIL FROM: command (typically the part after the @ character). If no domain name is specified or available, the one specified during HELO/EHLO is used as domain. Most ISPs distribute an authorized list of IP addresses that match their domains. If the IP address does not match the domain name, then the message is assumed to be forged.

Note: Prior to querying the DNS, the software checks the SPF_LOCAL mapping table for a match for domain. If a match is found there, it will be used first.

If a record found from the mapping table contains a **redirect=** domain clause, then the redirection to domain will be done as a DNS query, skipping the recursive and redundant mapping file check.

An example of a resulting TXT record is:

```
v=spf1 +mx a:colo.example.com/28 -all
```

The **v=spf1** token is required for SPF records supported by this RFC.

+mx checks MX records for domain and confirms that the source IP address for this SMTP connection matches one of the IP addresses given as a result of an MX query for domain. If there is a match, the + means that the result of this is **Pass**.

a:colo.example.com/28 checks for A records for **colo.example.com**, then confirm that the source IP address for this SMTP connection is in the same specified CIDR subnet as the A records, comparing only 28 bits (masked against 255.255.255.240). No qualifier character was specified, so it defaults to + meaning that a match results in a **Pass**.

Finally, **-all** matches everything else and results in **Fail**. For a complete description of SPF records, refer to RFC 4408 at <http://www.ietf.org/rfc/rfc4408.txt>.

SPF processing can have one of several results. Table 32-1 shows the results and their descriptions.

Table 32-1 SPF Processing Results

Result	Description
Pass	The lookup passed, meaning that an SPF record was found and the record validated the originating system as being authorized to use domain.
Fail	The lookup found a matching SPF record, however, the record explicitly denied authorization for the SMTP client to use domain during the SMTP transaction. The default behavior of Messaging Server's SPF implementation is to reject the SMTP command with a 5xx reply.
SoftFail	The lookup found a matching SPF record, and the record also denies authorization for the SMTP client to use domain. However, the denial is less strict and the record does not direct an outright failure. The default behavior of Messaging Server's implementation is to accept the message, but note the SoftFail in the Received-SPF: header for subsequent evaluation such as during Sieve processing.
Neutral	The SPF record makes no claim to the SMTP client's authorization to use domain. The message will be accepted. The specification requires that Neutral be treated the same as None .

Table 32–1 (Cont.) SPF Processing Results

Result	Description
None	No matching SPF record was found, therefore no SPF processing was done.
PermError	A permanent error was encountered during SPF processing, such as syntax errors in the SPF record, DNS failures while processing nested SPF records (due to include: mechanism or a redirect= modifier), or exceeding configured limits for SPF processing while processing nested SPF records. The default behavior is to reject the SMTP command with a 5xx reply.
TempError	A temporary error was encountered during SPF processing, most likely due to DNS timeouts querying SPF records. The default behavior is to reject the SMTP command with a 4xx reply.

After SPF processing has completed, a **Received-SPF:** header is written to the message documenting the result of the SPF processing. This header can then be queried during Sieve processing for subsequent consideration. Extensive debugging is available if the MTA option **MM_DEBUG** is enabled (>0). In Unified Configuration, run the **msconfig** command to set the option. In legacy configuration, set the option in the **option.dat** file.

SPF Limitations

SPF is only one tool to use to fight spam, and it does not address all issues. A spammer can easily create a domain and add an SPF TXT record that makes the domain seem legitimate. On the other hand, SPF is very effective for detecting forged email from established ISPs, although many TXT records allow the SPF to not fail.

SPF Pre-Deployment Considerations

It is important to have a very fast DNS server on your system because a DNS query for every message is required.

Setting up the Technology

The two steps to set up SPF technology are:

- Place channel options on the incoming TCP channel (typically the **tcp_local** channel, although there might be other channels if you allow channel switching from **tcp_local** to another channel). See [Table 32–2, "SPF Options"](#) for more information.
- Set up the options. In Unified Configuration, run the **msconfig** command. In legacy configuration, edit the **option.dat** file. See [Table 32–3, "SPF Limiting Options"](#) for more information.

Reference Information

This section provides reference information for the SPF channel options and the SPF MTA options. SPF support is implemented through four channel options applied to the incoming **tcp_*** channel (typically **tcp_local**). [Table 32–2](#) shows the options and their descriptions. In Unified Configuration, use the **msconfig edit channels** command to edit SPF channel options. In legacy configuration, edit the **imta.cnf** file.

Table 32–2 SPF Options

Option	Description
spfnone	Disables SPF processing
spfhelo	Enables SPF processing for the domain name specified as an argument to HELO or EHLO.
spfmailfrom	Enables SPF processing for the domain name provided for the originator envelope address after receiving the MAIL FROM:.
spfrcptto	Enables SPF process for the domain name provided for the originator envelope address after receiving the RCPT TO:. Processing is the same as spfmailfrom except that it is delayed in the SMTP transaction until after the RCPT TO: command has been issued and the recipient has otherwise been confirmed to be a valid recipient.

Note: **spfmailfrom** and **spfrcptto** are conflicting options and you should only specify one of these two options on the channel. You can, however, use **spfhelo** in conjunction with either **spfmailfrom** or **spfrcptto** to perform both kinds of SPF checks.

Additional support to establish limits on SPF processing and to control whether SMTP commands will be accepted, failed with a 4xx response (temporary failure), or failed with a 5xx response (permanent failure) for the various SPF results includes: **Fail**, **SoftFail**, **PermError**, and **TempError**.

Use the following MTA options to place limits on SPF processing. In Unified Configuration, set options with the **msconfig** command. In legacy configuration, edit the **option.dat** file.

Table 32–3 SPF Limiting Options

Option	Description
SPF_MAX_RECURSION	Specifies the number of recursions that will be allowed into nested SPF records due to include: or redirect= . Exceeding this limit will result in a PermError . Default: 10 (mandated by the RFC)
SPF_MAX_DNS_QUERIES	Specifies the number of mechanisms or modifiers that require DNS lookups (including include: , a: , mx: , ptr: , exists: , redirect= , and exp=). The limit is not counted as the number of actual DNS lookups, so one mechanism could lead to several DNS queries. Exceeding this limit will result in a PermError . Default: 10 (mandated by the RFC)
SPF_MAX_TIME	Specifies the number of seconds that will be allowed for the SPF processing to complete. Exceeding this value will result in a TempError . The default value is more generous than the RFC suggests. Default: 45

Additionally, the following MTA options can be configured to control the behavior of the SMTP server in response to SPF results of **Fail**, **SoftFail**, **PermError**, and **TempError**. In Unified Configuration, run the **msconfig** command. In legacy configuration, edit the **option.dat** file. For each of these results, the SMTP server can send back a 2xx (success) response, 4xx (temporary failure), or 5xx (permanent failure). Also, for **Fail** and **SoftFail**, the MTA can distinguish between an SPF result as the result of an "all" mechanism versus an otherwise explicitly referenced match. You can

then make a distinction between a particular result and the SPF record's default result. The valid values for any of these options is 2, 4, or 5. The values of 2, 4, or 5 correspond to 2xx, 4xx, or 5xx responses from the SMTP server as a result of getting that particular SPF status. So, for example, if `SPF_SMTP_STATUS_FAIL=2` and the SPF record explicitly blocks us with a "-a:192.168.1.44" (our IP address), then instead of responding with a 5xx response, we'll accept the address with a "250 OK" instead.

Table 32-4 SPF Failure and Error Options

Option	Description
<code>SPF_SMTP_STATUS_FAIL</code>	Used when the match of an SPF record is a "-" flagged mechanism other than "-all." Default: 5
<code>SPF_SMTP_STATUS_FAIL_ALL</code>	Used when the matching mechanism is "-all." Default: 5
<code>SPF_SMTP_STATUS_SOFTFAIL</code>	Used when the match of an SPF record is a "~" flagged mechanism other than "~all." Default: 2
<code>SPF_SMTP_STATUS_SOFTFAIL_ALL</code>	Used when the matching mechanism is "~all." Default: 2
<code>SPF_SMTP_STATUS_TEMPERROR</code>	Used when there is a temporary failure, usually related to DNS processing problems. Default: 4
<code>SPF_SMTP_STATUS_PERMERROR</code>	Used when there is a permanent failure, usually due to syntax or other technical errors found during SPF processing. (This is due to a non-local error.) Default: 5

Testing SPF by Using spfquery

You can use the `spfquery` testing utility to test SPF processing.

Note: `spfquery` does not test your SPF configuration. It tests what would be returned if you were to enable SPF processing.

Requirements: Must be run as a user who has access to run the Messaging Server binaries and access its libraries such as `root` or `mailsrv`, for example.

Location: `MessagingServer_home/bin/`

Syntax

```
spfquery [-i ip-address] [-s sender-email] [-h helo-domain] [-e none | neutral |
pass | fail | temperror | permerror] [-v] [-V] [?] domain
```

Table 32-5 shows the `spfquery` options and their descriptions.

Table 32–5 *spfquery Options*

Option	Description
-i <i>ip address</i>	Specifies the IP address to be used as the remote address for the SPF query. Default is 127.0.0.1 . This option can also be --ip-address .
-s <i>domain</i>	The email address that will be used as if it were specified as MAIL FROM: . Default: postmaster@domain . This option can also be --sender .
-h <i>helo-domain</i>	The domain name as if it were specified for the HELO domain. This domain is not verified itself, but instead provided as supplemental information for macro processing. Default value is the same as the value you specified for domain. This option can also be --helo-domain .
-e <i>result</i>	spfquery compares the result of the SPF processing with what is expected and if the result is different, a message is printed and spfquery exits with a non-zero return status. Result can be one of: none , neutral , pass , fail , softfail , temperror , or permerror . This option can also be --expect .
-v	Enables verbose output during SPF processing. This option can also be --verbose .
-V	Prints the current version of the SPF library. This option can also be --version .
-?	Prints this usage information. This option can also be --help .

Example with Debugging Enabled

```
# /opt/SUNWmsgsr/sbin/spfquery -v -i 192.168.1.3 11.spf1-test.example.com
Running SPF query with:
  IP address: 192.168.1.3
  Domain: 11.spf1-test.example.com
  Sender: postmaster@11.spf1-test.example.com (local-part: postmaster)
  HELO Domain: 11.spf1-test.example.com

15:30:04.33: -----
15:30:04.33: SPFcheck_host called:
15:30:04.33:   source ip = 192.168.1.3
15:30:04.33:   domain = 11.spf1-test.example.com
15:30:04.33:   sender = postmaster@11.spf1-test.example.com
15:30:04.33:   local_part = postmaster
15:30:04.33:   helo_domain = 11.spf1-test.example.com
15:30:04.33:
15:30:04.33: Looking up "v=spf1" records for 11.spf1-test.example.com
15:30:04.35:   DNS query status: Pass
15:30:04.35:   "v=spf1 mx:spf1-test.example.com -all"
15:30:04.35:
15:30:04.35: Parsing mechanism: " mx : spf1-test.example.com"
15:30:04.35:   Assuming a Pass prefix
15:30:04.35:   Processing macros in spf1-test.example.com
15:30:04.35:   Comparing against 192.168.1.3
15:30:04.35:   Looking for MX records for spf1-test.example.com
15:30:04.41:     mx02.spf1-test.example.com:
15:30:04.41:       192.0.2.22 - No match
15:30:04.41:       192.0.2.21 - No match
15:30:04.41:       192.0.2.20 - No match
15:30:04.41:       192.0.2.23 - No match
15:30:04.41:     mx01.spf1-test.example.com:
15:30:04.42:       192.0.2.13 - No match
```

```

15:30:04.42:          192.0.2.11 - No match
15:30:04.42:          192.0.2.12 - No match
15:30:04.42:          192.0.2.10 - No match
15:30:04.42:          mx03.spf1-test.example.com:
15:30:04.42:          192.0.2.32 - No match
15:30:04.42:          192.0.2.30 - No match
15:30:04.42:          192.0.2.31 - No match
15:30:04.42:          192.168.1.3 - Matched
15:30:04.42: Mechanism matched; returning Pass
15:30:04.42:
15:30:04.42: Parsing mechanism: "- all : " (not evaluated)
15:30:04.42:
15:30:04.42: SPFcheck_host is returning Pass
15:30:04.42: -----

```

Handling Forwarded Mail in SPF by Using the Sender Rewriting Scheme (SRS)

As described above, SPF is a mechanism that attempts to prevent email forgery by looking up special **TXT** records associated with the domain in the mail **FROM:** (envelope from) address. This operation, which can actually involve several DNS lookups, eventually produces a list of IP addresses that are authorized to send mail from the domain. The IP address of the SMTP client is checked against this list and if it isn't found, the message can be considered to be fraudulent.

SPF presents serious problems for sites that provide mail forwarding services such as universities (for their alumni) or professional organizations (for their members). A forwarder ends up sending out mail from essentially arbitrary senders, which can include senders who have implemented SPF policies and which, of course, do not list the IP addresses of the forwarding system or systems as being permitted to use addresses from their domain.

The Sender Rewriting Scheme, or SRS, provides a solution to this problem. SRS works by encapsulating the original sender's address inside a new address using the forwarder's own domain. Only the forwarder's own domain is exposed for purposes of SPF checks. When the address is used it routes the mail (usually a notification) to the forwarder, which removes the address encapsulation and sends the message on to the real destination.

Of course address encapsulation isn't exactly new. Source routes were defined in RFC 822 and provide exactly this sort of functionality, as does percent hack routing and bang paths. However, these mechanisms are all problematic on today's Internet since allowing their use effectively turns your system into an open relay.

SRS deals with this problem by adding a keyed hash and a timestamp to the encapsulation format. The address is only valid for some period of time, after which it cannot be used. The hash prevents modification of either the timestamp or the encapsulated address.

SRS also provides a mechanism for handling multi-hop forwarding without undue growth in address length. For this to work certain aspects of SRS address formatting have to be done in the same way across all systems implementing SRS.

The following MTA options have been added:

- **SRS_DOMAIN.** This must be set to the domain to use in SRS addresses. Email sent to this domain must always be routed to a system capable of SRS operations for the domain. SRS processing is handled as an overlay on top of normal address

processing so nothing prevents a site from using their primary domain as the SRS domain.

- **SRS_SECRETS.** This is a comma separated list of secret keys used to encode and decode SRS addresses. The first key on the list is used unconditionally for encoding. For decoding, each key is tried in order to generate a different hash value. The decoding operation proceeds if any of the hashes match.

The ability to use multiple keys makes it possible to change secrets without service disruption: Add a second key, wait for all previously issued addresses to time out, and then remove the first key.

- **SRS_MAXAGE.** Optionally specifies the number of days before a message times out. The default if the option isn't specified is 14 days.

Every system that handles email for the selected SRS domain must be configured for SRS processing and must have all three SRS options set identically.

Setting these options is sufficient to enable SRS address decoding. Encoding is another matter and should only be done to envelope **From:** addresses you know are associated with forwarding activity. SRS encoding is controlled by six new channel options: **addresssrs**, **noaddresssrs**, **destinationssrs**, **nodestinationssrs**, **sourcecsrs**, and **nosourcecsrs**.

Three conditions have to be met for SRS encoding to occur:

1. The current source channel has to be marked with **sourcecsrs**. (**nosourcecsrs** is the default).
2. The current destination channel has to be marked with **destinationssrs**. (**nodestinationssrs** is the default).
3. The current address, when rewritten, has to match a channel marked **addresssrs**. (**noaddress** is the default).

Encoding only occurs when all of these conditions are true. The simplest setup consists of pure forwarding where all messages enter and exit on the **tcp_local** channel and all non-local addresses need SRS handling. In such a setup, **tcp_local** would be marked with the three options **sourcecsrs**, **destinationssrs**, and **addresssrs**.

Finally, **imsimta test -rewrite** has been enhanced to show SRS encoding and decoding results for whatever address is input. For example, the address **foo@example.com** might produce the output similar to:

```
SRS encoding = SRS0=dnG=IS=example.com=foo@example.org
```

If this encoded address is rewritten it produces the following output:

```
SRS decoding = foo@example.com
```


imsimta test -rewrite also shows any errors that occur during SRS decoding.

Message Store Directory Layout

The "Message Store Directory Layout" and Table 33-1 show and describe the message store directory layout.

Message Store Directory Layout

Figure 33-1 Message Store Directory Layout



The message store consists of a number of mailbox databases and the user mailboxes. The mailbox databases consists of information about users, mailboxes, partitions, quotas and other message store related data. The user mailboxes contain the user's messages and folders. Mailboxes are stored in a *message store partition*, an area on a *disk partition* specifically devoted to storing the message store. See "Managing Message Store Partitions and Adding Storage" for details. Message store partitions are not the same as disk partitions, though for ease of maintenance, we recommend having one disk partition for each message store partition.

Mailboxes such as INBOX are located in the **store_root**. For example, a sample directory path might be:

```
<store_root>/partition/primary/=user/53/53/=mack1
```

Table 33–1 describes the message store directory.

Table 33–1 Message Store Directory Description

Location	Content/ Description
<i>MessagingServer_home</i>	Default for Oracle Communications Messaging Server: /opt/sun/comms/messaging64 The directory on the Messaging Server machine that holds the server program, configuration, maintenance, and information files.
<i>store_root</i>	<i>MessagingServer_home/data/store</i> Top-level directory of the message store. Contains the mboxlist , user , and partition subdirectories.
./store.expirerule	Contains the automatic message removal rules (expire rules). This optional file can be at different locations. See "Message Store Message Expiration."
<i>store_root/dbdata/snapshots</i>	Message store database backup snapshots that stored makes periodically.
<i>store_root/mboxlist/</i>	Contains mailbox database, database (Berkeley DB) that stores information about the mailboxes and quota information. annotate.db supports the ANNOTATE extension to IMAP, which permits clients and servers to maintain "meta data" for messages, or individual message parts, stored in a mailbox on the server. For example, you could use IMAP ANNOTATE to attach comments and other useful information to a message, or to attach annotations to specific parts of a message, marking them as seen or important, or a comment added. folder.db contains information about mailboxes, including the name of the partition where the mailbox is stored, the ACL, and a copy of some of the information in store.idx . There is one entry in folder.db per mailbox. quota.db contains information about quotas and quota usage. There is one entry in quota.db per user. lright.db is an index for the folders by acl lookup rights. peruser.db contains information about per-user flags. The flags indicate whether a particular user has seen or deleted a message. subscr.db contains information about user subscriptions.
<i>store_root/session/</i>	Contains active message store process information.
<i>store_root/user/</i>	Not used.
<i>store_root/partition/</i>	Contains the message store partitions. A default primary partition is created. Place any other partitions you define in this directory.
<i>store_root/partition/primary/=user/</i>	Contains all the user mailboxes in the subdirectory of the partition. The mailboxes are stored in a hash structure for fast searching. To find the directory that contains a particular user's mailbox, use the hashdir utility.
<i>.../=user/hashdir/hashdir/userid/</i>	The top-level mail folder for the user whose ID is <i>userid</i> . This contains the user's INBOX. For the default domain, <i>userid</i> is <i>uid</i> . For hosted domains, <i>userid</i> is <i>uid@domain</i> . A user's incoming messages are delivered to the INBOX here.

Table 33–1 (Cont.) Message Store Directory Description

Location	Content/ Description
<i>.../userid/folder</i>	A user-defined mailbox on the Messaging Server host.
<i>.../userid/store.idx</i>	An index that provides the following information about mail stored in the <i>/userid/</i> directory: number of messages, disk quota used by this mailbox, the time the mailbox was last appended, message flags, variable-length information for each message including the headers and the MIME structure, and the size of each message. The index also includes a backup copy of mbolist information for each user and a backup copy of quota information for each user.
<i>.../userid/store.usr</i>	Contains a list of users who have accessed the folder. For each user listed, contains information about the last time the user accessed the folder, the list of messages the user has seen, and the list of messages the user has deleted.
<i>.../userid/store.sub</i>	Contains information about user subscriptions.
<i>.../userid/store.exp</i>	Contains a list of message files that have been expunged, but not removed from disk. This file appears only if there are expunged messages.
<i>.../userid/ nn/</i> or <i>.../userid/folder/nn/</i>	<i>nn</i> is a hash directory that contains messages in the format <i>message_id.msg</i> ; <i>nn</i> can be a number from 00 to 99. <i>message_id</i> is also a number. Example: messages 1 through 99 are stored in the <i>.../00</i> directory. The first message is 1.msg , the second is 2.msg , third 3.msg , and so on. Messages 100 through 199 are stored in the 01 directory; messages 9990 through 9999 are stored in the 99 directory; messages 10000 through 10099 are in the 00 directory, and so on.

Shared Folders Overview

See "[Managing Shared Folders](#)" for shared folder tasks.

A *shared folder* is like any other mail folder except that users other than its owner can read, delete, or add messages to it, depending on the access rights they are granted. Messages can be added to shared folders by normal drag and drop, by Sieve filters, or by sending messages directly using the form: *uid+folder@domain*.

The example below shows the address for sending email to a *private shared folder* owned by **carol.fanning@example.com** called **crafts_club**:

```
carol.fanning+crafts_club@example.com
```


This example shows the address for sending email to a *public shared folder* called **tennis**:

```
public+tennis@example.com
```

Shared folders are useful for starting, sharing, and archiving an ongoing email conversation on a particular topic. For example, a group of software developers can create a shared folder for discussing development of a particular project called **mosaic_voices**. When a message is sent or dropped into the folder **mosaic_voices**, anyone who has permissions to access the shared folder (permissions can be granted to individuals or groups) can open this mailbox and read the message.

Shared folders are displayed in user's mailbox tree under a folder called **Shared Folders**. An example is shown below.

Figure 34-1 Example of Shared Mail Folder List as Seen from a Mail Client



There are two kinds of shared folders:

- **Private Shared Folder** - A shared folder created and owned by a specific user with access rights granted to other users or groups. The owner can grant access rights using Convergence or other mail clients that support shared folder creation. The mail administrator can also grant access rights using the **readership** command. Private shared folders appear in the **Shared Folders/User** mail folder directory.
- **Public Shared Folder** - A shared folder created by the mail administrator and not owned by a specific user. The mail administrator can grant access rights using the **readership** command. Public shared folders appear in the **Shared Folders/Public** mail folder directory.

For example, you might want a folder, such as **public+software_dev@example.com** for posting information about a special interest group inside the company. Interested employees would be granted access to this public folder.

Oracle Communications Messaging Server allows folders to be shared among users of different backend message stores. See "[Setting Up Distributed Shared Folders](#)" for details.

See "[Managing Shared Folders](#)" for examples of creating shared folders and granting access rights.

Monitoring LDAP Directory Server

The LDAP directory server (**slapd**) provides directory information for the messaging system. If **slapd** is down, the system will not work properly. If **slapd** response time is too slow, this will affect login speed and any other transaction that requires LDAP lookups.

Topics:

- [Symptoms of slapd Problems](#)
- [To Monitor slapd](#)

Symptoms of slapd Problems

- Client POP, IMAP, or Webmail Authentication fails or slower than expected.
- MTA not working properly

To Monitor slapd

- Check that **ns-slapd** process is running.
- Check **slapd** log files **access** and **errors** in **slapd-instance/logs/**
- Check the **ns-slapd** response time while searching for a user.
- See also "[immonitor-access](#)."

Monitoring System Performance

This information focuses on Oracle Communications Messaging Server monitoring, however, you also need to monitor the system on which the server resides. A well-configured server cannot perform well on a poorly-tuned system, and symptoms of server failure may be an indication that the hardware is not powerful enough to serve the email load. This information does not provide all the details for monitoring system performance as many of these procedures are platform specific and may require that you refer to the platform specific system documentation.

This chapter contains the following topics:

- [Monitoring End-to-end Message Delivery Times](#)
- [Monitoring CPU Usage](#)

Monitoring End-to-end Message Delivery Times

Email needs to be delivered on time. This may be a service agreement requirement, but also it is good policy to have mail delivered as quickly as possible. Slow end-to-end times could indicate a number of things. It may be that the server is not working properly, or that certain times of the day experience overwhelming message loads, or that the existing hardware resources are being pushed beyond their capacity.

Symptoms of Poor End-to-end Message Delivery Times

Mail takes a longer period of time to be delivered than normal.

To Monitor End-to-end Message Delivery Times

- Use any facility that sends a message and receives it. Compare the headers times between server hops, and times between point of origin and retrieval. See "[immonitor-access](#)."

Monitoring CPU Usage

High CPU usage is either a sign that there is not enough CPU capacity for the level of usage or some process is using up more CPU cycles than is appropriate.

Symptoms of CPU Usage Problems

Poor system response time. Slow logging in of users. Slow rate of delivery.

To Monitor CPU Usage

Monitoring CPU usage is a platform specific task. Refer to the relevant platform documentation.

Monitoring the Message Store

This information describes message store monitoring tasks. See "[Managing the Message Store and Mailboxes](#)" for conceptual information.

Topics:

- [General Message Store Monitoring Procedures](#)
- [Monitoring imapd, popd and httpd](#)
- [Monitoring stored](#)
- [Monitoring the State of Message Store Database Locks](#)
- [To Monitor Mailbox Quotas and Usage](#)
- [To Monitor Message Store Database Statistics with imcheck](#)
- [Gathering Message Store Counter Statistics by Using counterutil](#)

For More Information

- [Monitoring Disk Space](#)
- [Monitoring User Access to the Message Store](#)
- [Using Message Store Log Messages](#)

General Message Store Monitoring Procedures

This section outlines standard monitoring procedures for the message store. These procedures are helpful for general message store checks, testing, and standard maintenance.

Topics in this section:

- [Checking Hardware Space](#)
- [Checking Log Files](#)
- [Checking User IMAP/POP/Webmail Session by Using Telemetry](#)
- [Checking stored Processes](#)
- [Checking Database Log Files](#)
- [Checking User Folders](#)
- [Checking for Core Files](#)

Checking Hardware Space

A message store should have enough additional disk space and hardware resources. When the message store is near the maximum limit of disk space and hardware space, problems might occur within the message store.

Inadequate disk space is one of the most common causes of the mail server problems and failure. Without space to write to the message store, the mail server will fail. In addition, when the available disk space goes below a certain threshold, there will be problems related to message delivery, logging, and so forth. Disk space can be rapidly depleted when the clean up function of the **stored** process fails and deleted messages are not expunged from the message store.

For information on monitoring disk space, see "[Monitoring Disk Space](#)."

Checking Log Files

Check the log files to make sure the message store processes are running as configured. Oracle Communications Messaging Server creates a separate set of log files for each of the major protocols, or services, it supports: SMTP, IMAP, POP, and HTTP. You can look at the log files in the *MessagingServer_home/log/* directory. You should monitor the log files on a routine basis.

Be aware that logging can impact server performance. The more verbose the logging you specify, the more disk space your log files will occupy for a given amount of time. You should define effective but realistic log rotation, expiration, and backup policies for your server. For information about defining logging policies for your server, see "[Using Message Store Log Messages](#)."

Checking User IMAP/POP/Webmail Session by Using Telemetry

Messaging Server provides a feature called telemetry that can capture a user's entire IMAP, POP or HTTP session into a file. This feature is useful for debugging client problems. For example, if users complain that their message access client is not working as expected, this feature can be used to trace the interaction between the access client and Messaging Server.

To capture a POP session, create the following directory:

```
MessagingServer_home/data/telemetry/pop_or_imap_or_http/userid
```

To capture a POP session, create the following directory:

```
MessagingServer_home/data/telemetry/pop/userid
```

To capture an IMAP session, create the following directory:

```
MessagingServer_home/data/telemetry/imap/userid
```

To capture a Webmail session, create the following directory:

```
MessagingServer_home/data/telemetry/http/userid
```

Note: *userid* is "uid" for default domain and "uid@domain" for hosted domains.

Note that the directory must be owned or writable by the messaging server *userid*.

Messaging Server will create one file per session in that directory. Example output is shown below.

```
LOGIN redb 2003/11/26 13:03:21
>0.017>1 OK User logged in
<0.047<2 XSERVERINFO MANAGEACCOUNTURL MANAGELISTSURL MANAGEFILTERSURL
>0.003>* XSERVERINFO MANAGEACCOUNTURL {67}
```

```

http://redb@cuisine.blue.planet.com:800/bin/user/admin/bin/enduser
MANAGELISTSURL NIL MANAGEFILTERSURL NIL
2 OK Completed
<0.046<3 select "INBOX"
>0.236>* FLAGS (\Answered flagged draft deleted \Seen $MDNSent Junk)
* OK [PERMANENTFLAGS (\Answered flag draft deleted \Seen $MDNSent Junk \*)]
* 1538 EXISTS
* 0 RECENT
* OK [UNSEEN 23]
* OK [UIDVALIDITY 1046219200]
* OK [UIDNEXT 1968]
3 OK [READ-WRITE] Completed
<0.045<4 UID fetch 1:* (FLAGS)
>0.117>* 1 FETCH (FLAGS (\Seen) UID 330)
* 2 FETCH (FLAGS (\Seen) UID 331)
* 3 FETCH (FLAGS (\Seen) UID 332)
* 4 FETCH (FLAGS (\Seen) UID 333)
* 5 FETCH (FLAGS (\Seen) UID 334)
<etc>

```

You can gather command telemetry that does not include end-user information by using the **imap.logcommands msconfig** option (or in legacy configuration **local.imap.logcommands**). See the *Messaging Server Reference* for additional information.

To disable the telemetry logging, move or remove the directory that you created.

Checking stored Processes

The **stored** function performs a variety of important tasks such as deadlock and transaction operations of the message database, enforcing aging policies, and expunging and erasing messages stored on disk. If **stored** stops running, Messaging Server will eventually run into problems. If **stored** does not start when **start-msg** is run, no other processes will start.

- Check that the **stored** process is running. Run "**imcheck**."
- Check for the log file build up in *store_root/mboxlist*.
- Check for **stored** messages in the default log file *MessagingServer_home/log/default/default*.
- Check that the time stamps of the following files (in directory *MessagingServer_home/config*) are updated whenever one of the following functions are attempted by the **stored** process:

Table 37–1 *stored Operations*

stored Operation	Function
stored.ckp	Touched when a database checkpoint was initiated. Stamped approximately every 1 minute.
stored.lcu	Touched at every database log cleanup. Time stamped approximately every 5 minutes.
stored.per	Touched at every spawn of peruser db write out. Time stamped once an hour.

For more information on the **stored** process, see "**stored**." For additional information on monitoring the **stored** function, see "**Monitoring stored**."

Checking Database Log Files

Database log files refer to sleepycat transaction checkpointing log files (in directory *store_root/mboxlist*). If log files accumulate, then database checkpointing is not occurring. In general, there are two or three database log files during a single period of time. If there are more files, it could be a sign of a problem.

Checking User Folders

If you want to check the user folders, you might run the command **reconstruct -r -n** (recursive no fix) which will review any user folder and report errors. For more information on the **reconstruct** command, see "[Repairing Mailboxes and the Mailboxes Database \(reconstruct Command\)](#)."

Checking for Core Files

Core files only exist when processes have unexpectedly terminated. It is important to review these files, particularly when you see a problem in the message store. On Oracle Solaris, use **coreadm** to configure **core** file location.

Monitoring imapd, popd and httpd

These processes provide access to IMAP, POP and Webmail services. If any of these is not running or not responding, the service will not function appropriately. If the service is running, but is over loaded, monitoring will allow you to detect this and configure it more appropriately.

Topics in this section:

- [Symptoms of imapd, popd and httpd Problems](#)
- [To Monitor imapd, popd and httpd](#)

Symptoms of imapd, popd and httpd Problems

Connections are refused or system is too slow to connect. For example, if IMAP is not running and you try to connect to IMAP directly you will see something like this:

```
telnet 0 143 Trying 0.0.0.0... telnet: Unable to connect to remote host: Connection refused
```

If you try to connect with a client, you will get a message such as:

"Client is unable to connect to the server at the location you have specified. The server may be down or busy."

To Monitor imapd, popd and httpd

- Can be monitored with **watcher** and **msprobe**. See "[Automatic Restart of Failed or Unresponsive Services](#)" and "[Monitoring Using msprobe and watcher Functions](#)."
- Can be monitored with SNMP. If you have the SNMP set up, this is a very good way to monitor these processes. See "[SNMP Support](#)." The server information is in the Network Services Monitoring MIB.
- Check log files. Look in the directory *MessagingServer_home/log/service* where *_service_* can be HTTP or IMAP or POP. In that directory you will find a number of log files. One filename is the name of the *service* (imap, pop, http) and the others

are the name of the service plus a sequence number and a date concatenated to the service name. For example:

```
imap imap.29.1010221593 imap.31.1010394412 imap.33.1010567224
```

The file with just the service name is the latest log. The other ones are ordered by the sequence number (here 29, 31, 33) and the one with the highest sequence number is the next newest one. (See ["Using Message Store Log Messages."](#))

If a server was shut down you might see something like this:

```
imap.12.1065431243:[07/Oct/2003:01:15:43 -0700] gotmail-2 imapd[20525]: General
Warning: Sun Java System Messaging Server IMAP4 6.1 (built Sep 24 2003) shutting
down
```

- Can be checked with ["counterutil."](#) See ["Gathering Message Store Counter Statistics by Using counterutil."](#)
- Run the platform-specific command to verify that the `imapd`, `popd` and `httpd` processes are running. For example, in Oracle Solaris you can use the `ps` command and look for **imapd**, **popd** and **mshttpd**.
- You can set alarms for specified server performance thresholds by setting the server response configuration options described in ["Alarm Messages."](#)
- See ["immonitor-access."](#)

Monitoring stored

`"stored"` performs a variety of important tasks such as deadlock and transaction operations of the message database, enforcing aging policies, and expunging and erasing messages stored on disk. If **stored** stops running, the messaging server will eventually run into problems. If **stored** does not start when **start-msg** is run, no other processes will start. See `"stored"` for more information.

Topics in this section:

- [Symptoms of stored Problems](#)
- [To Monitor stored](#)

Symptoms of stored Problems

There are no outward symptoms.

To Monitor stored

- Check that the **stored** process is running. **stored** creates and updates a **pid** file in `MessagingServer_home/data/proc` called **store**. The **pid** file shows an **init** state when recovering and a **ready** state when ready. For example:

```
231: cat store
28250
ready
```

The number on the first line is the process ID of **stored**.

```
232: ps -eaf | grep stored
inetuser 28250 1 0 Jan 05 ? 8:44
/opt/SUNWmsgsr/lib/stored -d
```

- Check for log file build up in *MessagingServer_home/store/mboxlist*. Note that not every log file build up is caused by direct **stored** problems. Log files may also build up if **imapd** dies or there is a database problem.
- Check the timestamp on the following files in *MessagingServer_home/config*:
 - stored.ckp** - Touched when attempt at checkpointing is made. Should get time stamped every 1 minute.
 - stored.lcu** - Touched at every db log cleanup. Should get time stamped every 5 minutes.
 - stored.per** - Touched at every spawn of peruser db writeout. Should get time stamped every 60 minutes.
- Check for **stored** messages in the default log file *MessagingServer_home/log/default/default*
- Can be monitored with **watcher** and **msprobe**. See "[Automatic Restart of Failed or Unresponsive Services](#)" and "[Monitoring Using msprobe and watcher Functions](#)."

Monitoring the State of Message Store Database Locks

The state of database-locks is held by different server processes. These database locks can affect the performance of the message store. In case of deadlocks, messages will not be getting inserted into the store at reasonable speeds and the **ims-ms** channel queue will grow larger as a result. There are legitimate reasons for a queue to back up, so it is useful to have a history of the queue length in order to diagnose problems.

Topics in this section:

- [Symptoms of Message Store Database Lock Problems](#)
- [To Monitor Message Store Database Locks](#)

Symptoms of Message Store Database Lock Problems

Number of transactions are accumulating and not resolving.

To Monitor Message Store Database Locks

Use the command "[imcheck](#)" -s (used to be [counterutil -o db_lock](#)).

To Monitor Mailbox Quotas and Usage

You can monitor mailbox quota usage and limits by using the "[imquotacheck](#)" utility. The **imquotacheck** utility generates a report that lists defined quotas and limits, and provides information on quota usage.

For example, the following command lists all user quota information:

```
% imquotacheck
-----
Domain red.example.com (diskquota = not set msgquota = not set) quota usage
-----
diskquota      size(K)    %use    msgquota    msgs    %use    user
# of domains = 1
# of users = 705
no quota       50418     no quota  4392        ajonk
no quota        5         no quota  2           andrt
no quota       355518    no quota  2500        ansri
```

...

The following example shows the quota usage for user sorook:

```
% imquotacheck -u sorook
-----
quota usage for user sorook
-----
diskquota      size(K)    %use      msgquota    msgs      %use      user
no quota       1487                no quota    305                sorook
```

To list the usage of all users whose quota exceeds the least threshold in the rule file:

imquotacheck

To list quota information for a the domain **example.com**:

imquotacheck -d example.com

To send a notification to all users in accordance to the default rule file:

imquotacheck -n

To send a notification to all users in accordance to a specified *rulefile*, *myrulefile*, and to a specified mail template file, *mytemplate.file* (for more information, refer to "imquotacheck"):

imquotacheck -n -r myrulefile -t mytemplate.file

To list per folder usages for one user user1 (will ignore the rule file):

imquotacheck -u user1 -e

To Monitor Message Store Database Statistics with imcheck

Use **imcheck -s** to monitor database statistics including logs and transactions. See "imcheck."

Gathering Message Store Counter Statistics by Using counterutil

Topics in this section:

- [To Get a Current List of Available Counter Objects](#)
- [counterutil Output](#)
- [Gathering Alarm Statistics by Using counterutil](#)
- [IMAP, POP, HTTP, and MMP Connection Statistics by Using counterutil](#)
- [Disk Usage Statistics by Using counterutil](#)
- [Server Response Statistics](#)

To Get a Current List of Available Counter Objects

This utility provides statistics acquired from different system counters. (See "counterutil.")

Here is how to get a current list of available counter objects:

```
# counterutil -l
Listing registry (/opt/sun/comms/messaging64/data/counter/counter)
numobjects = 7
refcount = 20
```

```

created = 17/Mar/2015:14:10:03 +0000
modified = 24/Aug/2015:13:00:24 +0000
counterobjects:
  imapstat
  popstat
  alarm
  serverresponse
  diskusage
  httpstat
  mmpstat

```

Each entry represents a counter object and supplies a variety of useful counts for this object. In this section we will only be discussing the **alarm**, **diskusage**, **serverresponse**, **popstat**, **imapstat**, and **httpstat** counter objects. For details on **counterutil** command usage, refer to "[counterutil](#)."

counterutil Output

"**counterutil**" has a variety of flags. A command format for this utility may be as follows:

```
counterutil -oCounterObject-i 5 -n 10
```

where,

-oCounterObject represents the counter object **alarm**, **diskusage**, **serverresponse**, **popstat**, **imapstat**, and **httpstat**.

-i 5 specifies a 5 second interval.

-n 10 represents the number of iterations (default: infinity).

An example of **counterutil** usage is as follows:

```

# counterutil -o imapstat -i 5 -n 10
Monitor counterobject (imapstat)
registry /gotmail/iplanet/server5/msg-gotmail/counter/counter opened
counterobject imapstat opened
count = 1 at 972082466 rh = 0xc0990 oh = 0xc0968
global.currentStartTime [4 bytes]: 17/Oct/2000:12:44:23 -0700
global.lastConnectionTime [4 bytes]: 20/Oct/2000:15:53:37 -0700
global.maxConnections [4 bytes]: 69
global.numConnections [4 bytes]: 12480
global.numCurrentConnections [4 bytes]: 48
global.numFailedConnections [4 bytes]: 0
global.numFailedLogins [4 bytes]: 15
global.numGoodLogins [4 bytes]: 10446
...

```

Gathering Alarm Statistics by Using counterutil

These alarm statistics refer to the alarms sent by **stored**. The alarm counter provides the following statistics:

Table 37–2 counterutil alarm Statistics

Suffix	Description
alarm.countoverthreshold	Number of times crossing threshold.
alarm.countwarningsent	Number of warnings sent.

Table 37–2 (Cont.) counterutil alarm Statistics

Suffix	Description
alarm.current	Current monitored valued.
alarm.high	Highest ever recorded value.
alarm.low	Lowest ever recorded value.
alarm.timelastset	The last time current value was set.
alarm.timelastwarning	The last time warning was sent.
alarm.timereset	The last time reset was performed.
alarm.timestatechanged	The last time alarm state changed.
alarm.warningstate	Warning state (yes(1) or no(0)).

IMAP, POP, HTTP, and MMP Connection Statistics by Using counterutil

To get information on the number of current IMAP, POP, HTTP, and MMP connections, number of failed logins, total connections from the start time, and so forth, you can use the command **counterutil -oCounterObject-i 5 -n 10**. Where *CounterObject* represents the counter object **popstat**, **imapstat**, **httpstat**, or **mmpstat**. For **mmpstat**, we have modified the counter names to differentiate the services IMAP and POP since the MMP proxies both. The meaning of the **imapstat** suffixes is shown in Table 37–3. The **popstat** and **httpstat** objects provide the same information in the same format and structure.

Table 37–3 counterutil imapstat Statistics

Suffix	Description
currentStartTime	Start time of the current IMAP server process.
lastConnectionTime	The last time a new client was accepted.
maxConnections	Highest recorded number of concurrent TCP connections handled by IMAP server since the last counter reset.
numConnections	Total number of TCP connections successfully accepted by the current IMAP server. numConnections can include failed connections, but not always.
numCurrentConnections	Current number of active TCP connections.
numFailedConnections	Total number of failed TCP connections by the current IMAP server. This number accumulates until the server restart or reset by "counterutil." numFailedConnections counts connections abnormally terminated, including unsuccessful accepts and connections successfully accepted but which had an error later. An error message is logged when a connection failed with an expected error. You can check your IMAP log files for error messages such as the following: Unable to accept client connection: <error message> Socket error : <error message>
numFailedLogins	Number of failed system logins served by the current IMAP server.
numGoodLogins	Number of successful system logins served by the current IMAP server.

Disk Usage Statistics by Using counterutil

The command `counterutil -o diskusage` generates following information:

Table 37-4 counterutil diskusage Statistics

Suffix	Description
<code>diskusage.availSpace</code>	Total space available in the disk partition. The values are scaled to fit in the 4 byte counter. If you have a very large file system, the actual number will be divided by 1024 until it is small enough to fit in the 32-bit integer.
<code>diskusage.lastStatTime</code>	The last time statistic was taken.
<code>diskusage.mailPartitionPath</code>	Mail partition path.
<code>diskusage.percentAvail</code>	Disk partition space available percentage.
<code>diskusage.totalSpace</code>	Total space in the disk partition. The values are scaled to fit in the 4 byte counter. If you have a very large file system, the actual number will be divided by 1024 until it is small enough to fit in the 32-bit integer.

Server Response Statistics

The command `counterutil -o serverresponse` generates following information. This information is useful for checking if the servers are running, and how quickly they're responding.

Table 37-5 counterutil serverresponse Statistics

Suffix	Description
<code>http.laststattime</code>	Last time http server response was checked.
<code>http.responsetime</code>	Response time for the http.
<code>imap.laststattime</code>	Last time imap server response was checked.
<code>imap.responsetime</code>	Response time for the imap.
<code>pop.laststattime</code>	Last time pop server response was checked.
<code>pop.responsetime</code>	Response time for the pop.

Monitoring User Access to the Message Store

Oracle Communications Messaging Server provides the command, **imsconnutil**, which allows you to monitor user's message store access via IMAP, POP and http. You can also determine the last log in and log out of users. This command works on a per message store basis and will not work across message stores.

Note: Use of this function or other Messaging Server functions to monitor, read or otherwise access user's email may constitute a potential source of liability if used in violation of applicable laws or regulations or if used in violation of the customer's own policies or agreements.

This command requires root access by the system user (default: **mailsrv**), and you must set the configuration variables **local.imap.enableuserlist**, **local.http.enableuserlist**, **local.enablelastaccess** to 1.

To list users currently logged on via IMAP or any web mail client, use the following command:

imsconnutil -c

To list the last IMAP, POP, or Messenger Express access (log in and log out) of every user on the message store use:

imsconnutil -a

The following command does two things: 1) it determines whether the specified user is currently logged on via IMAP or Messenger Express or any client that connects via mshttp (note that this does not work for POP because POP users generally do not stay connected), and 2) it lists the last time the users have logged on and off:

imsconnutil -c -a -u user_ID

Note that a list of users can be input from a file, one user per line, using the following command:

imsconnutil -c -a -f filename

You can also specify a particular service (**imap** or **http**) using the **-s** flag. For example, to list whether a particular user ID is logged onto IMAP or not, use the following command:

imsconnutil -c -s imap -u user_ID

Note that the **-k** option may only work if IMAP IDLE is configured. For a complete description of the **imsconnutil** syntax, refer to "imsconnutil." Here is some example output:

```

$ imskonutil -a -u sorork
UID      IMAP last access    HTTP last access    POP last access
=====
ed  08/Jul/2003:10:49:05    10/Jul/2003:14:55:52  ---NOT-RECORDED---
$ imskonutil -c
IMAP
UID      TIME                AUTH                TO                FROM
=====
ed  17/Jun/2003:11:24:03    plain               172.58.73.45:193  129.157.12.73:2631
bil 17/Jun/2003:04:28:43    plain               172.58.73.45:193  129.158.16.34:2340
mia 17/Jun/2003:09:36:54    plain               172.58.73.45:193  192.18.184.103:3744
jay 17/Jun/2003:05:38:46    plain               172.58.73.45:193  129.159.18.123:3687
pau 17/Jun/2003:12:23:28    plaintext           172.58.73.45:193  192.18.194.83:2943
ton 17/Jun/2003:05:38:46    plain               172.58.73.45:193  129.152.18.123:3688
ani 17/Jun/2003:12:26:40    plaintext           172.58.73.45:193  192.18.164.17:1767
ani 17/Jun/2003:12:25:17    plaintext           172.58.73.45:193  129.150.17.34:3117
jac 17/Jun/2003:12:26:32    plaintext           172.58.73.45:193  129.150.17.34:3119
ton 17/Jun/2003:12:25:32    plaintext           172.58.73.45:193  192.18.148.17:1764
=====
10 users were logged in to imap.
Feature is not enabled for http.
-----

```

Message Archiving

This information describes archiving concepts for Oracle Communications Messaging Server. It does not provide instructions on how to set up an archiving system.

Topics:

- [Microsoft Exchange Envelope Journaling](#)
- [Archiving Overview](#)

Microsoft Exchange Envelope Journaling

Messaging Server is able capture sieve action to produce Microsoft Exchange's "envelope journaling" format. This format consists of a multipart MIME message where the first part contains envelope information in a semi-structured format and the second part is the actual message. This new format is produced by specifying a **:journal** option to capture:

```
capture :journal "trigger-address";
```

Exchange Journal Format Archiving for IMAP APPEND with LDAP Attributes

Support has been added to the archiving library to produce Microsoft Exchange Journal format archive messages. Note that this support extends to store compliance archiving of IMAP APPENDs as well as the archiving plugin.

In the case of the archiving plugin, the **STYLE** option accepts a value of 3, indicating that Microsoft Exchange Journal format messages should be produced. Additionally, two option file options have been added:

- **SOURCE_CHANNEL** (channel name; no default; required for Microsoft Exchange Journal format only)

The **SOURCE_CHANNEL** option specifies the name of the channel used to submit Microsoft Exchange Journal format messages. We recommend that you create a separate channel for this purpose so that such submissions are clearly identifiable in the logs.

- **DESTINATION** (string; no default; Microsoft Exchange Journal format only).

If set, this option specifies the address where Exchange Journal format archive messages are to be sent. If the option is not set archive messages are sent to the various capture attributes associated with the message's authorized sender, envelope from, and envelope recipient addresses.

In the case of store compliance archiving, there are three options:

- **store.archive.style** - Same semantics as the **STYLE** option file option.

- **store.archive.source_channel** - Same semantics as the **SOURCE_CHANNEL** option file option.
- **store.archive.destination** - Same semantics as the **DESTINATION** option file option.

Archiving Overview

A message archiving system saves all or specified incoming and outgoing messages on a system separate from Messaging Server. Sent, received, deleted, and moved messages can all be saved and retrieve in an archive system. Archived messages cannot be modified or removed by email users, so the integrity of incoming and outgoing messages is maintained. Message archiving is useful for compliance record keeping, but it is also useful for message store management. For example, some customers may use archiving to perform message back-up or to move older messages from more expensive message store storage to less expensive archive storage.

Archived messages can be accessed through a separate archiving software GUI client or through Messaging Server. If the messages are deleted from Messaging Server, then the archiving client can be used to search for and retrieve those deleted messages since archived messages are never deleted. Note, however, that archived messages are not stored in mailbox folders as they are in the Messaging Server.

The system can also be set up so that archived messages can be accessed from Messaging Server. For example, you can set up a system to archive messages over 2 years old. Instead of having message bodies reside in the message store, they would instead reside in the archive system. From the users standpoint, the message appears no different from a regular email message. The same header and subject information will appear (this is still stored in the message store storage), but the message body is downloaded from the archive server by the message store when needed. Thus, there may be a slight delay as messages are downloaded from the archive server. In addition, archived messages cannot be searched from the email client. Searching must be done from the archiving GUI.

Message Archiving Systems: Compliance and Operational

There are two types of archiving, compliance and operational. Compliance archiving is used when you have a legal obligation to maintain strict retrievable email record keeping. Selected email (selected by user(s), domain, channel, incoming, outgoing and so on) coming into the MTA is copied to the archive system before being delivered to the message store or the internet. Archiving can be set to occur either before or after spam and virus filtering.

Operational archiving is used for mail management purposes. For example:

- To reduce storage usage on the Messaging Server message store by moving less used (older) messages to an archiving system which uses lower cost storage.
- As an alternative for data backup.

Note that compliance and operational archiving are not exclusive. That is, you can set up your system so that it does both compliance and operational archiving.

Unified Messaging

With Oracle Communications Messaging Server, Oracle provides back-end software that lets you receive, store, and manage all types of messages--voice, fax, email, video--on one, off-the-shelf, cost-effective system.

The following sections describe Oracle's Unified Messaging solution:

- ["Using Messaging Server to Manage Unified Messaging"](#) is a white paper describing Oracle's unified messaging technology.
- ["Designing and Coding Your Unified Messaging Application"](#) describes how to design, configure, and code Messaging Server to support a unified messaging solution.

Using Messaging Server to Manage Unified Messaging

This document describes Oracle's solution to receive, store, and manage all types of messages--voice, fax, email, video--on one, off-the-shelf, cost-effective system. This paper is intended for telephony engineers and decision-makers interested in using standard Internet email software to manage voicemail, video, and fax communications.

For implementation details, see ["Designing and Coding Your Unified Messaging Application."](#)

What is the Challenge?

As broadband service providers compete to integrate internet communications (email and instant messaging) and voice communications (voice portal, voicemail, fax, and voice conferencing), they must reduce the costs of creating and maintaining those integrated services. To do this, telephone companies and their network equipment manufacturers (NEPs) must adopt open standards and use cost-effective, off-the-shelf storage technologies.

The Oracle Solution

Oracle's leadership in this area allows it to provide, with its partners, a solution that supports the convergence of these communication services, called unified messaging.

Unified messaging provides the following benefits:

- **Multi-modal access:** the ability to access different types of messaging services using different types of technology. For example, using the phone to access email messages or using the computer to access voices messages.

- A cost-effective, off-the-shelf technology for storing voice, email, video, and fax messages on the same system.
- The ability to manage all types of messages using the same administrative procedures. This includes expiring old messages, setting message space quotas, archiving, logging, usage profiles, and so on.
- Access to technologies like text-to-speech (TTS) and automated speech recognition (ASR).

Using open standards and best practices, Messaging Server offers a single back-end service that receives, stores, manages, and expires messages, no matter what their type.

Open Standards and Regulatory Requirements

In response to the need for open standards, engineering leaders have created the Voice Profile for Internet Mail (VPIM) specification, which defines the interfaces between voice front-ends and IMAP-based message stores. These standards, integrated into and implemented by Messaging Server, provide a substantial savings over traditional proprietary voice storage. Furthermore, unified messaging applications can take advantage of the efficient mailbox operations and system-performance knowledge provided by Messaging Server, leveraging the expertise and technology that have deployed hundreds of millions of Messaging Server mailboxes around the world.

Today's new legal and regulatory realities require an effective and unobtrusive mechanism for legal mail interception (LMI). Messaging Server enables service-provider personnel to apply the same LMI mechanisms already used on email to voice and fax interception and recovery. These mechanisms both capture messages on the network and archive messages stored on disk.

Architectural Overview of a Unified Messaging Application

The following sections define a high-level architecture for a unified messaging application using Messaging Server for communications storage, management, and notifications.



This architectural view follows the lifecycle of the message types (voice, fax, email), divided into the following stages:

1. An incoming message is deposited in the message store, and a message-waiting indicator is turned on.
2. The end user retrieves the message via the Telephone User Interface (TUI). The message-waiting indicator is turned off.
3. Alternatively, the end user retrieves the message via an IMAP messaging client (such as Thunderbird) or Convergence.
4. The Messaging Server message store administrator uses the Messaging Server's Mailbox Administration and Operations Management (MAOM) functions to administer and expire the message.

The sections that follow take snapshots of the communications workflow at each one of these stages.

Message Deposit

The following figure and state diagram show the message flow of a voice message or fax message when it first arrives and is deposited in the message store.

Figure 40–1 Message Deposit Function of Unified Messaging for Telecommunications**Figure 40–2 Message Deposit Function--State Diagram**

1. Receive the call.

The Public-Switched Telephone Network (PSTN) receives a phone call or FAX intended for the end user and passes it on to the Media Application Server.

A Media Application Server is a third-party system that services voice and FAX callers. It is similar to a home answering machine in that it picks up calls, plays a pre-recorded message, handles touch-tone interactions, and retrieves and stores voice messages and so on. In many cases, the Media Application Server sits on a voice trunk with many ports; some servers support thousands of ports. A port is equivalent to one telephone line.

2. Profile the call with LDAP data.

Using the end user's phone number, the Media Application Server looks up the LDAP entry for that number in Directory Server. Directory Server returns the LDAP profile to the Front-end Server.

The types of profile information retrieved can include:

- Status of the mailbox: Available, Disabled, Vacation
- Allowed services
- Current greeting for the user.

3. Retrieve and play voicemail greeting.

The Media Application Server retrieves the greeting voice file from the message store and plays it to the caller. For example: "Hello, you've reached..." It prompts the caller to leave a message.

4. Record the message and release caller.

The Media Application Server records the caller's message in binary file format, usually WAV format for voice and TIFF for FAX data. The caller hangs up.

5. Encode the voice/FAX message.

The Media Application Server creates an email message, attaches the caller's encoded message file to it, and addresses it to the user's mailbox.

- It creates an RFC2822 email message.
- It attaches the voice message file as a binary attachment (base64). This format conforms to the Voice Profile for Internet Mail (VPIM) standard.
- It addresses the message to the user's mailbox, using a standard email address--for example, 555-555-5555@example.com.

The voice message is now an email message with a large attachment.

6. Send email to message store.

The Media Application Server sends the email message via SMTP to the Messaging Server. The Messaging Server Message-Transfer Agent (MTA) accepts the email and routes it to the user's mailbox in the message store.

The message store, a component of Messaging Server, stores and manages user mailboxes.

7. Generate a new-message event notification.

This is to notify the system that a new message has arrived for this mailbox and is available for retrieval. The notification, and the actions taken in response to the notification--for example, turning on the message wait indicator--must be implemented by the customer.

Message Queue provides the infrastructure for producing and distributing event notifications.

Message Retrieval via Telephone User Interface

The following figures show the message flow and data states of a voice or fax message as the end user picks it up (listens to it or has the fax machine print it).

Figure 40-3 Message Retrieval via Telephone User Interface





Figure 40-4 Message Access and Retrieval via Telephone User Interface--State Diagram



These diagrams illustrate the following actions:

1. User dials in and system profiles user mailbox.

In this implementation of the Media Application Server, the end user dials into the voice mail system. Using the end user's phone number, the Telephone/FAX Media Application Server looks up the LDAP entry for that number in Directory Server. Directory Server with Access Manager authenticates the user.

After the authentication, the Media Application Server then looks at the user's Message Store server from the user's LDAP mailHost attribute. The system prompts the user for the password.

2. User enters password.

3. System retrieves new mailbox summary.

The Media Application Server performs an IMAP connection to the mailbox owner's Message Store. The IMAP connection may pass through a component called the Messaging Multiplexor (MMP) which routes connection requests to the appropriate message store. It then opens the user's folders to check for number of new messages (for example, number of new voicemail, fax and email messages).

The Media Application Server then plays the number of new messages of each type to the user over the phone.

Note: Steps 3 and 4 are separate procedures that are executed at the same time.

4. System retrieves and plays messages.

The mailbox caller selects either voicemail or fax mailbox. Using IMAP, the Media Application Server retrieves the message from its message store. The Media Application Server then plays each Message header. For voicemail, it plays the message.

- To play the voice message, the Media Application Server retrieves the RFC822 message. It decodes the attachment and plays the audio file (typically .WAV or mp3) over the phone line.
- To play fax messages, the Media Application Server sends information about the fax (for example, date/time, caller number, number of pages, urgency, etc.). Typically, mailbox owners will forward these fax messages to fax machines nearby to them.
- To play email messages, the Media Application Server sends info about the email (e.g. sender, subject, time/date, urgency). If text-to-speech has been implemented, the server will attempt to "read" the message to the caller. Some implementations allow the caller to forward the fax and the attachments to a nearby fax machine. After playing the message, the caller can delete, replay or forward the message.

The message itself remains in the user's mailbox in the message store, but the status of the message flag is changed to "seen/read."

5. User replays, deletes, forwards or archives the message.

6. System generates a message-read or message-deleted event notification.

This is to notify the system that the message has been read or deleted. The notification, and the actions taken in response to the notification--for example, turning off the message wait indicator--must be implemented by the customer.

Message Queue provides the infrastructure for producing and distributing event notifications.

Message Retrieval via PC


Messaging Server provides two ways to access messages through a PC: through an IMAP client such as Thunderbird or Outlook (with Connector) and through the , a web-based communication client, which uses an HTTP connection.

These will be described in separate sections.

Message Retrieval Through an IMAP Client

The following figure shows the message flow and data state of a voice or fax message as it is retrieved through an IMAP client.

Figure 40–5 IMAP Client Message Retrieval--Component and State Diagram



This diagram illustrates the following actions:

1. **User opens IMAP Client which connects to Messaging Server.**
2. **User enters the password and is authenticated by the MMP.**

The MMP finds the appropriate message store and sets up a direct connection between the store and the client.

3. **Message store sends message header information to the IMAP client.**
4. **User clicks the message header to open.**

Client issues an IMAP FETCH command and the message store returns the RFC822 email message with the MIME attachment containing the voice/FAX data. The format of the data could be in .WAV, MP3, TIFF or JPEG.

5. **User clicks the attachment.**

IMAP client launches the media player on the PC, then strips off the wrappers and encoding in the attachment, and sends the resulting binary file to the media player.


6. **User deletes, forwards, or archives the message.**

IMAP client sends command to the message store to delete, forward or archive the message.

7. **Message store generates an event notification indicating the changed status (deleted/forwarded/archived) of the message.**

Message Retrieval Through Convergence

The following figure shows the message flow and data state of a voice or fax message as it is retrieved through the Convergence client.

Figure 40–6 Convergence Message Retrieval--Component and State Diagram

This diagram illustrates the following actions:

1. **User logs in with password to the Convergence Server through a browser and retrieves the message headers.**

Convergence Server makes a HTTP request to the Webmail Server for the message headers. The Webmail Server requests and retrieves the headers via IMAP and returns it to Convergence in HTTP where the user can view it.

The Webmail Server translates HTTP commands to IMAP, and IMAP commands to HTTP. Note that the message store stores data in 7-bit format, and that http transfers data in 8-bit binary format.

2. **User clicks the message header to open.**

Again, Convergence Server makes a HTTP request to the Webmail Server for the message. The Webmail Server requests and retrieves the message via IMAP and returns it to Convergence in HTTP where the user can view it. The message does not contain the MIME attachment, but contains a link to the attachment.

3. **User clicks the attachment.**

Again, an HTTP request is made to the Webmail Server, which translates the request to IMAP and sends it to the message store.

4. **Retrieve the MIME attachment.**

The Webmail Server retrieves the MIME attachment, strips off the wrappers and encoding in the attachment, and sends the resulting base64 binary file via Convergence to the browser which launches media player.

5. **The media player plays and displays the attachment.**

6. **User deletes, forwards, or archives the message.**

Convergence sends command to the message store to delete, forward or archive the message.

7. **Message store generates an event notification indicating the changed status of the message (deleted/forwarded/archived).**

Designing and Coding Your Unified Messaging Application

This document describes how to design and configure Messaging Server to implement a unified messaging application.

For a high-level description of how Messaging Server features can integrate and support a unified messaging solution, see the following white paper:

- [Using Messaging Server to Manage Unified Messaging](#)

The following sections point out key areas you'll need to design and code to integrate the Messaging Server back-end with your UM application:

- [Planning the Message-Type Configuration](#)
- [Coding and Configuring Your UM System](#)
- [Mailbox Administration and Operations](#)
- [Delivering Notifications for Message Types](#)
- [Additional Unified Messaging Support Features](#)

Planning the Message-Type Configuration

To administer messages of different types, all components of the UM system must use the same message-type definitions and the same header fields to identify the messages.

Before you configure Messaging Server to support message types, you must

- Plan which message types you intend to use
- Decide on the definition for each message type
- Decide which header field to use

For example, if the application includes phone messages, you can define this message type as "multipart/voice-message" and use the Content-Type header field to identify message types.

You would then configure the Media Application Server to add the following header information to each phone message to be delivered to the message store:

```
Content-Type: multipart/voice-message
```

Next, you would configure the MTA and/or message store to recognize the **multipart/voice-message** message type.

Coding and Configuring Your UM System

This section takes a closer look at the message life cycle described in earlier sections. Here we focus on the stages in the life cycle where you must design, code, and configure components of the UM system. We give you an idea of the Media Server coding and Messaging Server configuration that you will perform to implement your UM system.

1. **Unanswered call is routed to the Media Server.**

Before the Media Server sends a message to Messaging Server, it must check the status of the user's mailbox to ensure that the user's mailbox is not full or busy.

The Media Server must be coded to check the `mailuserstatus` attribute in the user's entry in the directory server to see, for example, if the user's mailbox is full or not. You also must configure quota enforcement on the Message Store to enable this feature.

For an introduction to quota enforcement by message type, see "[Administering Quotas for Message Types](#)." For details about configuring quotas for the message store, see "[Managing Message Store Quotas](#)."

2. The Media Server encodes the voicemail as a binary attachment to an email message, which it sends to the MTA. The message is labeled by type.

(Of course, you must code the Media Server to perform the voicemail-to-email state transformation.)

When a message comes into the MTA via the Media Server, the first thing that must be done is to add a **Content-type** header to the message. The value of this header identifies the message type.

In this way, the messages can be managed according to their type. For example, voicemail types may have different quotas than text types.

You must define each message type with a unique identifier such as **multipart/voice-message**.

Content-type headers can be inserted in the message by the Media Server or by the MTA.

Typing by Media Server. When the Media Server constructs the email message to send to the MTA, it can be coded to add the **Content-type** header with the appropriate message-type value. For example, a voicemail message could require that **Content-type: multipart/voice-message** be added to the message.

Typing by MTA. The alternative is to set up your system so that the MTA adds the **Content-type** header. This can be done by configuring separate, non-standard ports for the various message types.

Typically, email comes into the MTA via port 25. Thus, for example, you can configure text email to go to the standard port 25, voicemail to port 225, FAX to port 325. You need to configure the Media Server to send the message to the appropriate port. Also, you need to configure the MTA to append the appropriate **Content-type** header to messages arriving at each port.

3. The MTA deposits the message in the message store, and an IMAP flag identifying the message type is appended to the message.

The message store reads the **Content-type** header to identify the message type.

You can configure the MTA or the message store to append a message-type flag to the message. You must define unique values for each message-type flag.

Messaging Server presents the message-type flag as a user flag to IMAP clients. (This flag cannot be modified by end users.)

Mapping the message type to a user flag allows mail clients to use simple IMAP commands to manipulate messages by message type.

4. Client retrieves mailbox status and displays the message together with its type.

From the user's perspective, the client email software may display an icon indicating the type of each message (if the client supports this feature). For example, a phone icon appears next to a voicemail message.

The IMAP SEARCH command, using the message-type flag as a keyword, retrieves a count of each type of message. The IMAP FETCH command retrieves the message headers with the message status and message-type flag name.

For examples, see ["Sample IMAP Sessions Using Message-Type Flags."](#)

5. **User clicks on the voicemail attachment, and the voicemail is played by a media player on the PC.**

For details, see steps 2 and 3 in ["Message Retrieval Through Convergence."](#)

6. **Message store generates a notification indicating the changed status of the message.**

If the status of a messages changes, the message store can generate a notification that can be retrieved by the email client or the Media Server. A notification can deliver a count of the messages in a user's mailbox for each message type and for each change in message status.

For example, a notification can deliver a count of all new (unread) voicemail messages and all new text messages in a user's mailbox. When the user listens to a voicemail, another notification can be generated indicating all voicemail messages that have been read and all text messages that have been read. In this case, the number of new voicemail messages is reduced by one.

You need to configure Messaging Server to determine how and when to generate notifications. Also, you need to write Media Server code to retrieve the notification and take appropriate action. For example, when a new message arrives in the store, you may wish to send an indicator to the customer's phone. For details, see ["Delivering Notifications for Message Types."](#)


Mailbox Administration and Operations

The Messaging Server message store can be configured to identify and manage message types.

The customer does not have to maintain different message types in individual mailbox folders. The message store can identify a message type, no matter where the message is stored. Thus, you can store heterogeneous message types in the same folder.

The following figure illustrates how an incoming message is identified by its type:

Figure 40–7 Message Management by the Messaging Server Message Store



In this diagram, the message store identifies the message type of an incoming message by reading the **Content-type** header. In addition, the message store appends an IMAP

flag identifying the message type (if the IMAP flag has not already been appended by the MTA).

Once message types have been configured, the message store lets you

- Set flags that allow IMAP commands to fetch and search for information about message types
- Configure quota roots that apply to each message type
- Write expire rules to expire and purge messages according to message type

The rest of this section describes the following topics:

- [Sample IMAP Sessions Using Message-Type Flags](#)
- [Administering Quotas for Message Types](#)
- [Expiring Messages by Message Type](#)

Sample IMAP Sessions Using Message-Type Flags

This section describes sample IMAP sessions using IMAP FETCH and IMAP SEARCH.

Example 1: IMAP FETCH Session

The following IMAP session fetches messages for the currently selected mailbox:

```
2 fetch 1:2 (flags rfc822)
* 1 FETCH (FLAGS (\Seen text) RFC822 {164}

Date: Wed, 8 July 2006 03:39:57 -0700 (PDT)
From: bob.smith@example.com
To: john.doe@example.com
Subject: Hello
Content-Type: TEXT/plain; charset=us-ascii

* 2 FETCH (FLAGS (\Seen voice_message) RFC822 {164}

Date: Wed, 8 July 2006 04:17:22 -0700 (PDT)
From: sally.lee@example.com
To: john.doe@example.com
Subject: Our Meeting
Content-Type: MULTIPART/voice-message; ver=2.0

2 OK COMPLETED
```

In the preceding example, two messages are fetched, one text message and one voice mail.

The **Content-type** header fields identify the message types. The message-type names are displayed as they were received in the incoming messages.

Example 2: IMAP SEARCH Session

The following IMAP session searches for voice messages for the currently selected mailbox:

```
3 search keyword voice_message
* SEARCH 2 4 6
3 OK COMPLETED
```

In the preceding example, messages 2, 4, and 6 are voice messages. The keyword used in the search, `voice_message`, is the flag name defined for voice messages.

Administering Quotas for Message Types

When you set a quota for a message type, you include that value in a quota root. A quota root specifies quotas for a user.

You can specify the following quotas for a user's mailbox tree:

- Quota values for specific folders in the user's mailbox
- Quota values for specific message types such as voice mail or text messages. A message type quota applies to messages of that type in all folders in the user's mailbox.
- A default quota value that applies to all folders and message types in the user's mailbox that are not explicitly assigned quotas.


Quotas can be configured for the number of messages allowed and for the maximum amount of disk storage used.

Example of a Message-Type Quota Root

Suppose a customer wants to configure separate quotas for text messages and voice messages in each user's mailbox. Yet another quota is to be set for the user's Archive folder.

The following figure illustrates this example:

Figure 40–8 Administering Quotas in the Message Store



This example sets the following quotas for a user:

- The storage quota for the Archive folder is 100M
- The storage quota for text message types is 10M
- The message quota for text message types is 2000
- The storage quota for voice message types is 10M
- The message quota for voice message types is 200
- The default mailbox storage quota is 40M

- The default mailbox message quota is 5000

This quota root permits greater storage in the Archive folder (100 M) than in all the other folders and message types combined (60 M). Also, no message limit is set for the Archive folder; in this example, only storage limits matter for archiving. The message types have both storage and number-of-message quotas.

The message-type quotas apply to the sum of all messages of those types, whether they are stored in the Archive folder or in any other folder.

The default mailbox quotas apply to all messages that are not text or voice message types and are not stored in the Archive folder. That is, the message-type quotas and Archive quota are not counted as part of the default mailbox quotas.

For example, the default mailbox quota would apply to a message that arrives in the user's INBOX without a Content-Type header and message-type definition. When that message is archived, the Archive folder storage quota would apply.

Guidelines for Specifying Multiple Quota Values

The following guidelines apply when you assign multiple quota values for a user:

- Quotas do not overlap. For example, when there is a quota for a particular message type or folder, messages of that type or messages in that folder are not counted toward the default quota. Each message counts toward one and only one quota.
- The total quota for the whole user mailbox equals the sum of the values of all the quotas specified by default, type, and folder.
- Message-type quotas take precedence over folder quotas. For example, suppose one quota is specified for a user's memos folder and another quota is specified for voice messages. Now suppose the user stores eight voice messages in the memos folder. The eight messages are counted toward the voice-mail quota and excluded from the memos folder quota.

Sample IMAP Session Returning Quota Root Values

When you run the `getquotaroot` IMAP command, the resulting IMAP session displays all quota roots for the user's mailbox, as shown here:

```
1 getquotaroot INBOX
* QUOTAROOT INBOXuser/joe user/joe/#text user/joe/#voice
* QUOTA user/joe (STORAGE 12340 20480 MESSAGE 148 5000)
* QUOTA user/joe/#text (STORAGE 1966 10240 MESSAGE 92 2000)
* QUOTA user/joe/#voice (STORAGE 7050 10240 MESSAGE 24 200)

2 getquotaroot Archive
* QUOTAROOT user/joe/Archive user/joe/#text user/joe/#voice
* QUOTA user/joe/Archive (STORAGE 35424 102400)
* QUOTA user/joe/#text (STORAGE 1966 10240 MESSAGE 92 2000)
* QUOTA user/joe/#voice (STORAGE 7050 10240 MESSAGE 24 200)
```

Expiring Messages by Message Type

The expire and purge feature allows the customer to move messages from one folder to another, archive messages, and remove messages from the message store, according to criteria the customer defines in expire rules. These tasks are performed with the `imexpire` utility.


Because the **imexpire** utility is run by the administrator, it bypasses quota enforcement.

The customer can write expire rules so that messages of different types are expired according to different criteria.

The expire feature is extremely flexible, offering many choices for setting expire criteria. This section describes one example in which text and voice messages are expired according to different criteria.

The following figure illustrates an example of expiring messages by type:

Figure 40–9 Expiring Messages by Message Type



In this example, text messages and voice mail are expired in different ways, and they follow different schedules, as follows:

- Text messages are moved from a user's inbox to the user's Archive folder one year after they arrive in the message store.
- Voice mail is moved from the inbox to the OldMail folder after two weeks. If the user saves a voice message, the saved date is reset, and the message is moved two weeks after the new date.
- Voice mail is moved from the OldMail folder to the Trash folder after 30 days. The user also can save a voice message in the OldMail folder, which postpones the removal of the message for another 30 days after the new saved date.
- Messages of all types are discarded seven days after they are moved to the Trash folder.

The expire rules move voice mail to Trash automatically. Text messages are moved to Trash when a user deletes them.

Sample Rules for Expiring Different Message Types

You can implement the example described in this section by writing the following expire rules:

```
TextInbox.folderpattern: user/%/INBOX
TextInbox.messageheader.Content-Type: text/plain
TextInbox.messagedays: 365
TextInbox.action: fileinto:Archive

VoiceInbox.folderpattern: user/%/INBOX
VoiceInbox.messageheader.Content-Type: multipart/voice-message
VoiceInbox.savedays: 14
VoiceInbox.action: fileinto:OldMail

VoiceOldMail.folderpattern: user/%/OldMail
VoiceOldMail.messageheader.Content-Type: multipart/voice-message
VoiceOldMail.savedays: 30
VoiceOldMail.action: fileinto:Trash

Trash.folderpattern: user/%/Trash
Trash.savedays: 7
Trash.action: discard
```

Delivering Notifications for Message Types

Messaging Server, together with Message Queue, can produce notifications that deliver status information about messages of different types, such as voice mail, text messages, fax data, and image data.

For example, suppose a new phone message arrives in a user's mailbox, as described in "[Message Deposit](#)."

You can configure Messaging Server to generate a new-message notification for the Message Queue service. You can *also* configure Messaging Server to identify particular message types, including voice mail.

Now, when the voicemail is deposited in the user's mailbox, the message store triggers a notification that says, essentially:

- "This user has a new message."
- "The new message is voicemail."

After Messaging Server delivers the new-message notification to the Message Queue service, Message Queue sends it to a consumer (client interface), which filters and delivers the message to its destination.

You can write your Message Queue client program to interpret notification messages by message type and deliver status information about each type. In this example, the client program, recognizing the new message is voice mail, could trigger the Message Wait Indicator to turn on the "New Messages" light on the end-user's phone.

Now suppose new messages of different types--say, voice messages and email--arrive in the user's mailbox.

Once you have configured message types, a new-message notification carries data that counts the number of each type--in this case, the number of new voice messages and new email (text) messages. Your Message Queue client program can deliver the count by message type to the Media Application Server, which would notify the user that there are, for example, seven new voice mail messages and four new text messages in the user's cell phone inbox.

Notifications for Particular Message States

The following notifications can be generated when a message changes state. For example, when a user reads a message or listens to voicemail, a **ReadMsg** notification can be generated. These notifications can carry information that tracks particular message types:

Table 40–1 Notification Message Descriptions

Notification Message	Description
NewMsg	New message was received by the system into the user's mailbox. Can contain message headers and body.
UpdateMsg	Message was appended to the mailbox by an IMAP operation. For example, the user copied an email message to the mailbox. Can contain message headers and body.
ReadMsg	Message in the mailbox was read. (In the IMAP protocol, the message was marked Seen .)
TrashMsg	Message was marked for deletion by IMAP or HTTP. The user may still see the message in the folder, depending on the mail client's configuration. The messages are to be removed from the folder when an expunge is performed.
DeleteMsg	Messages marked as Deleted are removed from the mailbox. This is the equivalent to IMAP expunge
PurgeMsg	Message expunged (as a result of an expired date) from the mailbox by the server process imexpire. This is a server side expunge, whereas <code>{{DeleteMsg}}</code> is a client side expunge. This is not a purge in the true sense of the word.
OverQuota	Operation failed because the user's mailbox exceeded one of the quotas (diskquota , msgquota). The MTA channel holds the message until the quota changes or the user's mailbox count goes below the quota. If the message expires while it is being held by the MTA, it will be expunged.
UnderQuota	Quota went back to normal from OverQuota state.

How Do You Implement Notifications for Message Types?

To implement notifications that count by message type, you need to do these things:

- Configure the message store to identify message types
- Configure a Messaging Server component, the JMQ notification plug-in, to produce notifications that identify message types
- Write a Message Queue client that retrieves, filters, and delivers the notification
- Design/write your Unified Messaging system (for example, the Media Application Server) to receive the notification and deliver it to the end user or other destination

Configuration Details

So far, you've seen the changes in message state that can trigger notifications. But how do the notifications carry information about message types?

We need to explain a few configuration details to show how these components work together. The following sections discuss the first two items listed above--the Messaging Server components.

Configuring the Message Store to Recognize Message Types

You use the Messaging Server **msconfig** or **configutil** utility to configure two options that enable the message store to identify message types:

- **store.message.enable** (same for both Unified Configuration and legacy configuration)
- **store.message.mtindex:x.contenttype** (Unified Configuration)
or
store.message.x (legacy configuration)

First, you set the **store.message.enable** option to on (**-v 1**) to enable message types.

Next, you define one **store.message.mtindex:x.contenttype**(Unified Configuration) or **store.message.x** (legacy configuration) option for each message type. For example, to identify four message types, you define four iterations of this option with four different values. You do these things:

- Set an integer value for variable *x*.
- Specify a text string that is the value of the message type used in your Unified Messaging system with the **Content-Type** header.

For example, to define a text message, you can enter:

```
configutil -o store.message.1 -v text/plain
```

To define a voicemail type, you can enter:

```
configutil -o store.message.2 -v multipart/voice-message
```

Now the message store can identify any message type with a **Content-Type** header value of **text/plain** or **voice-message**.

Also, the message store will identify **message.1** with text messages and **message.2** with voicemail. (You'll need to know this when you see how notification properties carry information about message types.)

Note: You can configure other **configutil** options to define IMAP flag names, quota root names, and even an alternate message header name (other than **Content-Type**).

For details, see the discussion on managing message types in the *Sun Java System Messaging Server Administration Guide*.

Configuring the JMQ Notification Plug-In to Generate Messages

You use the **configutil** utility to define options that configure the Messaging Server JMQ notification plug-in. The plug-in can produce a notification whenever a message changes state.

For each state change that should trigger a notification, you define a **configutil** option. For example, to enable notifications for new messages, you enter:

```
configutil -o local.store.notifyplugin.jmqnotify.NewMsg.enable -v 1
```

where **jmqnotify** is the name of the plug-in and **-v 1** enables notifications for this message.

To generate notifications for messages that the end user has read, you define this option:

```
local.store.notifyplugin.jmqnotify.ReadMsg.enable
```

and so on.

Note: To fully configure a JMQ notification plug-in, you must define several other **configutil** options. For details, see the discussion on configuring the JMQ Notification Plug-in to produce message for message queue in the *Sun Java System Messaging Server Administration Guide*.

How the Message Store and JMQ Notification Plug-in Work Together

Once you configure the message store's message-type feature and the JMQ notification plug-in, *both* components recognize and carry information about message types.

We've already discussed how the message store can enforce quotas and expire messages by message type.

Now, when a message changes state, the message store can generate a notification *and* the JMQ notification plug-in can automatically recognize the message type as well as the changed state.

Notification Properties for Message Types

Every notification carries additional information defined in properties. Different properties are present for different messages. For example, a **NewMsg** notification indicates the IMAP **uid** of the new message.

If message types are configured, the following properties are carried with notifications. These properties deliver a count of the messages in a particular state for each message type you have defined:

- **numMsgs $_{nn}$**
- **numSeen $_{nn}$**
- **numDeleted $_{nn}$**
- **numSeenDeleted $_{nn}$**

Suppose a new-message (**NewMsg**) notification is generated.

The Messaging Server JMQ notification function counts the number of new messages currently in the mailbox, by message type. Instead of sending one count with the **NewMsg** notification, an array specifying the count for each message type is sent.

The message-specific count is carried in the **numMsgs $_{nn}$** property and delivered with the notification.

The message-type number ($_{nn}$) identifies a particular type. For example, you can configure message type 2 (**store.message.2**) to identify voice messages, message type 3 (**store.message.3**) to identify text messages, and so on.

For **ReadMsg** and **TrashMsg** notifications, the number of messages seen (**numSeen $_{nn}$**) and the number marked as deleted (**numDeleted $_{nn}$**) are also counted by message type.

Table 40–2 describes these properties:

Table 40–2 Notification Properties for Message Types

Property	Data Type	Description
NumMsgsnm	MQInt32	<p>The total number of messages now in the mailbox, specified for each message type. If message types are configured, a numMsgsnm property carries a count for each message type nm.</p> <p>The numMsgs property is always sent; it counts the total number of all messages in the mailbox, including all types.</p> <p>For example, if 20 messages are currently in the mailbox, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification:</p> <pre>numMsgs=20 numMsgs3=10 numMsgs16=7</pre>
NumSeennm	MQInt32	<p>The total number of messages in the mailbox marked as seen (read), specified for each message type. If message types are configured, a numSeennm property carries a count for each message type nm.</p> <p>The numSeen property is always sent; it counts the total number of all messages marked as seen, including all types.</p> <p>For example, if 20 messages are marked as seen, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification:</p> <pre>numSeen=20 numSeen3=10 numSeen16=7</pre>
NumDeletednm	MQInt32	<p>The total number of messages in the mailbox marked as deleted, specified for each message type. If message types are configured, a numDeletednm property carries a count for each message type nm.</p> <p>The numDeleted property is always sent; it counts the total number of all messages marked as deleted, including all types.</p> <p>For example, if 20 messages are marked as deleted, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification:</p> <pre>numDeleted=20 numDeleted3=10 numDeleted16=7</pre>
NumSeenDeletednm	MQInt32	<p>The total number of messages in the mailbox marked as seen (read) and marked as deleted, specified for each message type. If message types are configured, a numSeenDeletednm property carries a count for each message type nm.</p> <p>The numSeenDeleted property is always sent; it counts the total number of all messages marked as seen and deleted, including all types.</p> <p>For example, if 20 messages are marked as seen and deleted, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification:</p> <pre>numSeenDeleted=20 numSeenDeleted3=10 numSeenDeleted16=7</pre>

Additional Unified Messaging Support Features

The following features have been added to Messaging Server 7.0 and may provide additional support for implementing Unified Messaging Solutions.

Set IMAP Flag Based on Header Value at Delivery

The IMAP flag setting is done through the **imap4flags** sieve extension. (Testing header values is, of course, a basic sieve capability.) This feature is fully specified in RFC 5232. The goal is to be able to represent compound message context states, such as **URGENT**, **VOICEMAIL**, **BROADCAST FAX** and so on via numeric IMAP Flags.

To apply these actions to multiple users, you probably want to make use of system, source channel, and destination channel, or perhaps domain sieve rather than user sieves.

Implementation Description: The LMTP and SMTP processes allow IMAP flags to be set upon delivery based on header value. Upon insertion of a message into the message store, a sieve rule in conformance with SIEVE-IMAPFLAGS-05 will modify the IMAP flag to represent the context of the message with an integer value from 1 to 32. The JMS will publish a notification event consequent to this event to the message queue.

In deployment, a SIEVE rule will perform a logical AND operation resulting in an IMAP flag and **X-HEADER** numeric value which would represent a context via a particular compound operation (for example, 17 to represent **URGENT** and **EMAIL** and so on). **EDITHEADER-09** is supported, which, with configuration changes, will allow this to be used during SMTP delivery to allow devices without the capability to construct an RFC3458 header to construct one via an **ADDHEADER** rule, which in turn will allow preprocessing of message type identification prior to arrival in the final mailbox.

Modifications to IMAP Commands to Provide Message Counts

Messaging Server provides quick message counts and message states to clients for improving response time. The IMAP **STATUS** and **SEARCH** commands have been modified to provide return value data representing counts for [RFC3458] message types.

IMAP Unauthenticate

The IMAP command **UNAUTHENTICATE** is supported. This will be advertised by the **XUNAUTHENTICATE CAPABILITY** and response. When invoked, the command will return the connection to an unauthenticated state, where **AUTHENTICATE** can be invoked without creating another connection to enable its reuse. This provides more effective use of IMAP pooling techniques via reuse of Network Sockets. The requirements of the IMAP **AUTHENTICATE** mechanism to create a new connection is bypassed and thereby enables connection reuse.

Modify IMAP APPEND to bypass quotas

Administrative users with appropriate privileges shall be able to bypass quota enforcement when they append messages to mailboxes with the IMAP **APPEND** command. The configuration option is **local.imap.adminbypassquota** which, when enabled, will bypass quota enforcement. Messages will be added to quota usage. Messages will not be rejected when the mailbox has exceeded its quota. Over quota warnings will still be delivered.

SMTP Future Release

Mail clients can indicate a future time for a message to be released for delivery up to one calendar year in advance of the current date. This support for RFC4865 has been added. The maximum values specified in the RFC for the option **HOLDFOR** shall be supported (+999999999 seconds). Equivalent **HOLDUNTIL** timestamp values shall be supported. This support will be enabled by placing the **futurerelease** channel option on the source channel used for initial message submission. The keyword shall take a single integer argument: the maximum number of seconds a message can be held.

Care should be used when enabling future release since it allows messages to be in effect stored in the MTA's queues. Future release should only be used for channels handling initial message submission and authentication should be required.

Messaging Server Command-Line Reference

This information describes the Oracle Communications Messaging Server **configure**, **configxml**, and **msconfig** commands that you use to manage a Unified Configuration. You need to be logged in either as **root** or **mailsrv** to run these commands. These commands are located by default in the *MessagingServer_home/bin* directory.

See "[Overview of Messaging Server Unified Configuration](#)" for a description of Unified Configuration.

Topics:

- [configure Command](#)
- [configxml Command](#)
- [msconfig Command](#)

configure Command

The **configure** command creates an initial runtime configuration for Messaging Server. It gives you a base working configuration from which you can make your specific customizations. The command is only meant to be run once. Subsequent running of this command overwrites the existing configuration. To modify your initial runtime configuration, use the "[msconfig Command](#)."

Syntax

```
configure [<options>]
```

Options

[Table 41–1](#) shows the options for the **configure** command:

Table 41–1 Options for configure Command

Option	Description
--debug	Provides additional debugging (primarily for LDAP operations)
--help,-?	Shows this help
--ignoreSendmail	Does not disable sendmail after configuration
--ldapport=port	Specifies an LDAP port, if you want to use other than port 389
--ldif	Does not modify LDAP directory, just writes LDIF

Table 41-1 (Cont.) Options for configure Command

Option	Description
<code>--noldap</code>	Runs without LDAP present (statefile only)
<code>--novalidate</code>	Skips most validation of user inputs
<code>--noxml</code>	Generates legacy configuration (does not use XML-based Unified Configuration); can also be used to replace a Unified Configuration with a freshly generated legacy configuration (fresh installation of Messaging Server, not an upgrade where the <code>configtoxml</code> command was run)
<code>--saveState=filename</code>	Specifies file where state information is saved
<code>--ssl</code>	Requires SSL when configuring LDAP
<code>--state=filename</code>	Uses silent state file to configure product
<code>--version,-V</code>	Shows product version
<code>--xml</code>	Generates Unified Configuration (XML)

Examples

- To configure an initial Unified Configuration:

```
configure
```

- To configure an initial legacy configuration:

```
configure --noxml
```

Note: A statefile can use the `XMLCONFIG=0` or `XMLCONFIG=1` option to specify a legacy or Unified Configuration, respectively. The `--xml` and `--noxml` command options override what is specified in the statefile.

configtoxml Command

The `configtoxml` command converts a legacy configuration to a Unified Configuration.

Syntax

```
configtoxml [<options>]
```

Options

Table 41-2 shows the options for the `configtoxml` command:

Table 41-2 Options for configtoxml Command

Option	Description
<code>-32,-64</code>	Installation is 32-bit (<code>-32</code>) or 64-bit (<code>-64</code>) Messaging Server. Default is 64-bit when the installation type cannot be inferred from the <code>SERVERROOT</code> or <code>ConfigRoot</code> environment variables or from the location of this script.

Table 41–2 (Cont.) Options for configtoxml Command

Option	Description
-f --force	<p>Ignores safety checks, enables running as non-root and permits overwriting of any pre-existing Unified Configuration files.</p> <p>Caution: Using this option may result in a non-functioning configuration. The restricted.cnf file must always be owned by root.</p>
-h --help	Shows this help.
-i <i>INSTANCE</i>	Inserts instance name in the generated configuration files. The default is ims .
-l <i>DIR</i> --location <i>DIR</i>	<p>Reads the legacy configuration files from the specified directory. The default to use is determined by the following:</p> <ol style="list-style-type: none"> 1. The ConfigRoot environment variable, if defined 2. The SERVERROOT/config/ path, if the SERVERROOT environment variable is defined 3. Location of ConfigRoot, determined from the location of this script 4. The OS-specific default path
-n --noactive	<p>Does not generate an active configuration and does not move the legacy configuration files to the ConfigRoot/legacy-config/ directory. The generated Unified Configuration files have the names config.xml, xpass.xml, and restricted.cnf, and are written to ConfigRoot. This option cannot be used in conjunction with the --output or --undo options.</p>
-o <i>CONFIG-FILEPASSWORD</i> <i>-FILERESTRICTED-FILE</i> --output <i>CONFIG-FILEPASSWORD</i> <i>-FILERESTRICTED-FILE</i>	<p>Directs the Unified Configuration file output to the designated files. By default, the files config.xml, xpass.xml, and restricted.cnf are written to the ConfigRoot or SERVERROOT/config/ directory. This option cannot be used in conjunction with the --noactive or --undo options.</p>
-r <i>ROLE</i> --role <i>ROLE</i>	Inserts the role name in the generated configuration files. The default is ims .
-y --yes	Pre-answer any confirmation questions with a "yes" response so that this script can be run without user intervention.
-u --undo	Removes any active Unified Configuration files and restores any legacy configuration files.

Key environment variables for this command are:

- **SERVERROOT** (The default is **/opt/sun/comms/messaging64/**.)
- **ConfigRoot** (The default is **/opt/sun/comms/messaging64/config/**.)

Example

The following example shows the **configtoxml** converting a legacy configuration to a Unified Configuration.

```
# bin/imsimta version
Oracle Communications Messaging Server 7u5-28.12 64bit (built Nov 5 2012)
libimta.so 7u5-28.12 64bit (built 15:58:11, May 23 2012)
Using /opt/sun/comms/messaging/config/imta.cnf (not compiled)
```

```

Linux host1.example.com 2.6.39-100.5.1.el5uek #1 SMP Tue Mar 6 20:25:25 EST 2012
x86_64 x86_64 x86_64 GNU/Linux

# bin/configtoxml
WARNING: This procedure will produce an active Unified Configuration which
        will override any existing legacy configuration.

Continue anyway [no]? yes
Creating the directory /opt/sun/comms/messaging/config/legacy-config/
Moving the processed legacy configuration files to
/opt/sun/comms/messaging/config/legacy-config/
# bin/imsimta version
Oracle Communications Messaging Server 7u5-28.12 64bit (built May Nov 5 2012)
libimta.so 7u5-28.12 64bit (built 15:58:11, Nov 5 2012)
Using /opt/sun/comms/messaging/config/config.xml (not compiled)
Linux host1.example.com 2.6.39-100.5.1.el5uek #1 SMP Tue Mar 6 20:25:25 EST 2012
x86_64 x86_64 x86_64 GNU/Linux

```

Notes on the configtoxml Command

- Stop Messaging Server before running the **configtoxml** command. Alternatively, use the **--noactive** switch to prevent writing out an active configuration.
- When generating an active Unified Configuration, the **configtoxml** command moves all the processed legacy configuration files to the **\$ConfigRoot/legacy-config** directory. The **--undo** option removes the Unified Configuration and restores the legacy configuration files.
- The **--undo** option leaves the Unified Configuration **restricted.cnf** password file in place.

msconfig Command

The **msconfig** command administers the Unified Configuration.

Syntax

You can invoke the **msconfig** command in "interactive" or "non-interactive" mode.

Caution: Only edit your Unified Configuration by running the **msconfig** command. This saves old configurations and allows for rollback. Do not hand-edit any of the Unified Configuration files. Oracle Support may occasionally edit these files to work around any issues pertaining to **msconfig**

- To invoke in non-interactive mode:

```
msconfig <command> <option> <value>
```

See [Table 41-3, "msconfig Commands"](#) for a list of commands.

The option syntax is *scope.optionname*, where *scope* can be multiple . or : delimited name components and *optionname* is a single name component (with no .). In this context, *scope* is the set of groups (and group names) used in the naming convention as described in "[Unified Configuration Option Names.](#)" The option name determines the semantics of the option, and the *scope* determines the part of the product to which

those semantics apply. Thus, an option with "role.base" scope applies to the entire product unless the same option is used with a more specific scope.

For example, consider the **sslnicknames** option. A setting for **base.sslnicknames** applies to the entire product, but a different one can be specified for different parts of the product. For example, **imap.sslnicknames** only applies to the IMAP server (while the rest of the servers would use **base.sslnicknames**).

The **.** delimiter appears before a name component of an option scope that is predetermined. The **:** delimiter appears before a name component of an option scope that is customer supplied and usually extensible.

The *scope* prefix is optional if unambiguous. There is also a higher-level scope of *instance* or *role* that you normally do not specify but that the **msconfig** command does display. For example:

```
role.channel:ims-ms.official_host_name
role.channel:ims-ms.backoff
role.channel:ims-ms.defragment
role.channel:ims-ms.fileinto
role.channel:ims-ms.maxjobs
role.channel:ims-ms.notices
role.channel:ims-ms.pool
```

Some scopes have associated names. For example, **schedule.task:expire.crontab** sets the **crontab** option in the **task** named **expire** in the **schedule** scope.

- To invoke in interactive mode:

```
msconfig
```

When run interactively, the default prompt for **msconfig** changes to the following:

```
msconfig>
```

After you have modified the configuration, the prompt changes to the following:

```
msconfig#
```

The **DEFAULT**, **INSTANCE**, and **ROLE** commands also affect the prompt. In default mode:

```
msconfig>
```

In Instance mode:

```
msconfig.instance>
```

In Role mode:

```
msconfig.role>
```

Options

Table 41–3 lists the commands for **msconfig**:

Table 41–3 *msconfig* Commands

Command	Description
DEFAULT	Uses option's default location
DIRECTORY [<i>filter</i>]	Shows available recipes, optional <i>filter</i> matches string

Table 41–3 (Cont.) msconfig Commands

Command	Description
DIFFERENCES [<i>m</i> [<i>n</i>]]	Compares configurations
EDIT <i>object</i>	Edits by using external editor
EXECUTE <i>command</i>	Executes single recipe command
EXIT	Exits msconfig utility with option to write
HELP [<i>topic</i> [<i>subtopic ...</i>]]	Shows this help
HISTORY	Lists previous saved configurations
INSTANCE	Stores options in instance
IMPORT <i>config</i> [<i>pass</i>]	Reads configuration from alternate file(s)
LOG	Acts as a synonym for history
QUIT	Exits msconfig utility immediately
REVERT [<i>n</i>]	Discards any changes and reloads configuration
ROLE	Stores options in role
RUN <i>recipe</i>	Runs specified recipe
SET <i>option</i> [<i>value</i>]	Sets option to the specified value
SHOW <i>option</i> [<i>namefilter</i> [<i>valuefilter</i>]]	Shows value of option, optional <i>valuefilter</i> matches string
UNSET <i>option</i>	Deletes option from the configuration
WRITE [- <i>remark</i> = <i>remark-string</i>]	Writes configuration changes, includes optional remark

Tip: The **msconfig help** command contains extensive documentation not only about running the **msconfig** command itself, but also on other Unified Configuration topics.

Table 41–4 describes the *object* types used with the **EDIT** sub-command. The **EDIT** sub-command places the specified *object* in a file in an appropriate textual form then invokes the editor specified by the **EDITOR** shell variable. After being edited, the file is read and any changes are incorporated into the configuration.

Table 41–4 object Types Used with the EDIT Sub-command

object Type	Description
ALIASES [<i>alias-name</i>]	MTA aliases
CHANNELS [<i>channel-name</i>]	MTA channels
CONVERSIONS	MTA conversion channel control entries
FILTER	Sieve filter
MAPPINGS [<i>mapping-name</i>]	MTA mappings
OPTION <i>option-name</i>	Specifies an option
REWRITES	MTA rewrite rule

Notes on the msconfig Command

- The **msconfig** command checks syntax and prevents potential misconfigurations due to, for example, entering channel options twice or using invalid (obsolete) channel options.
- Do not run the **msconfig** and **configutil** commands together.
- The **msconfig** command validates the correctness of each individual option value when it is set and the correctness of the entire configuration.
- The **msconfig** prevents writing an invalid configuration to the active configuration.
- Configuration validation is not performed by other Messaging Server processes.
- Schema errors do not prevent Messaging Server from running, however, XML syntax errors do prevent operation.
- If the value provided to the **msconfig** option contains special characters, each special character must be prefixed with the escape character "\" if you are using non-interactive mode. If you use **msconfig** in interactive mode to set the value that contains special characters, you do not need to prefix them with the escape character.

Example: To set the value of the **auth.searchfilter** option to:

```
(|(uid=%U)(mail=%o))
```

you can run the following in non-interactive mode:

```
msconfig set auth.searchfilter "\\(|\\(|(uid\\=%U\\)(mail\\=%o\\))\\)"
```

or you can run the following in interactive mode:

```
msconfig
msconfig> set auth.searchfilter "(|(uid=%U)(mail=%o))"
msconfig# write
```

Option Name Changes in Unified Configuration

- For some **configutil** options, simply dropping the **local.** or **service.** prefix results in the Unified Configuration name but not always. Consult the **msconfig** help documentation for a list of option names.

Tip: Use the **configutil -o legacy-name -H** command to see the Unified Configuration option name.

- Two options were "restructured" rather than just renamed in Unified Configuration: **local.store.notifyplugin** and the MMP **ServiceList**.
- Most MTA **option.dat** option names are unchanged in Unified Configuration.
- Most job controller and dispatcher option names are unchanged in Unified Configuration, with a few exceptions that were misleading.

Table 41–5 lists some of the structural name changes for **configutil** options.

Table 41–5 Structural Name Changes for **configutil** Options in Unified Configuration

Legacy Option	Unified Configuration Option
logfile.*.*	*.logfile.*
sasl.default.ldap.*	auth.*

Table 41–5 (Cont.) Structural Name Changes for configutil Options in Unified

Legacy Option	Unified Configuration Option
sasl.default.*	auth.*
metermaid.table.*.*	metermaid.local_table:*.*
metermaid.mtaclient.*	metermaid_client.*
metermaid.config.*	metermaid.*
alarm.*.*	alarm.system:*.*
store.quotaexceededmsg;lang-*	message_language:*.*quotaexceededmsg
gen.newuserforms;lang-*	message_language:*.*welcomemsg
schedule.*	schedule.task:*.*crontab
encryption.rsa.nssslpersonalityssl	base.sslnicknames
probe.*.*	msprobe.probe:*.*
service.http.proxy.adminpass.*	proxy:*.*httpadminpass (host-specific proxy password)
local.service.proxy.adminpass.*	proxy:*.*imapadminpass (host-specific proxy password)
local.store.notifyplugin.*.*	notifytarget:*.*
service.imap.capability.*	imap.capability_*

Using the msconfig Command in Edit Mode

You can use the **msconfig editobject** command to edit the specified object in a file in an appropriate textual form. The **msconfig edit** command invokes the editor specified by the EDITOR shell variable. You can then make the change, save it, exit, and your configuration is updated.

For example, suppose you want to set the **master_debug** option on the **tcp_local** channel. In a legacy configuration, you need to edit the **imta.cnf** file, locate the **tcp_local** channel block, add the **master_debug** option to it, and save the file. This process is much simplified when you use the **msconfig** command. The equivalent operation is the following command:

```
msconfig set channel:tcp_local.master_debug
```

Even if you did not know the preceding command, you could still use the **msconfig** command to perform this operation by invoking it in edit mode:

```
msconfig edit channels
```

You can also edit a single channel block by itself by running the following command:

```
msconfig edit channel tcp_local
```

The **msconfig** command provides the same editing capability for rewrites (**edit rewrites**), mappings (**edit mappings**) and aliases (**edit aliases**).

Channel-specific option files are mapped into the Unified Configuration as sub-elements of a general "options" channel option. These options appear at the bottom of each channel block when you run **edit channels**. The following example illustrates this point:

```
tcp_local identnonnumeric inner loopcheck maysaslserver maytlsserver mx \
  pool SMTP_POOL remotehost saslswitchchannel tcp_auth smtp sourcespamfilter1 \
  switchchannel
tcp_local-daemon
```

```
==  
trace_level=2
```


Part II

Improving Performance

Part II of the *Messaging Server System Administrator's Guide* describes how to improve Oracle Communications Messaging Server performance.

Part II contains the following chapters:

- [Messaging Server Tuning and Best Practices](#)
- [Tuning the mboxlist Database Cache in Unified Configuration](#)
- [Best Practices for Messaging Server and ZFS](#)

Messaging Server Tuning and Best Practices

This section describes tuning and best practices for Oracle Communications Messaging Server:

- Keep historical logging data
- **Configuring Messaging Server Log File Rotation**

Oracle Communications Messaging Server provides logging facilities for the Messaging Server MTA, the Message Store, and services. The logging facilities provide you with time-stamped and labeled information about your system's messaging services. Using log files, you can gather message statistics, perform trend determination, troubleshoot problems, and so forth.

While the system performs automatic rollovers to maintain the current log file, you must determine and manage log file rotation aspects such as how large a single log file may be, how large cumulative log files may be, how many log files to retain, and so forth.

This section focuses on configuring log file rotation for the Messaging Server service logs, such as the IMAP service.

For more information on managing Messaging Server log files, see the discussion on managing logging in the *Messaging Server System Administrator's Guide*.

Log File Options (Unified Configuration)

The following **msconfig** options pertain to log file rollover:

- `service.logfile.maxlogfilesize`
- `.service.logfile.maxlogsize`
- `.service.logfile.rollovertime`
- `.service.logfile.maxlogfiles`
where *service* is admin, pop, imap, imta, or http.
- **maxlogfilesize** sets the largest size for a given logfile. (The limit is 2 Gbytes, or 2147483648 bytes.)
- **maxlogsize** sets the maximum value for the sum of the log file sizes.
- **maxlogfiles** sets the number of svc.seqNum.timestamp files to keep.
- **rollovertime** sets the interval or age of a file before its get rotated.

You use the **msconfig** command to set these options. For example:

```
msconfig set imap.logfile.maxlogfilesize 1073741824
```

which sets the maximum size of an IMAP service log file to 1 Gbyte.

Log File Options (legacy configuration)

The following **configutil** options pertain to log file rollover:

- **logfile.service.maxlogfilesize**
- **logfile.service.maxlogsize**
- **logfile.service.rollovertime**
- **logfile.service.maxlogfiles**
where *service* is admin, pop, imap, imta, or http.
- **maxlogfilesize** sets the largest size for a given logfile. (The limit is 2 Gbytes, or 2147483648 bytes.)
- **maxlogsize** sets the maximum value for the sum of the log file sizes.
- **maxlogfiles** sets the number of svc.seqNum.timestamp files to keep.
- **rollovertime** sets the interval or age of a file before its get rotated.

You use the **configutil** command to set these options. For example:

```
configutil -o logfile.imap.maxlogfilesize -v 1073741824
```

which sets the maximum size of an IMAP service log file to 1 Gbyte.

Log File Examples

If the conditions at your site meet any one of the above controls, then the system rolls over the log file. Therefore, based on your site's traffic usage, you can set three of these **msconfig** options (or **configutil** options in legacy configuration) to be very large and unreachable while you set the fourth parameter to a value that forces the log to rollover in a manner that you want.

For example, most sites want to retain logs covering specific time periods or time spans, say one week. In order to keep a week's worth of data, set the **maxlogfiles** option equal to *N* times the frequency with which you are performing the log file rotations.

For instance, with the 1 Gbyte limit, you could rotate the files each hour (the rollover time is in seconds, so this is 3600 seconds) and just keep 168 copies of those individual files. Your option settings would then look like the following:

- **maxlogfilesize=1073741824**
- **maxlogsize=VERY_BIG_NUMBER**
- **rollovertime=3600**
- **maxlogfiles=168**

Such settings give you 168 hours (7 days) worth of service log files which are at most 1 Gbyte in size. For *VERY_BIG_NUMBER*, start with 180388626432 (168 x 1073741824).

Use your own custom settings if you prefer smaller log files, quicker log file rollover, and so forth. For example, you will get 336 log files if you set the rollover time to 1800 seconds, or each 30 minutes.

Note: Setting the following options to zero (0) (or less than zero) results in the default value being used:

maxlogfiles (default is 10)

maxlogsize (default is 20971520)

- Configure LMTP between front-end MTA and back-end store. For more information, see the discussion on LMTP delivery in the *Messaging Server System Administrator's Guide*.

- Configure **store.dbtmpdir** to use a memory backed file system to reduce I/O

Instructions for Messaging Server on Oracle Solaris:

1. Configure **store.dbtmpdir**.

```
cd /opt/sun/comms/messaging64/sbin
msconfig set store.dbtmpdir /tmp/msg-server/
```

or in legacy configuration:

```
configutil -o store.dbtmpdir -v "/tmp/msg-server/"
```

2. Start Messaging Server.

```
cd /opt/sun/comms/messaging64/bin
stop-msg
start-msg
```

Instructions for Messaging Server on Red Hat Linux:

1. Configure **store.dbtmpdir**.

```
cd /opt/sun/comms/messaging64/sbin
msconfig set store.dbtmpdir /dev/shm/msg-server/
```

or in legacy configuration:

```
configutil -o store.dbtmpdir -v "/dev/shm/msg-server/"
```

2. Start Messaging Server.

```
cd /opt/sun/comms/messaging64/sbin
stop-msg
start-msg
```

- Move **mboxlist** to separate ZFS or UFS file system with own LUN's

Instructions for Messaging Server on Oracle Solaris:

1. Stop Messaging Server - ensure **stored** process is stopped.

2. Rename **mboxlist** directory to **mboxlist.backup**.

```
cd /opt/sun/comms/messaging64/data/store/
mv mboxlist mboxlist.backup
```

3. Copy **mboxlist** database to *new location*.

```
cp -Rp mboxlist.backup/* <new location>
```

4. Ensure the directory permissions for *new location* are set to the Messaging Server user.

```
chown mailsrv:mail <new location>
```

5. Create symlink to *new location*.

```
ln -s <new location> mboxlist
```

6. Start Messaging Server.

- Put MTA out “front:”
<http://www.oracle.com/technetwork/systems/networking/preferred-deploy-mta-jsp-135699.html>
- Separate ZFS file system for email:
https://blogs.oracle.com/factotum/entry/messaging_server_and_mailstore_best
- Use **imslog.pl** to review MTA traffic and tune accordingly.
 - **imslog.pl** is available in the *MessagingServer_home/examples/unsupported/* directory
- MTA RBL lookup tuning in the discussion about performance tuning realtime blocklists (RBL) lookups in the *Messaging Server Installation and Configuration Guide*.
- Messaging Server Considerations in the discussion of performance tuning considerations for a Messaging Server architecture in the *Messaging Server Installation and Configuration Guide*.
- Reduce spam load.
- Cache tuning:
 - Don't set excessive store **mboxlist** cache size
 - The default 16 MB cache is sufficient for most medium sized environments
 - Larger environments should review the following documentation to determine the best cache size

For more information, see the discussion on **store.dbtmpdir**, **base.lockdir** (Unified Configuration) or **local.lockdir** (legacy configuration), and **store.dbcachesize** in the section about performance tuning considerations for a Messaging Server architecture in the *Messaging Server Installation and Configuration Guide*.
- Number of Messaging Server processes:
 - Set **imap.numprocesses**(Unified Configuration) or **service.imap.numprocesses** (legacy configuration) to (the number of cores / 4). Increase **imap.maxsessions** (Unified Configuration) or **service.imap.maxsessions** (legacy configuration) as needed.
 - Set **http.numprocesses** (Unified Configuration) or **service.http.numprocesses** (legacy configuration) to 1. Increase **http.numprocesses** (Unified Configuration) or **service.http.numprocesses** (legacy configuration) as needed.
 - Monitor the number of threads (LWPs). If during peak load times the number of LWPs in the **mshttpd** or **imapd** processes is constantly over 200, you might need to increase **numprocesses**.
- In a Convergence deployment, the webmail/mshttpd process should run on the same system as the Convergence instance:
 - Reduces network delays
 - Simplifies trouble-shooting

-
- Linux: Change the **IMTA_TMP** option in the MTA tailor file (**imta_tailor**) to use a **tmpfs** such as **/dev/shm**. The default value is **/tmp**, which is a **tmpfs** on Solaris but on Linux is a disk file system.

Tuning the mboxlist Database Cache in Unified Configuration

This chapter includes the following topics:

- [Setting the Mailbox Database Cache Size](#)
- [To Adjust the Mailbox Database Cache Size](#)
- [To Monitor the Mailbox Database Cache Size](#)

Setting the Mailbox Database Cache Size

Oracle Communications Messaging Server makes frequent calls to the mailbox database. For this reason, it helps if this data is returned as quickly as possible. A portion of the mailbox database is cached to improve Message Store performance. Setting the optimal cache size can make a big difference in overall Message Store performance. You set the size of the cache with the `store.dbcachesize` option.

You should use the `store.dbtmpdir` option to redefine the location of the mailbox temporary files to a `tmpfs`, that is, `/tmp/msgDBtmpdir`. On Linux, the `tmpfs` location is usually `/dev/shm`. The default `store.dbtmpdir` value (`/tmp/.xxx`) is not appropriate, so you should change it manually to use the `tmpfs` location. On reboot, the `tmpfs` will be cleared, therefore be sure that permissions are stored so that the correct directory location can be recreated.

Note: The defaults for the `store.dbtmpdir` option and the `local.lockdir` option are now `/tmp/encodedsubdirectory/store` and `/tmp/encodedsubdirectory/lockm` respectively. The `encodedsubdirectory` is `.mailuserMSinstall-location`, where `mailuser` is the Message Server `mailsrv` user and `MSinstall-location` is the install location of Messaging with slashes ("/") replaced by underscores ("_").

For example, if the `mailsrv` user is `mailuser` and Messaging Server is installed in the `/opt/sun/comms/messaging64` directory, then `/tmp/encodedsubdirectory` is `/tmp.mailuser_opt_sun_comms_messaging64/`.

Because of this, you do not need to set these options on a Solaris platform.

To define the location of the mailbox temporary files:

1. Create the directory, for example:

```
mkdir /tmp/msgDBtmpdir
```

2. Assign the appropriate ownership and permissions so that this directory is owned, readable, and writable by the **mailsrv** user. For example, if the **mailsrv** user name is **mailuser** and the group name is **mail**:

```
chown mailuser:mail /tmp/msgDBtmpdir  
chmod 700 /tmp/msgDBtmpdir
```

3. Set the **store.dbtmpdir** option, for example:

```
msconfig  
msconfig> set store.dbtmpdir /tmp/msgDBtmpdir  
msconfig# write  
msconfig> exit  
#
```

Note: Be sure to set the **store.dbtmpdir** option to a uniquely named subdirectory of a **tmpfs** file system such as **/tmp**. In Oracle Solaris Cluster environments, or any situation where it could be possible for multiple instances of Messaging Server to be running on the same system, it is essential that you set this to a name that is unique to that instance. You need to make sure that no two Messaging Server instances can ever try to use the same **tmpdir** directory on the same system, for example, **/tmp/msg-instance-dbtmp** rather than just **/tmp/mboxlist**.

The files stored in the **store.dbtmpdir** location are temporarily memory mapped files used by all processes connecting to the database. Due to their usage pattern, the pages of these files will most likely be in memory all the time. So setting this to be on a **tmpfs** will not really increase memory usage. What it does is save I/O. When the Oracle Solaris virtual memory system sees a memory mapped file on a **tmpfs**, it knows it does not really need to write the modified pages back to the file. So there is only one copy in memory and it saves I/O.

The mailbox database is stored in data pages. When the various daemons make calls to the database (**stored**, **imapd**, **popd**), the system checks to see if the desired page is stored in the cache. If it is, the data is passed to the daemon. If not, the system reads the desired page and writes it in the cache. If there is no clean page available, the system must write one page from the cache back to disk.

Lowering the number of disk read/writes helps performance, so setting the cache to its optimal size is important:

- If the cache is too small, the desired data will have to be retrieved from disk more frequently than necessary.
- If the cache is too large, dynamic memory (RAM) is wasted, and it takes longer to synchronize the disk to the cache.

Of these two situations, a cache that is too small will degrade performance more than a cache that is too large.

Cache efficiency is measured by **hit rate**. Hit rate is the percentage of times that a database call can be handled by cache. An optimally sized cache will have a 98 to 99 percent hit rate (that is, 98 to 99 percent of the desired database pages will be returned to the daemon without having to grab pages from the disk). The goal is to set the smallest cache so that it holds a number of pages such that the cache will be able to

return at least 98 to 99 percent of the requested data. If the direct cache return is less than 98 percent, then you need to increase the cache size.

To Adjust the Mailbox Database Cache Size

Use the **msconfig** command to set the size of the cache with the **store.dbcachesize** option, for example:

```
msconfig
msconfig> set store.dbcachesize 25165825
msconfig# write
msconfig> exit
#
```

It is important to tune the cache size to smallest size that will accomplish the desired hit rate.

The **store.dbcachesize** option controls the size of a shared memory segment used by all processes connected to the database, including **stored**, **imap**, **popd**, **imsbackup**, **imsrestore**, **ims_master**, **tcp_lmtp_server**, and so on. While the maximum value for **store.dbcachesize** is 2 GB, setting it to the maximum wastes memory. Instead, start with the default value of 64 MB and monitor the cache hit rate over a period of days. Increase the value only if the hit rate is under 98%.

Also consider the transaction checkpoint function (performed by **stored**). Set the **store.checkpoint.debug** option and refresh **stored** to see log messages to provide more exact data about transaction checkpoint function time. For example:

```
msconfig
msconfig> set -restricted store.checkpoint.debug 1
msconfig# write
msconfig> exit
refresh store
Refreshing store server 7585 ... done
```

This process must examine all buffers in the cache and hold a region lock during the checkpoint. Other threads needing the lock must wait.

To Monitor the Mailbox Database Cache Size

1. Use the **imcheck** command to measure the cache hit rate.

```
imcheck -s mpool > imcheck-s.out
```

2. Find the cache information section in the output file, for example:

```
2MB 513KB 604B Total cache size.
1 Number of caches.
1 Maximum number of caches
2MB 520KB Pool individual cache size.
```

Then there will be several blocks of output - a summary and one for each database file - look for these lines in each block:

```
0 Requested pages mapped into the process' address space.
55339 Requested pages found in the cache (99%).
```

In this case, the hit rate is 99 percent. This could be optimal or, more likely, it could be that the cache is too large. To test, lower the cache size until the hit rate moves to below 99 percent. When you hit 98 percent, you have optimized the DB cache

size. Conversely, if see a hit rate of less than 95 percent, then you should increase the cache size with the **store.dbcachesize** option.

As your user base changes, the hit rate can also change. Periodically check and adjust this option as necessary.

Best Practices for Messaging Server and ZFS

Oracle ZFS provides the following features that make it ideal for backing up the Oracle Communications Messaging Server message store:

- Snapshot backup
- Enables the use of less expensive SATA drives
- Built-in volume manager that enables you to grow file systems dynamically

Topics:

- [Before You Begin](#)
- [Configuration Recommendations for ZFS and Messaging Server](#)
- [To Configure ZFS and Messaging Server](#)
- [ZFS Administration Recommendations](#)
- [ZFS Reference](#)

Before You Begin

Before using ZFS to back up the Messaging Server message store, read "[Messaging Server and Tiered Storage Overview](#)," which describes message store operation, its performance characteristics, and how to plan for and allocate store partitions.

Configuration Recommendations for ZFS and Messaging Server

The basic recommendations for configuring ZFS and Messaging Server are:

1. Separating the Messaging Server **mbxlist** database, message file, and index cache files on different file systems
2. Configuring the index cache record file system to use the recordsize of 4 Kbytes
3. Disabling file access time record
4. Keeping ZFS pool space under 80 percent utilization to maintain pool performance

The following information provides more context on these recommendations.

mbxlist Database, Message File and Index Cache Files Overview

The **mbxlist** database is a sleepycat database that contains mailbox meta data. The index cache records (**store.idx** and **store.c***) contain meta information about mailboxes and messages. Messaging Server accesses and modifies this meta information, although the modifications tend to be more random and smaller.

The location of the index cache records is controlled by setting the **partition:partition_name.path** option, where *partition_name* is the name of the partition.

Each message file (*.msg) represents a single email. Each message file is written to disk once, never modified, read many times (for example, when a user accesses the email, when the messages are backed up, and so on) and may be deleted. By default, these files are stored with the index and cache files.

The location of message files is controlled by setting the **partition:partition_name.messagepath** option, where *partition_name* is the name of the partition.

Separating the message files from the index cache records to different partitions (and underlying file systems) enables you to configure the file system with properties appropriate for the access type.

Index Cache Record File System

The index recordsize is 128 bytes. Cache recordsize is usually less than 2 Kbytes. The **mboxlist** database maximum page size is 8 Kbytes. The default ZFS recordsize is 128 Kbytes. Reducing the recordsize to 4 Kbytes for these file systems can improve performance and reduce incremental snapshot backup size.

Access Time Record

The message store does not utilize the file access time. By disabling file access time updates, you reduce unnecessary overhead.

ZFS Pool Space Utilization

ZFS pool performance can degrade when a pool is very full. As the pool approaches 100 percent full, more time is needed to find free space and it is more likely that the free space is available only in small chunks.

Configure the disk usage alarm threshold **alarm.system:diskavail.threshold** option to at least 20, to receive a warning before the disk becomes full. (The default value is 10). To enable message throttling sooner, configure the **store.diskusagethreshold** to 90.

To Configure ZFS and Messaging Server

The following steps implement the previously discussed recommendations.

1. Separate the Messaging Server **mboxlist** database, message file, and index cache files on different file systems. For example:

```
zfs create store/mboxlist
zfs set mountpoint=/var/opt/sun/comms/messaging/store/mboxlist store/mboxlist
zfs create store/primary-idx
msconfig set partition:primary.path /store/primary-idx
zfs create store/primary-msg
msconfig set partition:primary.messagepath /store/primary-msg
```

2. Configure the index cache record file system to use the recordsize of 4 Kbytes. For example:

```
zfs set recordsize=8k store/primary-idx
zfs set recordsize=128k store/primary-msg
zfs set recordsize=8k store/mboxlist
```

Note:

- This setting controls the recordsize of newly created files only.
 - Do not set recordsize smaller than the system page size (8 Kbytes on SPARC and 4 Kbytes on Intel).
 - Set recordsize=128k on the **-msg** file system even though that is the default so that it does not accidentally get overridden by a setting on a parent file system at a later time.
-
-

The default recordsize of 128k is appropriate for the message store message file system.

3. Disable file access time record. For example:

```
zfs set atime=off store/mboxlist
zfs set atime=off store/primary-idx
zfs set atime=off store/primary-msg
```

4. Configure the disk usage alarm threshold **alarm.system:diskavail.threshold** option to at least 20, to receive a warning before the disk becomes full. (The default value is 10). For example:

```
msconfig set alarm.system:diskavail.threshold 20
```

5. To enable message throttling sooner, configure the **store.diskusagethreshold** option to 90. (The default is 99). For example:

```
msconfig set store.diskusagethreshold 90
```

ZFS Administration Recommendations

- Perform snapshot backup regularly. Back up the **mboxlist** database, index, and message file systems atomically by using the **zfs snapshot -r** command. Then use the **zfs send** and **receive** commands, or an enterprise-level backup solution to save the data. For example:

```
zfs snapshot -r store@now
zfs send store/mboxlist@now | ssh host2 zfs rcv store/mboxlist
zfs send store/primary-idx@now | ssh host2 zfs rcv store/primary-idx
zfs send store/primary-msg@now | ssh host2 zfs rcv store/primary-msg
```

- Perform incremental backups. You can use **zfs send -i** to perform incremental backups. Destroy the snapshots when they are not needed. For example:

```
zfs destroy -r store@now
```

Note: ZFS snapshots are for backing up and restoring the entire message store files system. You cannot back up and restore individual mailboxes. However, you can use **imsbackup** to back up the snapshot and **imsrestore** to restore the mailboxes.

ZFS Reference

http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide

http://www.solarisinternals.com/wiki/index.php/ZFS_Evil_Tuning_Guide

Part III

Troubleshooting

Part III of the *Messaging Server System Administrator's Guide* describes how to troubleshoot Oracle Communications Messaging Server components.

Part III contains the following chapters:

- [Troubleshooting the MTA](#)
- [Troubleshooting the Message Store](#)

Troubleshooting the MTA

This information describes common tools, methods, and procedures for troubleshooting the Message Transfer Agent (MTA) in Unified Configuration.

Topics:

- [Troubleshooting Overview](#)
- [Standard MTA Troubleshooting Procedures](#)
- [Common MTA Problems and Solutions](#)
- [General Error Messages](#)

A related topic, monitoring procedures can be found in "[Monitoring Messaging Server](#)."

Note: This information assumes that you are familiar with the MTA, both from a conceptual and administration perspective.

Troubleshooting Overview

One of the first steps in troubleshooting the MTA is to determine where to begin the diagnosis. Depending on the problem, you might look for error messages in log files. In other situations, you might check all the standard MTA processes, review the MTA configuration, or start and stop individual channels. Whatever approach you use, consider the following questions when troubleshooting the MTA:

- Did configuration or environmental problems prevent messages from being accepted (for example, disk space or quota problems)?
- Were MTA services such as the Dispatcher and the Job Controller present at the time the message entered the message queue?
- Did network connectivity or routing problems cause messages to be stuck or misrouted on a remote system?
- Did the problem occur before or after a message entered into the message queue?

This information addresses these questions in the subsequent sections.

Standard MTA Troubleshooting Procedures

This section outlines standard troubleshooting procedures for the MTA. Follow these procedures if a problem does not generate an error message, if an error message does not provide enough diagnostic information, or if you want to perform general wellness checks, testing, and standard maintenance of the MTA.

- [Check the MTA Configuration](#)
- [Check the Message Queue Directories](#)
- [Check the Ownership of Critical Files](#)
- [Check that the Job Controller and Dispatcher Are Running](#)
- [Check the Log Files](#)
- [Running a Channel Program Manually](#)
- [Starting and Stopping Individual Channels](#)
- [An MTA Troubleshooting Example](#)

Check the MTA Configuration

Test your address configuration by using the **imsimta test -rewrite** utility. With this utility, you can test the MTA's address rewriting and channel mapping without actually having to send a message. Refer to "[Message Transfer Agent Command-line Utilities](#)" for more information.

The utility will normally show address rewriting that will be applied as well as the channel to which messages will be queued. However, syntax errors in the MTA configuration will cause the utility to issue an error message. If the output is not what you expect, you may need to correct your configuration.

Check the Message Queue Directories

Check if messages are present in the MTA message queue directory, typically *MessagingServer_home/data/queue/*. Use command-line utilities like **imsimta qm** to check for the presence of expected message files under the MTA message queue directory. For more information, see "[imsimta qm](#)" and "[imsimta qm counters](#)."

If the **imsimta test -rewrite** output looks correct, check that messages are actually being placed in the MTA message queue subdirectories. To do so, enable message logging and check the log files in the directory *MessagingServer_home/log/*. For more information on MTA logging, see "[Managing MTA Message and Connection Logs](#)." You can track a specific message by its message ID to ensure that it is being placed in the MTA message queue subdirectories. If you are unable to find the message, you may have a problem with file disk space or directory permissions.

Check the Ownership of Critical Files

You should have selected a mail server user account (**mailsrv** by default) when you installed Oracle Communications Messaging Server. The following directories, subdirectories, and files should be owned by this account:

```
MessagingServer_home/data/queue/  
MessagingServer_home/data/log  
MessagingServer_home/data/tmp
```

Commands, like the ones in the following UNIX system example, can be used to check the protection and ownership of these directories:

```
ls -l -p -d /opt/sun/comms/messaging64/data/queue  
drwx----- 2 mailsrv mail 512 Sep 18 21:17 /opt/sun/comms/messaging64/data/queue  
  
ls -l -p -d /opt/sun/comms/messaging64/data/log  
drwx----- 2 mailsrv mail 2560 Oct 15 05:25 /opt/sun/comms/messaging64/data/log
```

```
ls -l -p -d /opt/sun/comms/messaging64/data/tmp
drwx----- 2 mailsrv mail 512 Sep 18 21:17 /opt/sun/comms/messaging64/data/tmp
```

Check that the files in *MessagingServer_home/data/queue* are owned by the MTA account by a using command such as the following UNIX system example:

```
ls -l -p -R /opt/sun/comms/messaging64/data/queue
```

Check that the Job Controller and Dispatcher Are Running

The MTA Job Controller handles the execution of the MTA processing jobs, including most outgoing (master) channel jobs.

Some MTA channels, such as the MTA's multi-threaded SMTP channels, include resident server processes that process incoming messages. These servers handle the slave (incoming) direction for the channel. The MTA Dispatcher handles the creation of such MTA servers. Dispatcher configuration options control the availability of the servers, the number of created servers, and how many connections each server can handle.

To check that the Job Controller and Dispatcher are present, and to see if there are MTA servers and processing jobs running, use the command **imsimta process**. Under idle conditions the command should result in **job_controller** and *dispatcher* processes. For example:

```
imsimta process
USER PID S VSZ RSS STIME TIME COMMAND
mailsrv2 308 S 50168 26896 23:05:00 00:01 /opt/sun/comms/messaging64/lib/tcp_smtp_
server
mailsrv2 4187 S 46880 17704 Feb_17 01:51 /opt/sun/comms/messaging64/lib/dispatcher
mailsrv2 5887 S 50160 26976 23:25:00 00:00 /opt/sun/comms/messaging64/lib/tcp_
smtp_server
mailsrv2 19018 S 47008 21640 Jun_20 01:21 /opt/sun/comms/messaging64/lib/job_
controller
```

If the Job Controller is not present, the files in the *MessagingServer_home/data/queue* directory get backed up and messages are not delivered. If you do not have a Dispatcher, then you are unable to receive any SMTP connections.

See "[imsimta process](#)" for more information.

You could also use the **imsimta qm jobs** command to list, channel by channel, all active and pending delivery processing jobs currently being managed by the Job Controller. Additional cumulative information is provided for each channel such as the number of message files successfully delivered and those requeued for subsequent delivery attempts. The command syntax is as follows:

```
jobs [-[no]hosts] [-[no]jobs] [-[no]messages] [channel-name]
```

If neither the Job Controller nor the Dispatcher is present, you should review the **dispatcher.log-*** or **job_controller.log-*** file in the *MessagingServer_home/data/log* directory.

If the log files do not exist or do not indicate an error, start the processes by using the "[start-msg](#)" command.

Note: You should not see multiple instances of the Dispatcher or Job Controller when you run **imsimta process**, unless the system is in the process of forking (**fork()**) child processes before it executes (**exec()**) the program that needs to run. However, the time frame during such duplication is very small.

Check the Log Files

If MTA processing jobs run properly but messages stay in the message queue directory, you can examine the log files to see what is happening. All MTA log files are created in the *MessagingServer_home/log* directory. Log file name formats for various MTA processing jobs are shown in [Table 45-1](#).

Table 45-1 MTA Log Files

File Name	Log File Contents
<i>channel_master.log-uniqueid</i>	Output of master program (usually client) for <i>channel</i> .
<i>channel_slave.log-uniqueid</i>	Output of slave program (usually server) for <i>channel</i> .
<i>dispatcher.log-uniqueid</i>	Dispatcher debugging. This log is created regardless if the Dispatcher DEBUG option is set. However, to get detailed debugging information, you should set the DEBUG option to a non-zero value.
imta	ims-ms channel error messages when there is a problem in delivery.
<i>job_controller.log-uniqueid</i>	Job controller logging. This log is created regardless if the Job Controller DEBUG option is set. However, to get detailed debugging information, you should set the DEBUG option to a non-zero value.
<i>tcp_smtp_server.log-uniqueid</i>	Debugging for the tcp_smtp_server . The information in this log is specific to the server, not to messages.
<i>return.log-uniqueid</i>	Debug output for the periodic MTA message bouncer job; this log file is created if the return_debug MTA option is set.

Note: Each log file is created with a unique ID (**uniqueid**) to avoid overwriting an earlier log created by the same channel. To find a specific log file, you can use the **imsimta view** utility. You can also purge older log files by using the **imsimta purge** command. However, by default this command is run on a regular basis (see "Pre-defined Automatic Tasks"). For more information, see "[imsimta purge](#)."

The *channelmaster.log-uniqueid* and *channelslave.log-uniqueid* log files are created in any of the following situations:

- There are errors in your current configuration.
- The **master_debug** or **slave_debug** options are set on the channel.
- If **mm_debug** is set to a non-zero value (**mm_debug > 0**).

For more information on debugging channel master and slave programs, see the *Messaging Server Reference*.

Running a Channel Program Manually

When diagnosing an MTA delivery problem, it is helpful to manually run an MTA delivery job, particularly after you enable debugging for one or more channels.

The command **imsimta submit** notifies the MTA Job Controller to run the channel. If debugging is enabled for the channel in question, **imsimta submit** creates a log file in the *MessagingServer_home/log* directory as shown in [Table 45-1, "MTA Log Files"](#).

The command **imsimta run** performs outbound delivery for the channel under the currently active process, with output directed to your terminal. This might be more convenient than submitting a job, particularly if you suspect problems with job submission itself.

Note: To manually run channels, the Job Controller must be running.

For information on syntax, options, and examples of **imsimta submit** and **imsimta run** commands, refer to "[Command Descriptions](#)."

Starting and Stopping Individual Channels

In some cases, stopping and starting individual channels may make message queue problems easier to diagnose and debug. Stopping a message queue allows you to examine queued messages to determine the existence of loops or spam attacks.

To Stop Outbound Processing (dequeueing) for a Specific Channel

1. Use the **imsimta qm stop** command to stop a specific channel. Doing so prevents you from having to stop the Job Controller and having to recompile the configuration. In the following example, the **conversion** channel is stopped:

```
imsimta qm stop conversion
```

2. To resume processing, use the **imsimta qm start** command to restart the channel. In the following example, the **conversion** channel is started:

```
imsimta qm start conversion
```

See "[imsimta qm](#)" for more information on **imsimta qm start** and **imsimta qm stop** commands.

Note: The command `imsimta qm start/stopchannel` might fail if run simultaneously for many channels at the same time. The tool might have trouble updating the `hold_list` and could report: **QM-E-NOTSTOPPED, unable to stop the channel; cannot update the hold list.** `imsimta qm start/stopchannel` should only be used sequentially with a few seconds interval between each run.

If you only want the channel to run between certain hours, use the following commands:

```
msconfig set job_controller.job_pool:DEFAULT.urgent_delivery
08:00-20:00
msconfig set job_controller.job_pool:DEFAULT.normal_delivery
08:00-20:00
msconfig set job_controller.job_pool:DEFAULT.nonurgent_delivery
08:00-20:00
```

To Stop Inbound Processing from a Specific Domain or IP Address (Enqueuing to a Channel)

You can run one of the following processes if you want to stop inbound message processing for a specific domain or IP address, while returning temporary SMTP errors to client hosts. By doing so, messages are not held on your system. Refer to the "PART 1. MAPPING TABLES."

- To stop inbound processing for a specific host or domain name, add the following access rule to the `ORIG_SEND_ACCESS` mapping table by running the `msconfig edit mapping` command:

```
ORIG_SEND_ACCESS

*|*@example.com|*|* $X4.2.1|$NHost$ temporarily$ blocked
```

By using this process, the sender's remote MTA holds messages on their systems, continuing to resend them periodically until you restart inbound processing.

- To stop inbound processing for a specific IP address, add the following access rule to the `PORT_ACCESS` mapping table by running the `msconfig edit mapping` command:

```
PORT_ACCESS

TCP|*|25|_IP_address_to_block_|* $N400$ can't$ connect$ now
```

When you want to restart inbound processing from the domain or IP address, be sure to remove these rules from the mapping tables and recompile your configuration. In addition, you may want to create unique error messages for each mapping table. Doing so enables you to determine which mapping table is being used.

An MTA Troubleshooting Example

This section explains how to troubleshoot a particular MTA problem step-by-step. In this example, a mail recipient did not receive an attachment to an email message. Note: In keeping with MIME protocol terminology, the "attachment" is referred to as a "message part" in this section. The aforementioned troubleshooting techniques are used to identify where and why the message part disappeared (See "Standard MTA Troubleshooting Procedures"). By using the following steps, you can determine the

path the message took through the MTA. In addition, you can determine if the message part disappeared before or after the message entered the message queue. To do so, you will need to manually stop and run channels, capturing the relevant files.

Note: The Job Controller must be running when you manually run messages through the channels.

Identify the Channels in the Message Path

By identifying which channels are in the message path, you can apply the **master_debug** and **slave_debug** options to the appropriate channels. These options generate debugging output in the channels' master and slave log files. In turn, the master and slave debugging information will assist in identifying the point where the message part disappeared.

1. Set the **log_message_id** MTA option to 1 by running the **msconfig set log_message_id 1** command. With this option set, you will see message ID: header lines in the **mail.log_current** file.
2. Run **imsimta cnbuild** to recompile the configuration.
3. Run **imsimta restartdispatcher** to restart the SMTP server.
4. Have the end user resend the message with the message part.
5. Determine the channels that the message passes through. While there are different approaches to identifying the channels, the following approach is recommended:
 1. On UNIX platforms, use the **grep** command to search for message ID: header lines in the **mail.log_current** file in directory *MessagingServer_home/log*.
 2. Once you find the message ID: header lines, look for the E (enqueue) and D (dequeue) records to determine the path of the message. Refer to ["Understanding the MTA Log Entry Format"](#) for more information on logging entry codes. See the following E and D records for this example:

```
29-Aug-2001 10:39:46.44 tcp_local conversion E 2 ...
29-Aug-2001 10:39:46.44 conversion tcp_intranet E 2 ...
29-Aug-2001 10:39:46.44 tcp_intranet D 2 ...
```

The channel on the left is the source channel, and the channel on the right is the destination channel. In this example, the E and D records indicate that the message's path went from the **tcp_local** channel to the **conversion** channel and finally to the **tcp_intranet** channel.

Manually Start and Stop Channels to Gather Data

This section describes how to manually start and stop channels. See ["Starting and Stopping Individual Channels."](#) By starting and stopping the channels in the message's path, you are able to save the message and log files at different stages in the MTA process. These files are later used to ["To Identify the Point of Message Breakdown."](#)

To Manually Start and Stop Channels

1. Set the **mm_debug** MTA option to 5 by running the **msconfig set mm_debug 5** command to provide substantial debugging information.
2. Add the **slave_debug** and **master_debug** options to the appropriate channels by running the **msconfig edit channels** command and modifying the appropriate channel definitions.

- a. Use the **slave_debug** option on the inbound channel (or any channel where the message is switched to during the initial dialog) from the remote system that is sending the message with the message part. In this example, the **slave_debug** option is added to the **tcp_local** channel.
 - b. Add the **master_debug** option to the other channels that the message passed through and were identified in "[Identify the Channels in the Message Path](#)." In this example, it would be added to the **conversion** and **tcp_intranet** channels.
 - c. Recompile the configuration by using **imsimta cnbuild** if running a compiled configuration.
 - d. Run the command **imsimta restart dispatcher** to restart the SMTP server.
3. Use the **imsimta qm stop** and **imsimta qm start** commands to manually start and stop specific channels. For more on information by using these channel options, see "[Starting and Stopping Individual Channels](#)."
 4. To start the process of capturing the message files, have the end user resend the message with the message part.
 5. When the message enters a channel, the message will stop in the channel if it has been stopped with the **imsimta qm stop** command. For more information, see Step 3 in this section.

- a. Copy and rename the message file before you manually run the next channel in the message's path. See the following UNIX platform example:

```
cp ZZ01K7LXW76T7O9TD0TB.00 ZZ01K7LXW76T7O9TD0TB.KEEP1
```

The message file typically resides in directory similar to *MessagingServer_home/data/queue/destination_channel/001*. The *destination_channel* is the next channel that the message passes through (such as: **tcp_intranet**). If you want to create subdirectories (like **001**, **002**, and so on) in the *destination_channel* directory, add the **subdirs** option to the channels.

- b. It is recommended that you number the extensions of the message each time you trap and copy the message to identify the order in which the message is processed.
6. Resume message processing in the channel and enqueue to the next destination channel in the message's path. To do so, use the **imsimta qm start** command.
 7. Copy and save the corresponding channel log file (for example: **tcp_intranet_master.log-***) located in the *MessagingServer_home/log* directory. Choose the appropriate log file that has the data for the message you are tracking. Make sure that the file you copy matches the timestamp and the subject header for the message as it comes into the channel. In the example of the **tcp_intranet_master.log-***, you might save the file as **tcp_intranet_master.keep** so the file is not deleted.
 8. Repeat steps 5 - 7 until the message has reached its final destination.

The log files you copied in Step 7 should correlate to the message files that you copied in Step 5. If, for example, you stopped all of the channels in the missing message part scenario, you would save the **conversion_master.log-*** and the **tcp_intranet_master.log-*** files. You would also save the source channel log file **tcp_local_slave.log-***.

In addition, you would save a copy of the corresponding message file from each destination channel: **ZZ01K7LXW76T7O9TD0TB.KEEP1** from the **conversion** channel and **ZZ01K7LXW76T7O9TD0TB.KEEP2** from the **tcp_intranet** channel.

9. Remove debugging options once the message and log files have been copied.
 - a. Remove the **slave_debug** and the **master_debug** options from the appropriate channels by running the **msconfig channels** command.
 - b. Reset the **mm_debug** MTA option and remove the setting for the **log_message_id** MTA option by running the **msconfig set mm_debug 0** and **msconfig set log_message_id 0** commands or by running the **msconfig unset mm_debug** and **msconfig unset log_message_id** commands.
 - c. Recompile the configuration by using **imsimta cnbuild** if running a compiled configuration.
 - d. Run the command **imsimta restart dispatcher** to restart the SMTP server.

To Identify the Point of Message Breakdown

1. By the time you have finished starting and stopping the channel programs, you should have the following files with which you can use to troubleshoot the problem:
 - a. All copies of the message file (for example: **ZZ01K7LXW76T7O9TD0TB.KEEP1**) from each channel program
 - b. A **tcp_local_slave.log-*** file
 - c. A set of **channel_master.log-*** files for each destination channel
 - d. A set of **mail.log_current** records that show the path of the message. All files should have timestamps and message ID values that match the message ID: header lines in the **mail.log_current** records.

Note that the exception is when messages are bounced back to the sender; these bounced messages will have a different message ID value than the original message.

2. Examine the **tcp_local_slave.log-*** file to determine if the message had the message part when it entered the message queue.

Look at the SMTP dialog and data to see what was sent from the client machine.

If the message part did not appear in the **tcp_local_slave.log-*** file, then the problem occurred before the message entered the MTA. As a result, the message was enqueued without the message part. If this the case, the problem could have occurred on the sender's remote SMTP server or in the sender's client machine.

3. Investigate the copies of the message files to see where the message part was altered or missing.

If any message file showed that the message part was altered or missing, examine the previous channel's log file. For example, you should look at the **conversion_master.log-*** file if the message part in the message entering the **tcp_intranet** channel was altered or missing.

4. Look at the final destination of the message.

If the message part looks unaltered in the **tcp_local_slave.log**, the message files (for example: **ZZ01K7LXW76T7O9TD0TB.KEEP1**), and the **channel_master.log-*** files, then the MTA did not alter the message and the message part is disappearing at the next step in the path to its final destination. If the final destination is the **ims-ms** channel (the Message Store), then you might download the message from the server to a client machine to determine if the message part is being dropped during or after this transfer.

If the destination channel is a `tcp_*` channel, then you need to go to the MTA in the message's path. Assuming it is an Messaging Server MTA, you will need to repeat the entire troubleshooting process (See ["Identify the Channels in the Message Path,"](#) ["Manually Start and Stop Channels to Gather Data,"](#) and this section). If the other MTA is not under your administration, then the user who reported the problem should contact that particular site.

Common MTA Problems and Solutions

This section lists common problems and solutions for MTA configuration and operation.

- [TLS Problems](#)
- [Changes to Configuration Files or MTA Databases Do Not Take Effect](#)
- [The MTA Sends Outgoing Mail but Does Not Receive Incoming Mail](#)
- [Dispatcher \(SMTP Server\) Won't Start Up](#)
- [Timeouts on Incoming SMTP Connections](#)
- [Messages Are Not Dequeued](#)
- [MTA Messages Are Not Delivered](#)
- [Messages are Looping](#)
- [Received Message is Encoded](#)
- [Server-Side Rules \(SSR\) Are Not Working](#)
- [Slow Response After Users Press Send Email Button](#)
- [Asterisks in the Local Parts of Addresses or Received Fields](#)
- [Abnormal Job Controller Terminations Seen in job_controller Logs](#)

TLS Problems

If, during SMTP dialog, the **STARTTLS** command returns the following error:

454 4.7.1 TLS library initialization failure

and if you have certificates installed and working for POP and IMAP access, check the following:

- Protections/ownerships of the certificates have to be set so **mailsrv** account can access the files.
- The directory where the certificates are stored need to have protections/ownerships set such that the **mailsrv** account can access the files within that directory.

After changing protections and installing certificates, you must run:

```
stop-msg dispatcher  
start-msg dispatcher
```

Restarting should work, but it is better to shut it down completely, install the certificates, and then start things back up.

Changes to Configuration Files or MTA Databases Do Not Take Effect

If changes to your configuration are not taking effect, check to see if you have performed the following steps:

1. Recompile the configuration (by running **imsimta cnbuild**).
2. Restart the appropriate processes (like **imsimta restart dispatcher**).
3. Re-establish any client connections.

The MTA Sends Outgoing Mail but Does Not Receive Incoming Mail

Most MTA channels depend upon a slave or channel program to receive incoming messages. For some transport protocols that are supported by the MTA (like TCP/IP and UUCP), you need to make sure that the transport protocol activates the MTA slave program rather than its standard server. Replacing the native **sendmail** SMTP server with the MTA SMTP server is performed as a part of the Messaging Server installation.

For the multi-threaded SMTP server, the startup of the SMTP server is controlled by the Dispatcher. If the Dispatcher is configured to use a **MIN_PROCS** value greater than or equal to one for the SMTP service, then there should always be at least one SMTP server process running (and potentially more, according to the **MAX_PROCS** value for the SMTP service). The **imsimta process** command may be used to check for the presence of SMTP server processes. See "[imsimta process](#)" for more information.

Dispatcher (SMTP Server) Won't Start Up

If the dispatcher won't start up, first check the **dispatcher.log-*** for relevant error messages. If the log indicates problems creating or accessing the **/tmp/.path.dispatcher.socket** file, then verify that the **/tmp** protections are set to 1777. This would show up in the permissions as follows:

```
drwxrwxrwt 8 root sys 734 Sep 17 12:14 tmp/
.
```

Also do an **ls -l** of the **.path.version-specific-name.dispatcher.socket** file and confirm the proper ownership. For example, if this is created by **root**, then it is inaccessible by **mailsrv**.

Do not remove the **.path.dispatcher.file** and do not create it if it's missing. The dispatcher will create the file. If protections are not set to 1777, the dispatcher will not start or restart because it won't be able to create/access the socket file. In addition, there may be other problems occurring not related to the Messaging Server.

Messaging Server: *MessagingServer_home/config/.var_opt_sun_comms_messaging64-config.dispatcher.socket*

Timeouts on Incoming SMTP Connections

Timeouts on incoming SMTP connections are most often related to system resources and their allocation. The following techniques can be used to identify the causes of timeouts on incoming SMTP connections.

To Identify the Causes of Timeouts on Incoming SMTP Connections

1. Check how many simultaneous incoming SMTP connections you allow. This is controlled by the **MAX_PROCS** and **MAX_CONNS** Dispatcher settings for the SMTP service. The number of simultaneous connections allowed is **MAX_**

PROCS*MAX_CONNS. If you can afford the system resources, consider raising this number if it is too low for your usage.

2. Another technique you can use is to open a TELNET session. In the following example, the user connects to **127.0.0.1** port 25. Once connected, 220 banner is returned. For example:

```
telnet 127.0.0.1 25
Trying 127.0.0.1...
Connected to 127.0.0.1.
Escape character is '^]'.
220 budgie.example.com --Server ESMTP (Sun Java System Messaging Server 6.1
(built May 7 2001))
```

If you are connected and receive a 220 banner, but additional commands (like **ehlo** and **mail from**) do not illicit a response, then you should run **imsimta test -rewrite** to ensure that the configuration is correct.

3. If the response time of the 220 banner is slow, and if running the **pstack** command on the SMTP server shows the following **iii_res*** functions (these functions indicate that a name resolution lookup is being performed):

```
febe2c04 iii_res_send (fb7f4564, 28, fb7f4de0, 400, fb7f458c, fb7f4564) +
42c febdfdcc iii_res_query (0, fb7f4564, c, fb7f4de0, 400, 7f) + 254
```

then it is likely that the host has to do reverse name resolution lookups, even on a common pair like **localhost/127.0.0.1**. To prevent such a performance slowdown, you should reorder your host's lookups in the **/etc/nsswitch.conf** file. To do so, change the following line in the **/etc/nsswitch.conf** file from:

```
hosts: dns nis [NOTFOUND=return] files
```

to:

```
hosts: files dns nis [NOTFOUND=return]
```

Making this change in the **/etc/nsswitch.conf** file can improve performance as fewer SMTP servers have to handle messages instead of multiple SMTP servers having to perform unnecessary lookups.

4. You can also put the **slave_debug** option on the channels handling incoming SMTP over TCP/IP mail, usually **tcp_local** and **tcp_intranet**. After doing so, review the most recent **tcp_local_slave.log-uniqueid** files to identify any particular characteristics of the messages that time out. For example, if incoming messages with large numbers of recipients are timing out, consider using the **expandlimit** option on the channel. Remember that if your system is overloaded and overextended, timeouts will be difficult to avoid entirely.

Messages Are Not Dequeued

Errors encountered during TCP/IP delivery are often transient; the MTA will generally retain messages when problems are encountered and retry them periodically. It is normal on large networks to experience periodic outages on certain hosts while other host connections work fine. To verify the problem, examine the log files for errors relating to delivery attempts. You may see error messages such as, "Fatal error from smtp_open." Such errors are not uncommon and are usually associated with a transient network problem. To debug TCP/IP network problems, use utilities like PING, TRACEROUTE, and NSLOOKUP.

The following example shows the steps you might use to see why a message is sitting in the queue awaiting delivery to **xtel.co.uk**. To determine why the message is not being dequeued, you can recreate the steps the MTA uses to deliver SMTP mail on TCP/IP.

```
% nslookup -query=mx example.com (Step 1)

Server: LOCALHOST
Address: 127.0.0.1

Non-authoritative answer:
example.com preference = 10, mail exchanger = mailhost.example.com (Step 2)

% telnet mailhost.example.com 25 (Step 3)
Trying... [10.1.1.1]
telnet: Unable to connect to remote host: Connection refused
```

1. Use the NSLOOKUP utility to see what MX records, if any, exist for this host. If no MX records exist, then you should try connecting directly to the host. If MX records do exist, then you must connect to the designated MX relays. The MTA honors MX information preferentially, unless explicitly configured not to do so. For more information, see the discussion on TCP/IP nameserver and MX record support in the *Messaging Server Reference*.
2. In this example, the DNS (Domain Name Service) returned the name of the designated MX relay for **xtel.co.uk**. This is the host to which the MTA will actually connect. If more than one MX relay is listed, the MTA will try each MX record in succession, with the lowest preference value tried first.
3. If you do have connectivity to the remote host, you should check if it is accepting inbound SMTP connections by using TELNET to the SMTP server port 25.

Note: If you use TELNET without specifying the port, you will discover that the remote host accepts normal TELNET connections. This does not indicate that it accepts SMTP connections; many systems accept regular TELNET connections but refuse SMTP connections and vice versa. Consequently, you should always do your testing against the SMTP port.

In the previous example, the remote host is refusing connections to the SMTP port. This is why the MTA fails to deliver the message. The connection may be refused due to a misconfiguration of the remote host or some sort of resource exhaustion on the remote host. In this case, nothing can be done to locally to resolve the problem. Typically, you should let the MTA continue to retry the message.

If you are running Messaging Server on a TCP/IP network that does not use DNS, you can skip the first two steps. Instead, you can use TELNET to directly access the host in question. Be careful to use the same host name that the MTA would use. Look at the relevant log file from the MTA's last attempt to determine the host name. If you are using host files, you should make sure that the host name information is correct. It is strongly recommended that you use DNS instead of host names.

If you test connectivity to a TCP/IP host and encounter no problems using interactive tests, it is quite likely that the problem has simply been resolved since the MTA last tried to deliver the message. You can re-run the **imsimta submit tcp_channel** on the appropriate channel to see if messages are being dequeued.

Creating a New Channel

In certain circumstances, a remote domain can break down and the volume of mail addressed to this server can be so great that the outgoing channel queue fills up with messages that cannot be delivered. The MTA tries to redeliver these messages periodically (the frequency and number of the retries is configurable using the **backoff** channel option) and under normal circumstances, no action is needed. However, if too many messages get stuck in the queue, other messages may not get delivered in a timely manner because all the channel jobs are working to process the backlog of messages that cannot be delivered.

In this situation, you can reroute these messages to a new channel running in its own job controller pool. This will avoid contention for processing and allow the other channels to deliver their messages. This procedure is described in the following procedure. Assume a domain called **example.org**.

To Create a New Channel

1. Create a new channel called **tcp_example-daemon** and add a new value for the **pool** option.

Channels are created in by running the **msconfig edit channels** command. The channel should have the same channel options on your regular outgoing **tcp_*** channel. Typically, this is the **tcp_local** channel, which handles all outbound (internet) traffic. Since **example.org** is out on the Internet, this is the channel to emulate. The new channel might look something like this:

```
tcp_example smtp nomx single_sys remotehost inner allowswitchchannel \  
dentnonenumeric subdirs 20 maxjobs 7 pool SMTP_example maytlserver \  
maysaslserver saslswitchchannel tcp_auth missingrecipientpolicy 0 \  
tcp_example-daemon
```

Note the new option-value pair **pool SMTP_example**. This specifies that messages to this channel will only use computer resources from the **SMTP_example** pool. There is a blank line before and after the new channel.

2. Add two rewrite rules by running the **msconfig edit rewrite** to direct email destined for **example.org** to the new channel.

The new rewrite rules look like this:

```
example.org $U$D@tcp_example-daemon  
.example.org $U$H$D@tcp_example-daemon
```

These rewrite rules direct messages to **example.org** (including addresses like **host1.example.org** or **hostA.host1.example.org**) to the new channel whose official host name is **tcp_example-daemon**. The rewriting part of these rules, **\$U\$D** and **\$U\$H\$D**, retain the original addresses of the messages. **\$U** copies the user name from original address. **%** is the separator---the **@** between the username and domain. **\$H** copies the unmatched portion of host/domain specification at the left of dot in pattern. **\$D** copies the portion of domain specification that matched.

3. Define a new job controller pool called **SMTP_example**.

Run the **msconfig set job_controller.job_pool:SMTP_example.job_limit 10** command to create a new pool called **SMTP_example** with a **job_limit** of 10. You can verify the addition of the new pool by running the **msconfig show job_controller.job_pool** command which will show output similar to the following:

```
msconfig show job_controller.job_pool  
role.job_controller.job_pool:DEFAULT.job_limit = 10  
role.job_controller.job_pool:DEFAULT.urgent_delivery = help
```

```
role.job_controller.job_pool:IMS_POOL.job_limit = 2
role.job_controller.job_pool:SMTP_POOL.job_limit = 10
role.job_controller.job_pool:SMTP_example.job_limit = 10
```

This creates a message resource pool called **SMTP_example** that allows up to 10 jobs to be simultaneously run. See "[The Job Controller](#)" for details on jobs and pools.

4. Restart the MTA.

Issue the commands: **imsimta cnbuild; imsimta restart**

This recompiles the configuration and restarts the job controller and dispatcher.

In this example, a large quantity of email from your internal users is destined for a particular remote site called **example.org**. For some reason, **example.org**, is temporarily unable to accept incoming SMTP connections and thus cannot deliver email. (This type of situation is not a rare occurrence.)

As email destined for **example.org** comes in, the outgoing channel queue, typically **tcp_local**, will fill up with messages that cannot be delivered. The MTA tries to redeliver these messages periodically (the frequency and number of the retries is configurable using the **backoff** options) and under normal circumstances, no action is needed.

However, if too many messages get stuck in the queue, other messages may not get delivered in a timely manner because all the channel jobs are working to process the backlog of **example.org** messages. In this situation, you may wish reroute **example.org** messages to a new channel running in its own job controller pool (see "[The Job Controller](#)"). This will allow the other channels to deliver their messages without having to contend for processing resources used by **example.org** messages. Creating a new channel to address this situation is described in the following information.

MTA Messages Are Not Delivered

In addition to message transport problems, there are two common problems which can result in unprocessed messages in the message queues:

1. The queue cache is not synchronized with the messages in the queue directories. Message files in the MTA queue subdirectories that are awaiting delivery are entered into an in-memory queue cache. When channel programs run, they consult this queue cache to determine which messages to deliver in their queues. There are circumstances where there are message files in the queue, but there is no corresponding queue cache entry.

- a. To check if a particular file is in the queue cache, you can use the **imsimta cache -view** utility. If the file is not in the queue cache, then the queue cache needs to be synchronized.

The queue cache is normally synchronized every four hours. If required, you can manually resynchronize the cache by using the command **imsimta cache -sync**. Once synchronized, the channel programs will process the originally unprocessed messages after new messages are processed. If you want to change the default (4 hours), you should modify the job_controller configuration by running the **msconfig set job_controller.synch_timeperiod** command where *timeperiod* reflects how often the queue cache is synchronized. The *timeperiod* must be greater than 30 minutes. In the following example, the queue cache synchronization is modified to 2 hours by running the following command:

```
msconfig set job_controller.synch_time 02:00
```

You can run **imsimta submitchannel** to clear out the backlog of messages after running **imsimta cache -sync**. Clearing out the channel may take a long time if the backlog of messages is large (greater than 1000).

For summarized queue cache information, run **imsimta qm -maint dir -database -total**.

- b. If after synchronizing the queue cache, messages are still not being delivered, you should restart the Job Controller. To do so, use the **imsimta restart job_controller** command.

Restarting the Job Controller causes the message data structure to be rebuilt from the message queues on disk.

Caution: Restarting the Job Controller is a drastic step and should only be performed after all other avenues have been thoroughly exhausted.

Refer to "[The Job Controller](#)" for more information on the Job Controller.

2. Channel processing programs fail to run because they cannot create their processing log file. Check the access permissions, disk space and quotas.

Messages are Looping

If the MTA detects that a message is looping, that message will be sidelined as a .HELD file. See "[Diagnosing and Cleaning up .HELD Messages](#)." Certain cases can lead to message loops which the MTA can not detect.

The first step is to determine why the messages are looping. You should look at a copy of the problem message file while it is in the MTA queue area, MTA mail log entries (if you have the **logging** channel option enabled in your MTA configuration for the channels in question) relating to the problem message, and MTA channel debug log files for the channels in question. Determining the From: and To: addresses for the problem message, seeing the Received: header lines, and seeing the message structure (type of encapsulation of the message contents), can all help pinpoint which sort of message loop case you are encountering.

Some of the more common cases include:

1. A postmaster address is broken. The MTA requires that the postmaster address be a functioning address that can receive email. If a message to the postmaster is looping, check that your configuration has a proper postmaster address pointing to an account that can receive messages.
2. Stripping of Received: header lines is preventing the MTA from detecting the message loop. Normal detection of message loops is based on Received: header lines. If Received: header lines are being stripped (either explicitly on the MTA system itself, or on another system like a firewall), it can interfere with proper detection of message loops. In these scenarios, check that no undesired stripping of Received: header lines is occurring. Also, check for the underlying reason why the messages are looping. Possible reasons include: a problem in the assignment of system names or a system not configured to recognize a variant of its own name, a DNS problem, a lack of authoritative addressing information on the system in question, or a user address forwarding error.

3. Incorrect handling of notification messages by other messaging systems are generating reencapsulated messages in response to notification messages. Internet standards require that notification messages (reports of messages being delivered, or messages bouncing) have an empty envelope From: address to prevent message loops. However, some messaging systems do not correctly handle such notification messages. When forwarding or bouncing notification messages, these messaging systems may insert a new envelope From: address. This can then lead to message loops. The solution is to fix the messaging system that is incorrectly handling the notification messages.

Diagnosing and Cleaning up .HELD Messages

If the MTA detects a serious problem having to do with delivery of a message, the message is stored in a file with the suffix **.HELD** in *MessagingServer_*
home/data/queue/channel. For example:

```
% ls
ZZ0HXZ00G0EBRBCP.HELD
ZZ0HY200C006LGHU.HELD
ZZ0HYA006LP6603H.HELD
ZZ0HZ7003EQQSE37.HELD
```

.HELD files can occur due to three major reasons:

- Looping messages. The MTA detected that the messages were looping via build-up of one or another sort of **Received:** header lines).
- User or domain status set to **hold**. These are messages that are, by intent of the MTA administrator, intentionally being side-lined, typically while some maintenance procedure is being performed, (for example, while moving user mailboxes).
- Suspicious messages. Messages that met some suspicious threshold and were held for later manual inspection by the MTA administrator. Messages can be **.HELD** due to exceeding a configured maximum number of envelope recipients (see the **holdlimit** channel option in the *Messaging Server Reference*), due to running the "imsimta qclean," "clean" or "hold" commands based on some suspicion of the message(s) in question, or due to use of a **hold** action in a Sieve script.

Messages .HELD Due to Looping

Messages bouncing between servers or channels are said to be looping. Typically, a message loop occurs because each server or channel thinks the other is responsible for delivery of the message. Looping messages usually have a great many ***Received:** header lines. The **Received:** header lines illustrate the exact path of the message loop. Look carefully at the host names and any recipient address information (for example, **for recipient** clauses or **ORCPT recipient** comments) appearing in such header lines. One cause of such message loops is user error.

For example, an end users might set an option to forward messages on two separate mail hosts to one another. On their **example.com** account, the users enable mail forwarding to their **example.edu** account. And, forgetting that they have enabled this setting, they set mail forwarding on their **example.edu** account to their **example.com** account.

A loop can also occur with a faulty MTA configuration. For example, MTA Host X thinks that messages for **mail.example.com** go to Host Y. However, Host Y thinks that Host X should handle messages for **mail.example.com**. As a result, Host Y returns the mail to Host X.

In these cases, the message is ignored by the MTA and no further delivery is attempted. When such a problem occurs, look at the header lines in the message to determine which server or channel is bouncing the message. Fix the entry as needed.

Another common cause of message loops is the MTA receiving a message that was addressed to the MTA host using a network name that the MTA does not recognize (has not been configured to recognize) as one of its own names. The solution is to add the additional name to the list of names that your MTA recognizes as *its own*. The MTA's thresholds for determining that a message is looping are configurable; see the **MAX_*RECEIVED_LINES** MTA options

(<http://download.oracle.com/docs/cd/E19566-01/819-4429/index.html>). Also note that the MTA may optionally be configured. See the **HELD_SNDOPR** MTA option to generate a syslog notice whenever a message is forced into **.HELD** state due to exceeding such a threshold. If syslog messages of **Received count exceeded; message held.** are present, then you know that this is occurring.

You can resend the **.HELD** message by using the **imsimta qm "release"** command or by following these steps:

Note: The **imsimta qm "release"** command is the preferred method.

1. Rename the **.HELD** extension to any 2 digit number other than 00. For example, **.HELD** to **.06**.

Note: Before renaming the **.HELD** file, be sure that the message is not likely to continue looping.

2. Run **imsimta cache -sync**.

Running this command updates the cache.

3. Run **imsimta submitchannel** or **imsimta runchannel**.

You might need to perform these steps multiple times, since the message may again be marked as **.HELD**, because the **Received:** header lines accumulate. If the problem still exists, the ***.HELD** file is recreated under the same channel with as before. If the problem has been addressed, the messages are dequeued and delivered.

If you determine that the messages can simply be deleted with no attempt to deliver them, see **"clean."**

Messages .HELD Due to User or Domain hold Status

Messages that are **.HELD** due to a user or domain status of **hold**, and only messages **.HELD** for such a reason, are normally stored in the hold channel's queue area. That is, **.HELD** message files in the hold channel's queue area can be assumed to be **.HELD** due to user or domain status.

Messages .HELD Due to a Suspicious Characteristic

Messages **.HELD** due to some suspicious characteristic exhibit that characteristic. The characteristic could be anything that the site has chosen to characterize as **suspicious**. MTA Administrators should stay aware of these configuration choices and actions. However, if you are not the only or original administrator of this MTA, then check the MTA configuration for any configured use of the **holdlimit** channel option (see the discussion on expansion of multiple addresses in the *Messaging Server Reference*), any use of the **\$H** flag in address-based ***_ACCESS** mapping tables, or any use of the **hold**

action in any system Sieve file, or any channel level Sieve filters configured and named by use of **sourcefilter** or **destinationfilter** channel options. See the discussion on the filter file location in the *Messaging Server Reference*. Additionally, ask any fellow MTA administrators about any manual command-line message holds (through, for instance, an **imsimta qm clean** command) they might have recently performed. Application of a Sieve filter **hold** action, whether from a system Sieve filter or from users' personal Sieve filters, may optionally be logged. See the discussion on the LOG_FILTER global MTA option in the *Messaging Server Reference* for more information.

Received Message is Encoded

Messages sent by the MTA are received in an encoded format. For example:

```
Date: Wed, 04 Jul 2001 11:59:56 -0700 (PDT)
From: "Desdemona Vilalobos" <Desdemona@example.com>
To: santosh@example.edu
Subject: test message with 8bit data
MIME-Version: 1.0
Content-type: TEXT/PLAIN; CHARSET=ISO-8859-1
Content-transfer-encoding: QUOTED-PRINTABLE
```

```
2=00So are the Bo=F6tes Void and the Coal Sack the same?="
```

These messages appear unencoded when read with the MTA decoder command **imsimta decode**. See the *Messaging Server Reference* for more information.

The SMTP protocol only allows the transmission of ASCII characters (a seven-bit character set) as set forth by RFC 821. In fact, the unnegotiated transmission of eight-bit characters is illegal through SMTP, and it is known to cause a variety of problems with some SMTP servers. For example, SMTP servers can go into compute bound loops. Messages are sent over and over again. Eight-bit characters can crash SMTP servers. Finally, eight-bit character sets can wreak havoc with browsers and mailboxes that cannot handle eight-bit data.

An SMTP client used to only have three options when handling a message containing eight-bit data: return the message to the sender as undeliverable, encode the message, or send it in direct violation of RFC 821. But with the advent of MIME and the SMTP extensions, standard encodings exist that can encode eight-bit data by using the ASCII character set.

In the previous example, the recipient received an encoded message with a MIME content type of TEXT/PLAIN. The remote SMTP server (to which the MTA SMTP client transferred the message) did not support the transfer of eight-bit data. Because the original message contained eight-bit characters, the MTA had to encode the message.

Server-Side Rules (SSR) Are Not Working

A filter consists of one or more conditional actions to apply to a mail message. Since the filters are stored and evaluated on the server, they are often referred to as server-side rules (SSR).

This section includes information on the following SSR topics:

- [Testing Your SSR Rules](#)
- [Common Syntax Problems](#)

See also "[To Debug User-level Filters.](#)"

Testing Your SSR Rules

- To check the MTA's user filters, run the following command:

```
imsimta test -rewrite -debug -filter <user>@<domain>
```

In the output, look for the following information:

```
mmc_open_url called to open ssrc: <user>@ims-ms
URL with quotes stripped: ssrc: <user>@ims-ms
Determined to be a SSRD URL.
Identifier: <user>@ims-ms-daemon
Filter successfully obtained.
```

- In addition, you can add the **slave_debug** option to the **tcp_local** channel to see how a filter is applied. The results are displayed in the **tcp_local_slave.log** file. Be sure to set **mm_debug** to 5 by running the **msconfig set mm_debug 5** command to get sufficient debugging information.

Common Syntax Problems

- If there is a syntax problem with the filter, look for the following message in the **tcp_local_slave.log**-* file:

Error parsing filter expression:...

- If the filter is good, then filter information is at the end of the output.
- If the filter is bad, then the following error is at the end of the output:

Address list error - 4.7.1 Filter syntax error: desdaemona@example.com

Also, if the filter is bad, then the SMTP RCPT TO command returns a temporary error response code:

```
RCPT TO: <user>@<domain>
452 4.7.1 Filter syntax error
```

Slow Response After Users Press Send Email Button

If users are experiencing delays when they send messages, undersized message queue disks could be responsible for reduced disk input/output. When users press the SEND button on their email client, the MTA will not fully accept receipt of the message until the message has been committed to the message queue. See the discussion on disk sizing for MTA message queues in the *Messaging Server Installation and Configuration Guide* for more information.

Asterisks in the Local Parts of Addresses or Received Fields

The MTA now checks for 8-bit characters (instead of just ASCII characters) in the local parts of addresses as well as the received fields it constructs and replaces them with asterisks.

Abnormal Job Controller Terminations Seen in job_controller Logs

The Job Controller is essentially an in-memory database. Unlike other parts of the MTA, it doesn't have queues or transactions with which to contend. It listens for activity coming in on various network connections and updates its database accordingly.

Consequently, if the Job Controller fails, it is most likely a resource allocation failure (resource exhaustion). The only significant resource the Job Controller uses, especially when under stress, is memory. Therefore, allocate the right amount of memory for the machine that contains the Job Controller. See the discussion on planning a messaging server sizing strategy in the *Messaging Server Installation and Configuration Guide* for details on memory utilization.

General Error Messages

When the MTA fails to start, general error messages appear at the command line. In this section, common general error messages will be described and diagnosed.

Note: To diagnose your own MTA configuration, use the **imsimta test -rewrite-debug** utility to examine your MTA's address rewriting and channel mapping process. This utility enables you to check the configuration without actually sending a message. See "[Check the MTA Configuration.](#)"

MTA subcomponents might also issue other error messages that are described in the MTA command-line utilities and configuration information in the *Messaging Server Reference*, "[Configuring POP, IMAP, and HTTP Services](#)," and "[About MTA Services](#)." This section includes the following types of errors:

- [Errors in mm_init](#)
- [Compiled Configuration Version Mismatch](#)
- [Swap Space Errors](#)
- [File Open or Create Errors](#)
- [Illegal Host/Domain Errors](#)
- [Errors in SMTP channels, os_smtp_* errors](#)

Errors in mm_init

An error in **mm_init** generally indicates an MTA configuration problem. If you run the **imsimta test -rewrite** utility, these errors are displayed. Other utilities such as **imsimta cnbuild**, or a channel, a server, or a browser might also return such an error.

Commonly encountered mm_init errors include:

- [bad equivalence for alias...](#)
- [cannot open alias include file...](#)
- [duplicate aliases found...](#)
- [duplicate host in channel table...](#)
- [duplicate mapping name found...](#)
- [mapping name is too long...](#)
- [error initializing ch_facility compiled character set version mismatch](#)
- [error initializing ch_facility no room in...](#)
- [local host alias or proper name too long for system...](#)
- [no equivalence addresses for alias...](#)

- [no official host name for channel...](#)
- [official host name is too long](#)

bad equivalence for alias...

The right-hand side of an alias file entry is improperly formatted.

cannot open alias include file...

A file included into the alias file cannot be opened.

duplicate aliases found...

Two alias file entries have the same left hand side. You need to find and eliminate the duplication. Look for an error message that says **error line #XXX** where **XXX** is a line number. You can fix the duplicated alias on the line.

duplicate host in channel table...

This error message indicates that you have two channel definitions in the MTA configuration that both have the same official host name.

Check your MTA configuration for any channel definitions with duplicate official host names.

duplicate mapping name found...

This message indicates that two mapping tables have the same name, and one of the duplicate mapping tables needs to be removed.

Note: A blank line should precede and follow any line with a mapping table name. However, no blank lines should be interspersed among the entries of a mapping table.

mapping name is too long...

This error means that a mapping table name is too long and needs to be shortened.

error initializing ch_facility compiled character set version mismatch

If you see this message, you need to recompile and reinstall your compiled character set tables through the command **imsimta chbuild**. See "[imsimta chbuild](#)" for more information.

error initializing ch_facility no room in...

This error message generally means that you need to resize your MTA character set internal tables and then rebuild the compiled character set tables with the following commands:

```
imsimta chbuild -noimage -maximum -option
imsimta chbuild
```

Verify that nothing else needs to be recompiled or restarted before making this change. See "[imsimta chbuild](#)" for more information.

local host alias or proper name too long for system...

This error indicates that a local host alias or proper name is too long (the optional right hand side in the second or subsequent names in a channel block).

no equivalence addresses for alias...

An entry in the alias file is missing a right hand side (translation value).

no official host name for channel...

This error indicates that a channel definition block is missing the required second line (the official host name line). See the MTA configuration and command-line utilities information in the *Messaging Server Reference* for more information on channel definition blocks. A blank line is required before and after each channel definition block, but a blank line must not be present between the channel name and official host name lines of the channel definition.

official host name is too long

The official host name for a channel (second line of the channel definition block) is limited to 128 octets in length. If you are trying to use a longer official host name on a channel, shorten it to a place holder name, and then use a rewrite rule to match the longer name to the short official host name. You might see this scenario if you work with the *l* (local) channel host name. For example:

```
<Original l Channel:>
!delivery channel to local /var/mail store
l subdirs 20 viaaliasrequired maxjobs 7 pool LOCAL_POOL
walleroo.pocofronitas.thisnameismuchtoolongandreallymakesnosensebutitisan
example.monkey.gorilla.orangutan.antidiexampleblismentarianism.newt.salaman
der.lizard.gecko.komododragon.com
```

```
<Create Place Holder:>
!delivery channel to local /var/mail store
l subdirs 20 viaaliasrequired maxjobs 7 pool LOCAL_POOL
newt
```

```
<Create Rewrite Rule:>
newt.salամander.lizard.gecko.komododragon.com $U%D@newt
```

When using the *l* (local) channel, you need to use a REVERSE mapping table. See the MTA configuration information in the *Messaging Server Reference* for information on usage and syntax.

Compiled Configuration Version Mismatch

One of the functions of the **imsimta cnbuild** utility is to compile MTA configuration information into an image that can be quickly loaded. The compiled format is quite rigidly defined and often changes substantially between different versions of the MTA. Minor changes might occur as part of patch releases.

When such changes occur, an internal version field is also changed so that incompatible formats can be detected. The MTA components halt with the "Compiled Configuration Version Mismatch" error when an incompatible format is detected. The solution to this problem is to generate a new, compiled configuration with the command **imsimta cnbuild**.

Also, use the **imsimta restart** command to restart any resident MTA server processes, so they can obtain updated configuration information.

Swap Space Errors

To ensure proper operation, it is important to configure enough swap space on your messaging system. The amount of required swap space will vary depending on your

configuration. A general tuning recommendation is that the amount of swap space should be at least three times the amount of main memory.

An error message such as the following indicates a lack of swap space:

```
jbc_channels: chan_execute [1]: fork failed: Not enough space
```

You might see this error in the Job Controller log file. Other swap space errors will vary depending on your configuration.

Use the following commands to determine how much swap space you have left and determine how much you have used:

swap -s (at the time MTA processes are busy), **ps -elf**, or **tail /var/adm/messages**

File Open or Create Errors

To send a message, the MTA reads configuration files and creates message files in the MTA message queue directories. Configuration files must be readable by the MTA or any program written against the MTA's SDKs. During installation, proper permissions are assigned to these files. The MTA utilities and procedures which create configuration files also assign permissions. If the files are protected by the system manager, other privileged user, or through some site-specific procedure, the MTA may not be able to read configuration information. This results in "File open" errors or unpredictable behavior. The **imsimta test -rewrite** utility reports additional information when it encounters problems reading configuration files. See "[imsimta test](#)."

If the MTA appears to function from privileged accounts but not from unprivileged accounts, then file permissions in the MTA table directory are likely the cause of the problem. Check the permissions on configuration files and their directories. See "[Check the Ownership of Critical Files](#)."

"File create" errors usually indicate a problem while creating a message file in an MTA message queue directory. See "[Check the Message Queue Directories](#)" to diagnose file creation problems.

Illegal Host/Domain Errors

You might see this error when an address is provided to the MTA through a browser. Or, the error may be deferred and returned as part of an error return mail message. In both cases, this error message indicates that the MTA is not able to deliver mail to the specified host. To determine why the mail is not being sent to the specified host, follow these troubleshooting procedures:

- Verify that the address in question is not misspelled, is not transcribed incorrectly, or does not use the name of a host or domain that no longer exists.
- Run the address in question through the **imsimta test -rewrite** utility. If this utility also returns an "illegal host/domain" error on the address, then the MTA has no rewrite rules or other configurations to handle the address. Verify that you have configured MTA correctly, that you answered all configuration questions appropriately, and that you have kept your configuration information up to date.
- If **imsimta test -rewrite** does not encounter an error on the address, then MTA is able to determine how to handle the address, but the network transport will not accept it. You can examine the appropriate log files from the delivery attempt for additional details. Transient network routing or name service errors should not result in returned error messages, though it is possible for badly misconfigured domain name servers to cause these problems.

- If you are on the Internet, check that you have properly configured your TCP/IP channel to support MX record lookups. Many domain addresses are not directly accessible on the Internet and require that your mail system correctly resolve MX entries. If you are on the Internet and your TCP/IP is configured to support MX records, you should have configured the MTA to enable MX support. See the discussion on TCP/IP connection and DNS lookup support in the *Messaging Server Reference* for more information. If your TCP/IP package is not configured to support MX record lookups, then you will not be able to reach MX-only domains.

Errors in SMTP channels, `os_smtp_*` errors

Errors such as the following are not necessarily MTA errors: `os_smtp_*` errors like `os_smtp_open`, `os_smtp_read`, and `os_smtp_write` errors. These errors are generated when the MTA reports a problem encountered at the network layer. For example, an `os_smtp_open` error means that the network connection to the remote side could not be opened. The MTA may be configured to connect to an invalid system because of addressing errors or channel configuration errors. The `os_smtp_*` errors are commonly due to DNS or network connectivity problems, particularly if this was a previously working channel or address. `os_smtp_read` or `os_smtp_write` errors are usually an indication that the connection was aborted by the other side or due to network problems.

Network and DNS problems are often transient in nature. The occasional `os_smtp_*` error is usually nothing to be concerned about. However, if you are consistently seeing these errors, it could be an indication of an underlying network problem.

To obtain more information about a particular `os_smtp_*` error, enable debugging on the channel in question. Investigate the debug channel log file that will show details of the attempted SMTP dialogue. In particular, look at the timing of when a network problem occurred during the SMTP dialogue. The timing could suggest the type of network or remote side issue. In some cases, you might also want to perform network level debugging (for example, TCP/IP packet tracing) to determine what was sent or received.

Troubleshooting the Message Store

This information provides guidelines for troubleshooting your message store as well as recovery procedures for when the message store becomes corrupted or unexpectedly shuts down.

Topics:

- [Repairing Mailboxes and the Mailboxes Database \(reconstruct Command\)](#)
- [Reduced Message Store Performance](#)
- [Convergence Not Loading Mail Page](#)
- [Command Using Wildcard Pattern Does Not Work](#)
- [Unknown/invalid Partition](#)
- [User Mailbox Directory Problems](#)
- [Store Daemon Not Starting](#)
- [User Mail Not Delivered Due to Mailbox Overflow](#)
- [IMAP Events Become Slow](#)

See also:

- [Using Message Store Log Messages](#)
- [Upgrading the Message Store](#)

Repairing Mailboxes and the Mailboxes Database (reconstruct Command)

If one or more mailboxes become corrupt, use the "reconstruct" utility to rebuild the mailboxes or the mailbox database. You can use this utility to recover from almost any form of data corruption in the mail store. See "[Error Messages Signifying that reconstruct is Needed](#)" and "reconstruct" for more details.

Reduced Message Store Performance

Message store problems can occur if the **mbxlist** database cache is too small. Specifically, Message store performance can slow to unacceptable levels and can even dump core. Refer to the discussion on performance tuning considerations in the *Messaging Server Installation and Configuration Guide*.

Red Hat Linux - Messaging Server Patch 120230-08 IMAP, POP and HTTP Servers Not Starting Due to Over Sessions Per Process

After installing this patch, when you try to start Oracle Communications Messaging Server, the IMAP, POP and HTTP servers do not start and may send the following example error logs:

```
http server - log:
[29/May/2006:17:44:37 +051800] usg197 httpd[6751]: General Critical: Not enough
file
descriptors to support 6000 sessions per process; Recommend ulimit -n 12851 or 87
sessions per process.
pop server - log:
[29/May/2006:17:44:37 +051800] usg197 popd[6749]: General Critical: Not enough
file
descriptors to support 600 sessions per process; Recommend ulimit -n 2651 or 58
sessions per process.
imap server - log:
[29/May/2006:17:44:37 +051800] usg197 imapd[6747]: General Critical: Not enough
file descriptors to support 4000 sessions per process; Recommend ulimit -n 12851
or 58 sessions per process.
```

Set the appropriate number of file descriptors for all three server sessions. Additional file descriptors are available by adding a line similar to the following to the `/etc/sysctl.conf` file and using `sysctl -p` to reread that file:

```
fs.file-max = 65536
```

You must also add a line like the following to the `/etc/security/limits.conf` file:

```
* soft nofile 65536
* hard nofile 65536
```

Convergence Not Loading Mail Page

If users accessing their mail with web clients, like Convergence, cannot load pages, the problem might be that the data is getting corrupted after compression. This can sometimes happen if the system has deployed an outdated proxy server. To solve this problem, try setting the `msconfighttp.gzipstatic` and `http.gzipdynamic` options to `0` to disable data compression. If this solves the problem, you may want to update the proxy server.

Command Using Wildcard Pattern Does Not Work

Some UNIX shells may require quotes around wildcard options and some will not. For example, the C shell tries to expand arguments containing wild cards (`*`, `?`) as files and will fail if no match is found. These pattern matching arguments may need to be enclosed in quotes to be passed to commands like `mboxutil`.

For example:

```
mboxutil -l -p user/usr44*
```

works in the Bourne shell, but fails with `tsch` and the C shell. These shells would require the following:

```
mboxutil -l -p "user/usr44*"
```

If a command using a wildcard pattern does not work, verify whether you need to use quotes around wildcards for that shell.

Unknown/invalid Partition

A user can get the message "Unknown/invalid partition" in Messenger Express if their mailbox was moved to a new partition that was just created and Messaging Server was not refreshed or restarted. This problem only occurs on new partitions. If you now add additional user mailboxes to this new partition, you will not have to do a refresh/restart of Messaging Server.

User Mailbox Directory Problems

A user mailbox problem exists when the damage to the message store is limited to a small number of users and there is no global damage to the system. The following guidelines suggest a process for identifying, analyzing, and resolving a user mailbox directory problem:

1. Review the log files, the error messages, or any unusual behavior that the user observes.
2. To keep debugging information and history, copy the entire *store_root/mboxlist/* user directory to another location outside the message store.
3. To find the user folder that might be causing the problem, run the command **reconstruct -r -n**. If you are unable to find the folder using **reconstruct**, the folder might not exist in the **folder.db**.

If you are unable to find the folder using the **reconstruct -r -n** command, use the **"hashdir"** command to determine the location.

4. Once you find the folder, examine the files, check permissions, and verify the proper file sizes.
5. Use **reconstruct -r** (without the **-n** option) to rebuild the mailbox.
6. If **reconstruct** does not detect a problem that you observe, you can force the reconstruction of your mail folders by using the **reconstruct -r -f** command.
7. If the folder does not exist in the **mboxlist** directory (*store_root/mboxlist/*), but exists in the **partition** directory (*store_root/partition/*), there might be a global inconsistency. In this case, you should run the **reconstruct -m** command.
8. If the previous steps do not work, you can remove the **store.idx** file and run the **reconstruct** command again.

Caution: You should only remove the **store.idx** file if you are sure there is a problem in the file that the **reconstruct** command is unable to find.

9. If the issue is limited to a problematic message, you should copy the message file to another location outside of the message store and run the command **reconstruct -r** on the **mailbox/** directory.
10. If you determine the folder exists on the disk (*store_root/partition/* directory), but is apparently not in the database (*store_root/mboxlist/* directory), run the command **reconstruct -m** to ensure message store consistency.

For more information on the **reconstruct** command, see ["Repairing Mailboxes and the Mailboxes Database \(reconstruct Command\)."](#)

Store Daemon Not Starting

If **stored** does not start and returns the following error message:

```
MessagingServer_home/bin/start-msg
MessagingServer_home: Starting STORE daemon ...Fatal error: Cannot find group in
name service
```

This indicates that the UNIX group configured in **local.servergid** cannot be found. **Stored** and others need to set their **gid** to that group. Sometimes the group defined by **local.servergid** gets inadvertently deleted. In this case, create the deleted group, add **mailsrv** to the group, change ownership of the *instance_root* and its files to **mailsrv** and the group.

User Mail Not Delivered Due to Mailbox Overflow

The message store has a hard limit of two gigabytes for a **store.idx** file, which is equivalent to about one million messages in a single mailbox (folder). If a mailbox grows to the point that the **store.idx** file will attempt to exceed two gigabytes, the user will stop receiving any new email. In addition, other processes that handle that mailbox, such as **imapd**, **popd**, **mshttpd**, could also experience degraded performance.

If this problem arises, you will see errors in **mail.log_current** such as this:

```
05-Oct-2005 16:09:09.63 ims-ms Q 7 ... System I/O error. Administrator, check
server log for details. System I/O error.
```

In addition, the MTA log file will have an errors such as the following:

```
[05/Oct/2005:16:09:09 +0900] jmail ims_master[20745]: Store Error: Unable to
append cache for user/admin: File too large
```

You can determine this problem conclusively by looking at the file in the user's message store directory, or by looking in the **imta** log file to see a more detailed message.

The immediate action is to reduce the size of the file. Either delete some mail, or move some of it to another mailbox. You could also use **mboxutil -r** to rename the folder out of the way, or **mboxutil -d** to delete the folder (see ["mboxutil"](#)).

Long-term, you will need to inform the user of mailbox size limitations, implement an aging policy (see ["Configuring Message Expiration \(Tasks\)"](#)), a quota policy (see ["Message Store Quota Overview"](#)), set a mailbox limit by setting **store.maxmessages** (see the *Messaging Server Reference*), set up an archiving system, or make an adjustment to keep the mailbox size under control.

IMAP Events Become Slow

Symptom: After working fine for a short period of time, many IMAP events become unreasonably slow, with some events taking over a second.

Diagnosis: You have the Event Notification Service (ENS) plugin, **libibiff**, configured, but ENS is not running or not reachable. See ["Administering Event Notification Service"](#) for ENS details.

Solution: If you want ENS notifications, verify that the ENS is enabled and configured correctly. If you do not want ENS notifications, make sure that **libibiff** is not being loaded. Typical incorrect configuration:

```
notifytarget = /opt/sun/comms/messaging/lib/libibiff
ens.enable = 0
```

Instead, use one of the following configurations:

```
notifytarget =
ens.enable = 0
```

or

```
notifytarget = /opt/sun/comms/messaging/lib/libibiff
ens.enable = 1
```


Part IV

Java Message Queue Notifications

Part IV of the *Messaging Server System Administrator's Guide* describes the JMQ notifications and services.

Part IV contains the following chapters:

- [Java Message Queue Notification](#)
- [Java Message Queue Notification Messages and Properties](#)
- [Configuring a JMQ Notification Service \(Tasks and Examples\)](#)
- [Enabling JMQ Event Notification \(Example\)](#)

Java Message Queue Notification

Topics:

- [Two Notification Messaging Services](#)
- [Notification Plug-ins](#)
- [Benefits of Using JMQ Notification](#)
- [JMQ Notification](#)

For more information JMQ notification, see the following topics:

- [Enabling JMQ Event Notification \(Example\)](#)
- [Configuring a JMQ Notification Service \(Tasks and Examples\)](#)
- [Java Message Queue Notification Messages and Properties](#)

The Oracle Communications Messaging Server notification plug-in enables you to deliver notification messages to a messaging service or event service. The messaging service sends the notifications to consumers (client interfaces), which filter and deliver the messages to specified users.

For example, when new email arrives in a user's mailbox, the notification plug-in delivers a notification message to the messaging service. The message consumer, a component of the messaging service, receives the notification and sends it to the user's email client (such as Convergence or Mozilla Thunderbird). The email client can then display a pop-up on the user's computer screen: "You have received a new message."

Another example: if a user's mailbox exceeds its quota, the notification plug-in produces an over-quota notification message. The message consumer sends a warning to the user and to an administrator who needs to be informed of the event.

Two Notification Messaging Services

You can configure Messaging Server to deliver notifications to two different messaging services:

- [Oracle GlassFish Message Queue](#)
- [Event Notification Service](#)

The Message Queue service implements the Java Messaging Service (JMS) specification, providing a message broker, interfaces to create clients that produce or consume messages, and administrative services and control. Message Queue follows the JMS standard for routing and delivery functions, protocols, and message formats.

The Event Notification Service is a component bundled with Messaging Server and Calendar Server. It is a proprietary service that uses a publish/subscribe architecture for sending and receiving event notifications.

You can configure a notification producer for Message Queue, for the Event Notification Service, or for both services.

Note: This information describes how to configure notifications for Message Queue only.

For information about the Event Notification Service, see:

- [Administering Event Notification Service](#)
- [Messaging Server Specific Event Notification Service Information](#)
- [Event Notification Service C API Reference](#)

Note: ENS and JMQ event notification services are "best effort" network services. A best effort network service does not provide any guarantees that the data is delivered. More specifically, under high load situations, ENS and JMQ silently discard notifications. In addition, setting JMQ transport to "reliable" does not prevent notifications from being silently discarded. In fact, "reliable" mode may actually increase the load on the JMQ broker and thus increase the chances for dropped notifications.

Notification Plug-ins

To enable Messaging Server to produce notifications to Message Queue or the Event Notification Service, you must configure a plug-in for that service:

- The JMQ notification plug-in enables you to deliver notification messages to the Message Queue broker.
- The iBiff plug-in allows you to publish notification events to the Event Notification Service.

For information on how to load the iBiff plug-in and configure the Event Notification Service, see "[Administering Event Notification Service](#)."

Benefits of Using JMQ Notification

The JMQ notification plug-in, with Message Queue, provides the following benefits:

- Message Queue implements the JMS standard.
- With Message Queue, you can produce messages to a *topic* or a *queue*, or to both of these delivery methods. For a brief definition, see "[Publishing to a Topic or a Queue](#)."
- Message Queue offers enhanced load balancing during message distribution, especially when messages are produced to a queue.
- The JMQ notification plug-in allows you to configure up to five notification plug-ins. The different plug-ins can produce messages to a topic, to a queue, to the Event Notification Service, and so on. For details, see "[Using Multiple JMQ Notification Plug-ins](#)."

- You can configure notifications conditionally for a specific set of users instead of automatically sending notifications for all users, greatly reducing the total number of messages sent.
- You can configure notifications for different sets of users to be sent to different Message Queue systems, enhancing the scalability of the notification service.
- Message Queue provides reliable delivery of notifications. For example, if you configure the JMQ notification plug-in to produce messages with the persistent flag enabled, the message remains in the Message Queue broker until a consumer receives it. The message is saved so that, if a server goes down, the message can be retrieved and made available to the appropriate consumer.

The following sections describe these benefits in more detail:

- [Publishing to a Topic or a Queue](#)
- [Using Multiple JMQ Notification Plug-ins](#)
- [Configuring Options for a Notification Plug-in](#)
- [Configuring Conditional Notifications for Specified Users](#)
- [Sending Conditional Notifications to Distributed Message Queue Destinations](#)

Publishing to a Topic or a Queue

A topic and queue use different messaging delivery patterns. Each one can be configured in a Message Queue service.

Topic. When a message producer sends a message to a topic, a publish/subscribe architecture is used. In this broadcast pattern, a producer sends a message to a topic destination. Any number of consumers can be subscribed to this topic destination. Each consumer subscribed to the topic gets its own copy of the message. If no consumers are subscribed to the topic, the message is discarded.

The Event Notification Service also uses a publish/subscribe architecture. It is similar to the topic pattern defined in Message Queue.

Queue. When a message producer sends a message to a queue, a point-to-point architecture is used. In this pattern, a producer sends a message to a queue destination from which only one consumer can receive it. If several consumers are waiting for messages from the queue, only one subscriber will receive the message. If no consumers are waiting, the message is held until either the message times out, or a consumer expresses an interest in the queue.

Producing messages to a queue allows you to spread the message load across multiple consumers.

Using Multiple JMQ Notification Plug-ins

You can configure from one to five notification plug-ins.

The `libjmqnotify` notification plugins are built-in. In Unified Configuration, you use the `msconfig` command to specify options for a plug-in and to point the plug-in to the library of executable code.

If you specify more than one plug-in, each plug-in produces notification messages independently of the others. For example, if two plug-ins are configured with a delete-message option, and a message is deleted from a user's mailbox, both plug-ins will produce a notification message.

By configuring multiple plug-ins, you can take advantage of different message-distribution patterns for different purposes. For example, you could configure three different plug-ins to produce messages

- To a queue (using Message Queue)
- To a topic (using Message Queue)
- To the Event Notification Service

Configuring Options for a Notification Plug-in

For each plug-in you configure, you must define a separate set of configuration options.

The options determine two kinds of information:

- The kinds of notification messages to produce. For example, enabling the **loguser** option causes a notification message to be sent whenever a user logs in or out.
- Configuration information needed by Message Queue. For example, the **jmqhost** option identifies the IP address of the host where the Message Queue broker is running.

For instructions on how to configure a plug-in, see "[Configuring a JMQ Notification Service \(Tasks and Examples\)](#)."

Configuring Conditional Notifications for Specified Users

Suppose only a few users in your deployment use notifications. For example, you could have a million-mailbox deployment where 10,000 users have voice-mail notifications sent to their inboxes. Enabling notifications for all users would add unnecessary message traffic to the system.

You can configure notifications to be conditional, so that notifications are generated only for the specified users who require them. The conditional use of notifications can greatly reduce the total number of notifications sent, thus reducing the overall load on the system.

To set up conditional notifications, you do the following:

- Disable notifications for all users
- Enable notifications for those users with the LDAP attribute **mailEventNotificationDestination** added to their directory entries
- Ensure that this LDAP attribute is cached in queued mail messages

When a notification event occurs, Messaging Server looks up the user entry in the directory. If the LDAP attribute is present, a notification is sent to Message Queue. If not, no notification is generated.

For instructions on how to configure this feature, see "[To Enable Conditional Notifications for Specified Users](#)."

Sending Conditional Notifications to Distributed Message Queue Destinations

You can configure notifications for different sets of users to be sent to different Message Queue destinations in a distributed Message Queue environment. For example, for one set of users, notifications can be routed to one Message Queue host. For a second set, notifications can be routed to another Message Queue host.

This feature enhances the scalability of the notification system by distributing notification messages to multiple Message Queue destinations. It operates by extending conditional notifications, described in "[Configuring Conditional Notifications for Specified Users](#)." Notifications are produced for users with the LDAP attribute **mailEventNotificationDestination** in their directory entries. You can associate this attribute with different names; depending on the name, the notification messages are sent to a particular Message Queue destination.

For instructions on how to configure this feature, see "[To Configure Conditional Notifications to Be Sent to Different Message Queue Destinations](#)."

JMQ Notification

Oracle Communications Messaging Server can use Oracle GlassFish Server Message Queue, a standards-based messaging service, to send event notifications. GlassFish Message Queue is provided as a shared component when you install Messaging Server or other Communications Suite products.

This section includes the following topics:

- [Setting JMQ Notifications in Unified Configuration](#)
- [Where to Go for More Information](#)

Setting JMQ Notifications in Unified Configuration

To set JMQ notifications in legacy configurations, you use the string **libjmqnotify** in the **local.store.notifyplugin** option. For example, the following command configures the instance name **jmql**:

```
configutil -o local.store.notifyplugin -v 'libjmqnotify$jmql$'
```

In Unified Configuration, the **local.store.notifyplugin** option is replaced with the **notifytarget:*.notifytype** option, where ***** is the instance name. Thus, for an instance name **jmql** that uses JMQ as its event notification service, you would set **notifytarget:jmql.notifytype** to **jmql** as in the following example:

```
msconfig
msconfig> set notifytarget:jmql.notifytype jmq
msconfig# write -remark "notifytarget jmql"
msconfig> quit
```

Note: The **libibiff** and **libjmqnotify** notification plugins are built-in.

Where to Go for More Information

The following page is a quick-start guide on getting JMQ notifications working along with a sample program to show the output:

- [Enabling JMQ Event Notification \(Example\)](#)

The following pages describe how to configure a JMQ notification plug-in to produce messages to be consumed by clients in a Message Queue service:

- [Java Message Queue Notification](#)
- [Configuring a JMQ Notification Service \(Tasks and Examples\)](#)

- Java Message Queue Notification Messages and Properties

Java Message Queue Notification Messages and Properties

This information describes JMQ notification messages and their properties.

This page contains the following topics:

- [Notification Messages](#)
- [Rules and Guidelines for Notification Messages](#)
- [Notifications for Particular Message Types](#)
- [Default Values of the notifytarget Options](#)
- [Notification Message Properties](#)

For related topics on JMQ notification, see the following pages:

- [Enabling JMQ Event Notification \(Example\)](#)
- [Java Message Queue Notification](#)
- [Configuring a JMQ Notification Service \(Tasks and Examples\)](#)

Notification Messages

Notification messages can be generated for various kinds of events that occur in the message store. For example, when a user logs in, a **login** message can be produced and delivered to the Message Queue broker.

An **msconfig** option specifies each kind of message to be produced. You determine which events generate messages by configuring various options. The options are referenced by one or more JMQ Notification targets.

All messages are delivered to a topic or a queue, depending on whether the destination type is set to **topic** or **queue**. For information on how to configure the Message Queue destination, see "[To Configure an Instance of the JMQ Notification Plug-in](#)."

Each message is identified by the following message header:

MQ_MESSAGE_TYPE_HEADER_PROPERTY

The JMQ Notification plug-in supports the messages shown in [Table 48-1](#).

For a list of the options that enable these messages, see [Table 48-2, "notifytarget Options and Their Default Values"](#).

Table 48–1 JMQ Notification Messages

Notification Message	Description
deletemsg	Messages marked as "Deleted" are removed from the mailbox. This is the equivalent to IMAP expunge.
expungemsg	Messages are expunged from the mailbox.
login	User logged in from IMAP, HTTP, or POP. (This message is enabled with the notifytarget:*.loguser option.)
logout	User logged out from IMAP, HTTP, or POP. (This message is enabled with the notifytarget:*.loguser option.)
msgflags	Message flags on a message have been changed. The old and new flags are carried with this message.
newmsg	New message was received by the system into the user's mailbox. Can contain message headers and body.
overquota	Operation failed because the user's mailbox exceeded one of the quotas (diskquota, msgquota). The MTA channel holds the message until the quota changes or the user's mailbox count goes below the quota. If the message expires while it is being held by the MTA, it is expunged.
purgemsg	Message expunged (as a result of an expired date) from the mailbox by the server process imexpire . This is a server side expunge, whereas deletemsg is a client side expunge. This is not a purge in the true sense of the word.
readmsg	Message in the mailbox was read. (In the IMAP protocol, the message was marked Seen.)
trashmsg	Message was marked for deletion by IMAP or HTTP. The user may still see the message in the folder, depending on the mail client's configuration. The messages are to be removed from the folder when an expunge is performed.
underquota	Quota went back to normal from overquota state.
updatemsg	Message was appended to the mailbox by an IMAP operation. For example, the user copied an email message to the mailbox. Can contain message headers and body.

Rules and Guidelines for Notification Messages

The following rules and guidelines apply to the supported notification messages:

- The text of most notification messages is a single blank space. (The blank space is used because Message Queue does not permit an empty message body.) The exceptions are as follows:
 - The **newmsg**, **updatemsg**, and **expungemsg** messages can include a message header when configured with the **maxheadersize** option. You must set **maxheadersize** to a value greater than zero.
 - **newmsg** and **updatemsg** messages can include a message body when configured with the **maxbodysize** option. You must set **maxbodysize** to a value greater than zero.

For **newmsg** and **updatemsg**, by default the message body is not delivered (is turned off). This prevents overloading Message Queue. No other messages include a message body.

- Notification messages can be generated for changes to the **INBOX** alone, or to the **INBOX** and all other folders. The following option allows for **INBOX** only (value = 0), or for both the **INBOX** and all other folders (value = 1):

```
notifytarget:<jmqnotify>.noninbox
```

The default setting is to generate messages from the **INBOX** only (value = 0).

There is no mechanism to select folders. All folders are included when the variable is enabled (value = 1).

- The **newmsg** notification is issued only after the message is deposited in the user mailbox (as opposed to “after it was accepted by the server and queued in the message queue”).
- Messages are not generated for POP3 client access.
- All messages can be suppressed by issuing **XNOTNOTIFY**. For example, an IMAP script used for housekeeping only (the users are not meant to be notified) might issue it to suppress all messages.

Notifications for Particular Message Types

Notifications can deliver status information about messages of different types, such as text messages, voice mail, and image data. Users often expect these heterogeneous message types to be stored in the same mail folder. For example, a user may want new text messages and voice mail to arrive in the user's cell phone inbox.

To configure these message types, you use options such as **store.message.type.enable**. For information about configuring and managing message types, see ["Managing Message Types in the Message Store."](#)

Once the message types have been configured, JMQ notification messages can identify the particular message types. You can write your Message Queue client to interpret notification messages by message type and deliver status information about each type to the mail client.

For example, suppose new messages of different types arrive in a user's mailbox. A **newmsg** notification message can carry data to tell the user that, for example, there are seven new voice mail messages and four new text messages in the user's inbox.

The following notification messages can carry information that tracks particular message types:

```
newmsg
updatemsg
readmsg
trashmsg
deletemsg
purgemsg
overquota
underquota
```

The JMQ notification function counts the number of messages currently in the mailbox, by message type. Instead of sending one count, an array specifying the count for each message type is sent with the notification message.

The message-specific count is carried in the **nummsgs** property and delivered with the notification message. For **readmsg** and **trashmsg** notification messages, the number of messages seen (**numseen**) and the number marked as deleted (**numdeleted**) are also counted by message type.

Note: The Event Notification Service does not support message types. Use a JMQ notification plug-in to deliver information about message types.

Setting Different Options on a per-messagetype Basis

You can set separate **notifytarget** configuration options, including sending events to a different JMQ host, port, or type, on a per-messagetype basis. When doing so, do not create a per-message type instance, as it is created automatically based on the base instance. You only need to set the options you want to override from the base instance.

For example, assume that you want to use per-messagetype settings and your **notifytarget** configuration is the following:

```
notifytarget = /opt/sun/comms/messaging64/lib/libjmqnotify$msgjmqnotify
```

To set this configuration, you need to put the value in quotes to prevent the "\$" from being interpreted by the shell. Be sure to double check the configuration after setting it.

The options for that base instance would be of the form:

```
notifytarget.msgjmqnotify.<option-name> = <option-value>
```

To override the notifytarget handling of type 1 messages, use options of the form:

```
notifytarget.msgjmqnotify.1.<option-name> = <option-value>
```

Even though there is no "**msgjmqnotify.1**" instance created explicitly in the **notifytarget** option, it is created automatically from the "**msgjmqnotify**" instance because **messagetype 1** is defined. The **notifytarget.msgjmqnotify.1.option-name** settings are used to override the corresponding options from the **notifytarget.msgjmqnotify.option-name** settings.

Default Values of the notifytarget Options

The notification messages and the configuration information needed by Message Queue are configured with the options shown in [Table 48–2](#). For complete definitions of the options, see the *Messaging Server Reference*.

Table 48–2 *notifytarget* Options and Their Default Values

Option	Default Value
notifytarget:*.annotatmsg	0 — Disabled
notifytarget:*.changeflag	0 — Disabled
notifytarget:*.copymsg	0 — Disabled
notifytarget:*.deletemsg	1 — Enabled
notifytarget:*.destinationtype	"topic"
notifytarget:*.enabled	1 — Turned on by default
notifytarget:*.expungemsg	0 — Disabled
notifytarget:*.jmqhost	"127.0.0.1"
notifytarget:*.jmqport	7676
notifytarget:*.jmqpwd	"guest"

Table 48–2 (Cont.) notifytarget Options and Their Default Values

Option	Default Value
notifytarget:*.jmqueue	"JES-MS"
notifytarget:*.jmqtopic	"JES-MS"
notifytarget:*.jmquser	"guest"
notifytarget:*.ldapdestination	null — Turned off by default
notifytarget:*.loguser	1 — Enabled
notifytarget:*.maxbodysize	0 — Disabled
notifytarget:*.maxheadersize	0 — Disabled
notifytarget:*.msgflags	0 — Disabled
notifytarget:*.msgtypes	0 — Disabled
notifytarget:*.newmsg	1 — Enabled
notifytarget:*.noninBox	0 — Disabled
notifytarget:*.overquota	1 — Enabled
notifytarget:*.persistent	1 — Enabled
notifytarget:*.priority	4
notifytarget:*.purgemsg	1 — Enabled
notifytarget:*.readmsg	1 — Enabled
notifytarget:*.ttl	0 — Indicates that messages never time out
notifytarget:*.underquota	1 — Enabled
notifytarget:*.updatemsg	1 — Enabled

Notification Message Properties

Every message carries additional information defined in properties. Different properties are present for different messages. For example, **newmsg** indicates the IMAP **uid** of the new message.

Standard Notification Message Properties

Table 48–3 describes the standard notification message properties. These properties are present in all JMS messages.

Table 48–3 Standard Notification Message Properties

Property	Data Type	Description
hostname	ConstMQString	The host name of the machine that generated the message.
pid	MQInt32	ID of the process that generated the message.
process	ConstMQString	Specifies the name of the process that generated the message.
timestamp	MQFloat64	Specifies the number of milliseconds since the epoch (midnight GMT, January 1, 1970).

Properties Specific to Particular Notification Messages

Table 48–4 describes the properties carried with particular notification messages.

Each message includes a subset of the properties shown in the table below. For a list of the properties associated with each message, see Table 48–5, "Properties Carried with Each Notification Message".

Table 48–4 Properties Specific to Particular Notification Messages

Property	Data Type	Description
attrn	ConstMQString	Annotation attribute name.
client	ConstMQString	The IP address of the Message Queue client associated with the message.
diskquota	MQInt32	The disk space quota, in kilobytes, for the user associated with the message. The value is set to -1 to indicate no quotas.
diskquotaused	MQInt32	The amount of disk space used by the user associated with the message, in kilobytes.
entryn	ConstMQString	Annotation entry name.
hdrLen	MQInt32	The size of the message header. Note that this might not be the size of the header in the message body, because it might have been truncated.
imapUid	MQInt32	The IMAP uid property associated with the message.
lastUid	MQInt32	The last IMAP uid value used in the mailbox.
mailboxName	ConstMQstring	The message-store mailbox name associated with the event. The mailboxName has one of the following formats (where uid is the user's unique identifier): uid - identifies the inbox of a user in the default (primary) domain. uid@domain - identifies the inbox of a user in a hosted domain. uid/mailboxname - identifies the top-level mailbox of a user in the default domain. uid@domain/mailboxname - identifies the top-level mailbox of a user in a hosted domain. uid/foldername/mailboxname - identifies a mailbox in a folder of a user in the default domain. uid@domain/foldername/mailboxname - identifies a mailbox in a folder of a user in a hosted domain.
msgflags	ConstMQString	List of current message flags.
msgquota	MQInt32	The user's quota for the maximum number of messages. The value is set to -1 to indicate no quotas.
newflags	ConstMQString	The flags set for the user's mailbox message after they were changed by the current operation. This property is always present, together with oldflags , when a MsgFlags notification message is produced. For the syntax and values for newflags , see "Syntax for newflags and oldflags Properties," below this table.
numDeleted	MQInt32	The number of messages in the mailbox marked as deleted. This number counts the messages deleted by the mailbox owner. If other users have access to the mailbox, their actions in the mailbox are not included in this count. (However, the other users' actions can trigger notifications such as DeleteMsg).

Table 48–4 (Cont.) Properties Specific to Particular Notification Messages

Property	Data Type	Description
numDeleted_{mm}	MQInt32	The total number of messages in the mailbox marked as deleted, specified for each message type. If message types are configured, a numDeleted_{mm} property carries a count for each message type <i>mm</i> . The numDeleted property is always sent; it counts the total number of all messages marked as deleted, including all types. For example, if 20 messages are marked as deleted, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification: numDeleted=20numDeleted3=10numDeleted16=7
numMsgs	MQInt32	The total number of messages now in the mailbox.
numMsgs_{mm}	MQInt32	The total number of messages now in the mailbox, specified for each message type. If message types are configured, a numMsgs_{mm} property carries a count for each message type <i>mm</i> . The numMsgs property is always sent; it counts the total number of all messages in the mailbox, including all types. For example, if 20 messages are currently in the mailbox, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification: numMsgs=20numMsgs3=10numMsgs16=7
numSeen	MQInt32	The number of messages in the mailbox marked as seen (read). This number counts the messages read by the mailbox owner. If other users have access to the mailbox, their actions in the mailbox are not included in this count. (However, the other users' actions can trigger notifications such as ReadMsg).
numSeen_{mm}	MQInt32	The total number of messages in the mailbox marked as seen (read), specified for each message type. If message types are configured, a numSeen_{mm} property carries a count for each message type <i>mm</i> . The numSeen property is always sent; it counts the total number of all messages marked as seen, including all types. For example, if 20 messages are marked as seen, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification: numSeen=20numSeen3=10numSeen16=7
numSeenDeleted	MQInt32	The number of messages in the mailbox marked as seen (read) and marked as deleted. This number counts the messages marked as read and deleted by the mailbox owner. If other users have access to the mailbox, their actions in the mailbox are not included in this count. (However, the other users' actions can trigger notifications such as ReadMsg and DeleteMsg).
numSeenDeleted_{mm}	MQInt32	The total number of messages in the mailbox marked as seen (read) and marked as deleted, specified for each message type. If message types are configured, a numSeenDeleted_{mm} property carries a count for each message type <i>mm</i> . The numSeenDeleted property is always sent; it counts the total number of all messages marked as seen and deleted, including all types. For example, if 20 messages are marked as seen and deleted, 10 are of type 3, 7 are of type 16, and the rest are not of any recognized type, the following properties and counts are carried with the notification: numSeenDeleted=20numSeenDeleted3=10numSeenDeleted16=7
oldflags	ConstMQString	The flags set for the user's mailbox message before they were changed by the current operation. This property is always present, together with newflags , when a msgflags notification message is produced. For the syntax and values for oldflags , see " Syntax for newflags and oldflags Properties ," below this table.
quotaRoot	ConstMQString	This can be a user name, folder name, or message type.

Table 48–4 (Cont.) Properties Specific to Particular Notification Messages

Property	Data Type	Description
size	MQInt32	The size of the message. Note that this may not be the size of message body, since the body is typically a truncated version of the message.
uidlist	ConstMQString	List of UIDs.
uidValidity	MQInt32	The IMAP uid validity property.
userid	ConstMQString	The userid associated with the message.

Note: Subscribers should allow for undocumented properties when parsing the message reference. This allows for future compatibility when new properties are added.

Syntax for newflags and oldflags Properties

The **newflags** and **oldflags** properties are 5--character strings. The string must have the following values:

- If the /answered flag is set, the first character is "A". If not, it is blank (" ").
- If the /flagged flag is set, the second character is "F". If not, it is blank (" ").
- If the /deleted flag is set, the third character is "D". If not, it is blank (" ").
- If the /seen flag is set, the fourth character is "S". If not, it is blank (" ").
- If the /draft flag is set, the fifth character is "R". If not, it is blank (" ").

Properties Carried with Each Notification Message

Table 48–5 shows the properties associated with each notification message.

For example, to see which properties apply to a **trashmsg** message, look in the column header for "readmsg, trashmsg." A **trashmsg** message can use **mailboxName**, **numMsgs**, **uidValidity**, **numSeen**, and **numDeleted** (in addition to the standard properties).

Table 48–5 Properties Carried with Each Notification Message

Property	new msg, copy msg, update msg	read msg, trash msg	delete msg, purge msg	msg flags	login, logout	overquota, under quota	change flag	annotate msg	expunge msg	create	delete	rename
attrn	No	No	No	No	No	No	No	Yes	No	No	No	No
client	No	No	No	No	Yes	No	No	No	No	No	No	No
diskquota	No	No	No	No	No	Yes	No	No	No	No	No	No
diskquotaused	No	No	No	No	No	Yes	No	No	No	No	No	No
entryn	No	No	No	No	No	No	No	Yes	No	No	No	No
fromUidValidity	Yes	No	No	No	No	No	No	No	No	Yes	No	Yes
hdrLen	Yes	No	No	Yes	No	No	No	No	No	No	No	No
hostname	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
imapUid	Yes	No	Yes	Yes	No	No	No	Yes	No	No	No	No

Table 48–5 (Cont.) Properties Carried with Each Notification Message

Property	new msg, copy msg, update msg	read msg, trash msg	delete msg, purge msg	msg flags	login, logout	overquota, under quota	change flag	annotate msg	expunge msg	create	delete	rename
<i>lastUid</i>	No	No	Yes	No	No	No	No	No	Yes	No	No	No
<i>mailboxName</i>	Yes	Yes	Yes	Yes	No	No	No	Yes	No	Yes	Yes	Yes
<i>modseq</i>	No	No	No	No	No	No	Yes	Yes	Yes	No	No	No
<i>msgflags</i>	No	No	No	No	No	Yes	Yes	No	No	No	No	No
<i>msgquota</i>	No	No	No	No	No	Yes	No	No	No	No	No	No
<i>newflags</i>	No	No	No	Yes	No	No	No	No	No	No	No	No
<i>newName</i>	No	No	No	No	No	No	No	No	No	No	No	Yes
<i>numDeleted</i>	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
<i>numDeletedn</i>	Yes*	Yes*	Yes*	No	No	No	No	No	No	No	No	No
<i>numMsgs</i>	Yes	Yes	Yes	No	No	Yes	No	No	Yes	No	No	No
<i>numMsgsn</i>	Yes*	Yes*	Yes*	No	No	No	No	No	No	No	No	No
<i>numSeen</i>	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
<i>numSeenn</i>	Yes*	Yes*	Yes*	No	No	No	No	No	No	No	No	No
<i>numSeenDeleted</i>	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
<i>numSeenDeletedn</i>	Yes*	Yes*	Yes*	No	No	No	No	No	No	No	No	No
<i>oldflags</i>	No	No	No	Yes	No	No	No	No	No	No	No	No
<i>operation</i>	No	No	No	No	No	No	Yes	No	No	No	No	No
<i>Owner</i>	No	Yes	No	No	No	No	No	No	No	No	No	No
<i>peruser_flags</i>	No	No	No	No	No	No	Yes	No	No	No	No	No
<i>pid</i>	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
<i>process</i>	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
<i>quotaRoot</i>	No	No	No	No	No	Yes	No	No	No	No	No	No
<i>size</i>	Yes	No	No	No	No	No	No	No	No	No	No	No
<i>system_flags</i>	No	No	No	No	No	No	Yes	No	No	No	No	No
<i>timestamp</i>	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
<i>uidlist</i>	No	No	No	No	No	No	Yes	No	Yes	No	No	No
<i>uidValidity</i>	Yes	Yes	Yes	Yes	No	No	Yes	No	Yes	Yes	No	Yes
<i>unchangedsince</i>	No	No	No	No	No	No	Yes	No	No	No	No	No
<i>userid</i>	No	Yes	No	No	Yes	Yes	No	No	No	No	No	No

Note: The *numDeleted_n*, *numMsgs_n*, *numSeen_n*, and *numSeenDeleted_n* properties are carried with notifications only if message types are defined in the message store.

Configuring a JMQ Notification Service (Tasks and Examples)

This information describes how to configure a JMQ notification plug-in.

Topics:

- [Planning for Your JMQ Notification Service](#)
- [To Configure an Instance of the JMQ Notification Plug-in](#)
- [Specifying Notification Messages that Use More Than One msconfig Option](#)
- [To Configure New-Message and Updated-Message Notifications with Message Headers and Message Bodies](#)
- [To Configure Deleted-Message Notifications with Message Headers](#)
- [Configuring Notifications for Changes in Message Status](#)

For related topics on JMQ notification, see the following pages:

- [Enabling JMQ Event Notification \(Example\)](#)
- [Java Message Queue Notification](#)
- [Java Message Queue Notification Messages and Properties](#)

Planning for Your JMQ Notification Service

A JMQ notification plug-in is only one part of a Message Queue service. A Message Queue service also includes clients that consume the event messages and the Message Queue infrastructure (the broker, administration components, and so on).

The following steps outline the tasks you should perform to create a Message Queue service that supports Messaging Server:

1. Design your notification message service. Define the notification messages needed for your Oracle Communications Messaging Server installation. The planning and design phases of your message-service development lifecycle lie outside the scope of this information. However, you should answer the following design questions before you configure the JMQ notification plug-in:
 - Which message events do you need to produce notifications? For a list of the available notification messages, see "[Notification Messages](#)."
 - Do you intend to produce messages to a queue, a topic, or both?
 - Do you intend to use the proprietary Event Notification Service as well as the Message Queue service? The answers to these questions help you decide

whether to configure one instance of the notification plug-in or multiple instances, and to determine how to configure each instance.

2. Install, configure, and deploy the Message Queue product. For information about installing Message Queue, see the *Sun GlassFish Message Queue 4.4 Update 1 Installation Guide*. For information about configuring and deploying Message Queue, see the *Sun GlassFish Message Queue 4.4 Administration Guide*.
3. Write one or more Message Queue clients that consume the JMQ notification messages. The clients must conform to the requirements for a Message Queue Java application programming interface (API). For information about writing Message Queue clients in Java, see the *Sun GlassFish Message Queue 4.4 Developer's Guide for Java Clients*.
4. Configure and enable the JMQ notification plug-in for producing notification messages. The remainder of this information describes how to configure the notification plug-in.
5. Configure and start the runtime Message Queue clients. For information about deploying the runtime Message Queue clients, see the *Sun GlassFish Message Queue 4.4 Administration Guide*.

To Configure an Instance of the JMQ Notification Plug-in

Note: The `libibiff` and `libjmqnotify` notification plugins are now built-in.

In this procedure, you first configure the message events that produce notifications. Next, you specify the information needed by Message Queue. Finally (in step 10), you configure the instance name by specifying an option after the name of the plug-in library:

```
msconfig set notifytarget:<instance>.notifytype jmq
```

Before You Begin

You should install, configure, and deploy the following products:

- Messaging Server
- Oracle GlassFish Message Queue (see the *Messaging Server Release Notes* for more information)

Note: Most of the configuration options that you configure in the following steps are optional. For a list of their default values, see the *Messaging Server Reference*.

The `msconfig` options used by the default instance have names of the following form:

```
notifytarget:<instance>.<option>
```

1. Configure the notification event message options. For each kind of notification event message you want to include in the instance, use the **notifytarget**:`instance.event` option. For example, to enable notifications for new messages, type:

```
msconfig set notifytarget:<jmqnotify>.newmsg 1
```

where *jmqnotify* is the name of the default instance and **1** enables notifications for this event. A value of **0** disables notifications for this event. For a list of all the JMQ notification messages, see "[Notification Messages](#)." A few notification messages use more than one option to enable the message with additional features. For example, some messages can carry message headers in the notification text. For instructions on how to configure these messages, see "[Syntax for newflags and oldflags Properties](#)."

Note: You must configure options separately for each instance you configure. Thus, if you configure two instances, named **jmq1** and **jmq2**, and you want to enable new-message notifications for both instances, you must set that option for both instances:

```
msconfig set notifytarget:jmq1.newmsg 1
msconfig set notifytarget:jmq2.newmsg 1
```

2. Specify the host where the Message Queue destination (broker) is running. For example:

```
msconfig set notifytarget:<instance>.jmqhost "127.0.0.1"
```

3. Specify the port for the Message Queue broker. For example:

```
msconfig set notifytarget:<instance>.jmqport "7676"
```

4. Specify the user ID of the Message Queue user authorized to produce messages to the service. For example:

```
msconfig set notifytarget:<instance>.jmquser "guest"
```

5. Specify the password of the Message Queue user. You need to do this in interactive mode, for example:

```
msconfig
msconfig> set -prompt notifytarget:<instance>.jmqpwd
Password:
Verify:
msconfig# write
msconfig> exit
```

6. Specify whether the destination is a "**topic**" or a "**queue**". For example:

```
msconfig set notifytarget:<instance>.destinationtype "queue"
```

7. Specify the destination *topic* or *queue* name. For example, type one of the following commands:

```
msconfig set notifytarget:<instance>.jmqqueue "JES-MS"
```

or

```
msconfig set notifytarget:<instance>.jmqtopic "JES-MS"
```

The **jmqqueue** and **jmqtopic** options are synonymous and mutually exclusive. You can only use one of these options in one instance. "**JES-MS**" is an example name of the queue or topic to which messages are sent.

- Specify the Message Queue priority assigned to messages produced by this instance. For example:

```
msconfig set notifytarget:<instance>.priority 3
```

The default value of the **priority** option is 4.

- Specify the length of time (in milliseconds) that messages are retained by the Message Queue broker. For example:

```
msconfig set notifytarget:<instance>.ttl 100
```

This example specifies that a message is retained by the Message Queue service for 100 milliseconds before being either delivered or discarded. A value of 0 means that a message is retained permanently. It does not time out.

- Specify the message persistence. For example:

```
msconfig set notifytarget:<instance>.persistent 1
```

The **1** specifies that persistent messages are used in the Message Queue service. Allowed values are 1 (persistent) and 0 (non-persistent).

- Configure the instance name. To configure a single instance with the default name, type the following:

```
msconfig set notifytarget:<instance>.notifytype jmq
```

Note that **jmqnotify** is the default instance name.

To configure a different instance name such as **jmq42**, type the following command:

```
msconfig set notifytarget:jmq42.notifytype jmq
```

The **msconfig** options read by the **jmq42** instance would have names like:

```
notifytarget:jmq42.<option>
{code:none}
```

h2. To Configure Multiple Instances

Configure a separate set of JMQ notification parameters for each instance you intend to create.

For example, suppose you configure two instances named `{{jmq1}}` and `{{jmq2}}`. Assume you want to enable new-message notifications for both instances and purged-message notifications for the `{{jmq2}}` instance only. In this case, you would set the following options:

```
{code:none}
msconfig set notifytarget:jmq1.newmsg 1
msconfig set notifytarget:jmq2.newmsg 1
msconfig set notifytarget:jmq2.purgemsg 1
```

You also must specify options that enable the instances to communicate with the Message Queue service. For step-by-step instructions for configuring the all the notification options, see "[To Configure an Instance of the JMQ Notification Plug-in.](#)"

- Configure the instance names. To configure two instances named **jmq1** and **jmq2**, type the following command:

```
msconfig set notifytarget:jmq1.notifytype jmq
msconfig set notifytarget:jmq2.notifytype jmq
```

In this example, the first instance builds its configuration from options with the name **jm1**:

```
notifytarget:jm1.<option>
```

The second instance builds its configuration from options with the name **jm2**:

```
notifytarget:jm2.<option>
```

Specifying Notification Messages that Use More Than One `msconfig` Option

For most notification messages, you specify the message by running a single **msconfig notifytarget** command.

However, the following notification messages are (or can be) configured with more than one **msconfig notifytarget** command:

- **newmsg**
- **updatemsg**
- **deletemsg**
- **msgflags**

The following procedures describe how to set up these notification messages.

To Configure New-Message and Updated-Message Notifications with Message Headers and Message Bodies

You can add the message headers and message bodies to the text of notification messages sent when there are new or updated email messages.

Including message headers and message bodies is optional. You can include both features, one feature only, or neither feature. The default is to send messages without message headers or message bodies.

1. Specify the new-message or updated-message notification:

```
msconfig set notifytarget:<instance>.newmsg 1
msconfig set notifytarget:<instance>.updatemsg 1
```

2. Specify the **maxheadersize** option with a value greater than zero:

```
msconfig set notifytarget:<instance>.maxheadersize 1024
```

The default value of **maxheadersize** is **0**, which sends no header information with the message.

3. Specify the **maxbodysize** option with a value greater than zero:

```
msconfig set notifytarget:<instance>.maxbodysize 1024
```

The default value of **maxbodysize** is **0**, which sends no body with the message.

To Configure Deleted-Message Notifications with Message Headers

You can add the message headers to the text of notification messages sent when email messages are deleted.

Including message headers is optional. The default is to send notifications without message headers.

1. Enable notifications to be sent when email messages are deleted:

```
msconfig set notifytarget:<instance>.deletemsg.enable 1
```

2. Specify the **expungemsg** option:

```
msconfig set notifytarget:<instance>.expungemsg 1
```

Setting the value to **1** enables message headers to be carried with deleted-message notifications. The default value of **expungemsg** is **0**, which prohibits deleted-message notifications from carrying header information. You must configure the **expungeheaders** option to enable **deletemsg** messages to carry message headers.

3. Specify the **maxheadersize** option with a value greater than zero, as in the following example:

```
msconfig set notifytarget:<instance>.maxheadersize 1024
```

The default value of **maxheadersize** is **0**, which sends no header information with the message.

Configuring Notifications for Changes in Message Status

You can configure a notification message to be sent when an email message has changed status.

A message-flag notification is produced whenever a status flag has changed because the email message was:

- Answered
- Flagged
- Deleted
- Seen (read)
- Draft

When a message-flag notification is sent, the notification carries the following properties:

- The flags set for the email message before its status changed
- The flags set for the email message after its status changed

This information is carried in two properties, **oldflags** and **newflags**, which are 5--character strings.

For a description of the values of these two properties, see ["Syntax for newflags and oldflags Properties."](#)

To Enable Notifications When Message-Flags Have Changed

To enable message-flag notifications, you must configure the following **msconfig** option:

- **notifytarget:<instance>.msgflags 1**

This option enables the IMAP server and message store to identify and track the changing values of the status flags so that this information can be delivered in notification messages. This option is instance specific. If an instance does not use message-flag notifications, be sure that this option is disabled (default).

To Enable Conditional Notifications for Specified Users

This task enables you to configure notifications to be sent for specific users who require notifications rather than for all users in your deployment. The conditional use of notifications can greatly reduce the total number of notifications sent, thus reducing the overall load on the system.

For information about how conditional notifications work, see "[Configuring Conditional Notifications for Specified Users](#)."

Before You Begin

Follow all the steps in "[To Configure an Instance of the JMQ Notification Plug-in](#)."

Take These Steps

1. Disable notifications for all users:

```
msconfig set notifytarget:<instance>.enable 0
```

2. Specify the name of the LDAP attribute which, if present in the user's LDAP entry, enables notifications for that user and specify the instance to use:

```
msconfig set notifytarget:<instance>.ldapdestination  
mailEventNotificationDestination
```

3. Add the attribute to the LDAP entries of the users who require notifications. For example:

```
mailEventNotificationDestination: jmqnotify
```

where **jmqnotify** is the default instance name.

4. Ensure that this LDAP attribute is cached in enqueued messages and carried in LMTP deliveries.

```
msconfig set LDAP_SPARE_1=mailEventNotificationDestination  
msconfig set SPARE_1_SEPARATOR 259
```

5. Rebuild the configuration file and restart Messaging Server.

```
stop-msg  
imsimta cnbuild  
start-msg
```

To Configure Conditional Notifications to Be Sent to Different Message Queue Destinations

This task enables notifications for different sets of users to be sent to different Message Queue destinations in a distributed Message Queue environment. For example, for one set of users, notifications can be routed to one Message Queue host; for a second set, notifications can be routed to another Message Queue host.

This task begins with the similar steps as the preceding task, "[To Enable Conditional Notifications for Specified Users](#)." It extends that task to allow for multiple Message Queue destinations.

For more information about this feature, see "[Sending Conditional Notifications to Distributed Message Queue Destinations](#)."

Before You Begin

Follow **Step 1, Configure the notification message options**, in "[To Configure an Instance of the JMQ Notification Plug-in](#)."

Take These Steps

1. Disable notifications for all users:

```
msconfig set notifytarget:<instance>.enable 0
```

2. Specify the name of the LDAP attribute which, if present in the user's LDAP entry, enables notifications for that user and specify the instance to use:

```
msconfig set notifytarget:<instance>.ldapdestination  
mailEventNotificationDestination
```

3. Add the attribute to the LDAP entries of the users who require notifications. For example:

```
mailEventNotificationDestination: mqdestination1
```

Here the LDAP attribute is set to the value **mqdestination1**. This value can be any string. However, this value must match the JMQ destination name (instance) in the options that you set to provide configuration information needed by Message Queue. These options are listed in the next step.

4. Use the **msconfig set** command to configure the following JMQ configuration options:

```
notifytarget:<jmq_destination_name>.jmqhost  
notifytarget:<jmq_destination_name>.jmqport  
notifytarget:<jmq_destination_name>.jmquser  
notifytarget:<jmq_destination_name>.jmqpwd  
notifytarget:<jmq_destination_name>.destinationtype
```

```
notifytarget:<jmq_destination_name>.jmqtopic  
or  
notifytarget:<jmq_destination_name>.jmqqueue
```

```
notifytarget:<jmq_destination_name>.priority  
notifytarget:<jmq_destination_name>.ttl  
notifytarget:<jmq_destination_name>.persistent
```

For example:

```
msconfig set notifytarget:mqdestination1.jmqhost "127.0.0.1"
```

For details about how to configure these options, see **Steps 2 through 10** in "[To Configure an Instance of the JMQ Notification Plug-in](#)."

5. For each different JMQ destination that you want to create for different sets of users, repeat **Steps 3 and 4**, in this procedure. For example, to create two additional JMQ destinations named **mqdestination2** and **mqdestination3**:

- a. Use these strings in the users' **mailEventNotificationDestination** LDAP attributes. For example, for all users whose notifications go to the second destination, add:

```
mailEventNotificationDestination: messagequeuedestination2
```

For all users whose notifications go to the third destination, add:

```
mailEventNotificationDestination: messagequeuedestination3
```

- b. Set each JMQ configuration option (such as **jmqhost**) to its specific value. For example:

```
msconfig set notifytarget:mqdestination2.jmqhost "127.0.0.2"
```

```
msconfig set notifytarget:jmqdestination3.jmqhost "127.0.0.3"
```

and so on for each option.

6. Restart Messaging Server.

```
stop-msg  
start-msg
```

Enabling JMQ Event Notification (Example)

This example contains the following sections:

- [Event Notifications in Messaging Server Overview](#)
- [Enable Java Message Queue \(JMQ\)](#)
- [Configure and Enable the Messaging Server jmqnotify Plugin](#)
- [Verify the JMQ Broker](#)
- [Related Information](#)

Details on JMQ integration can be found in the following pages:

- [Java Message Queue Notification](#)
- [Configuring a JMQ Notification Service \(Tasks and Examples\)](#)
- [Java Message Queue Notification Messages and Properties](#)

Event Notifications in Messaging Server Overview

Oracle Communications Messaging Server supports two mechanisms for event notifications. One mechanism is the ENS event notification plugin. The second is the JMQ (Java Message Queue) event notification plugin.

The following is a quick-start guide on getting JMQ notifications working along with a sample program to show the output.

Enable Java Message Queue (JMQ)

1. Modify the JMQ configuration file `/etc/imq/imqbrokerd.conf`.

```
replace:
```

```
AUTOSTART=NO
```

```
with:
```

```
AUTOSTART=YES
```

2. Start Java Message Queue.

```
/etc/init.d/imq start
```

3. Reset the **admin/guest** password and add the **jesuser** account.

```
cd /usr/bin
```

```
imqusermgr update -u admin -p password
Are you sure you want to update user admin? (y/n) y
imqusermgr update -u guest -p guest
Are you sure you want to update user guest? (y/n) y
imqusermgr add -u jesuser -g user -p password
User repository for broker instance: imqbroker
User jesuser successfully added.
```

Configure and Enable the Messaging Server jmqnotify Plugin

1. Enable the appropriate Messaging Server settings. In this example, the target is **jmqnotify**.

```
cd /opt/sun/comms/messaging64/bin
msconfig set notifytarget:jmqnotify.notifytype jmq
msconfig set notifytarget:jmqnotify.newmsg 1
msconfig set notifytarget:jmqnotify.updatemsg 1
msconfig set notifytarget:jmqnotify.deletemsg 1
msconfig set notifytarget:jmqnotify.maxheadersize 1024
msconfig set notifytarget:jmqnotify.jmqhost "127.0.0.1"
msconfig set notifytarget:jmqnotify.jmqport "7676"
msconfig set notifytarget:jmqnotify.jmquser "jesuser"
msconfig set notifytarget:jmqnotify.destinationtype "queue"
msconfig set notifytarget:jmqnotify.jmqqueue "jesms"
msconfig set notifytarget:jmqnotify.priority 3
msconfig set notifytarget:jmqnotify.ttl 1000
msconfig set notifytarget:jmqnotify.persistent 1
```

2. Set the password for **jmquser** option, which in this example, is **jesuser**.

```
msconfig
msconfig> set -prompt notifytarget:jmqnotify.jmqpwd
Password:
Verify:
msconfig# write
msconfig> exit
```

3. Restart Messaging Server.

```
stop-msg
start-msg
```

4. Verify that the configuration is working by checking for an entry such as the following in the *MessagingServer_home/log/imap log*:

```
[29/Oct/2012:15:45:38 -0700] ipg-test3 imapd[18062]: General Notice: JMQ
notifications enabled: jmqnotify
```

Verify the JMQ Broker

You can verify that the JMQ software is operational at any time by running the following command:

```
cd /usr/bin
imqcmd query bkr -u admin
Password: password
Querying the broker specified by:
```

```

Host          Primary Port
-----
localhost    7676

Version              4.4 Update 1
Instance Name       imqbroker
Broker ID
Primary Port       7676
Broker is Embedded  false
Instance Configuration/Data Root Directory /var/imq

Current Number of Messages in System 36
Current Total Message Bytes in System 19692

Current Number of Messages in Dead Message Queue 0
Current Total Message Bytes in Dead Message Queue 0

Log Dead Messages      false
Truncate Message Body in Dead Message Queue true

Max Number of Messages in System unlimited (-1)
Max Total Message Bytes in System unlimited (-1)
Max Message Size       70m

Auto Create Queues    true
Auto Create Topics    true
Auto Created Queue Max Number of Active Consumers unlimited (-1)
Auto Created Queue Max Number of Backup Consumers 0

Cluster ID
Cluster is Highly Available false
Cluster Broker List (active) mq://10.133.153.173:7676/
Cluster Broker List (configured)
Cluster Master Broker
Cluster URL

Log Level            WARNING
Log Rollover Interval (seconds) 604800
Log Rollover Size (bytes) 268435456

```

Successfully queried the broker.

You can also run this command:

```

imqcmd -u admin list dst
Listing all the destinations on the broker specified by:

```

```

-----
Host          Primary Port
-----
localhost    7676

-----
-----
Name      Type   State   Producers      Consumers      Msgs
          Total Wildcard Total Wildcard Count Remote UnAck
Avg Size
-----
jesms     Queue RUNNING 24    -      0    -      0    0    0
0.0

```

```
mq.sys.dmq Queue RUNNING 0 - 0 - 0 0 0
0.0
```

Successfully listed destinations.

Related Information

Use the following links to find more information about producing a Java program to listen and process Messaging Server events:

- Java Message Service Tutorial at:
<http://docs.oracle.com/javaee/1.3/jms/tutorial/>
- Java Message Server Programming Interface at:
<http://www.oracle.com/technetwork/java/jms/index.html>
- Java Message Queue Developer's Guide for Java Clients at:
<http://docs.oracle.com/cd/E19148-01/819-4469/index.html>

Part V

Managing the Message Store and Mailboxes

The message store is the component of Oracle Communications Messaging Server that contains the user mailboxes as well as the servers that provide IMAP, POP, and HTTP access to the mailboxes.

Note: Unless otherwise noted, the message store (specifically, the **stored** process) should be up and running while performing management and maintenance tasks described in "[Message Store Command Reference](#)."

The size of the message store increases as the number of users, mailboxes, and log files increase. You can control the size of the message store by specifying limits on the size of mailboxes, by specifying limits on the total number of messages allowed, and by setting aging policies for messages in the store.

Depending on the number of users your server supports, the message store might require one physical disk or multiple physical disks. There are two ways to integrate this additional disk space into your system. The easiest way is to add additional message store partitions (see "[Managing Message Store Partitions and Adding Storage](#)"). Likewise, if you are supporting multiple hosted domains, you might want to dedicate a server instance to a single, large domain. With this configuration, you can designate a store administrator for a particular domain. You can also expand the message store by adding more partitions.

Part V contains the following chapters:

- [Managing Mailboxes](#)
- [Backing Up and Restoring the Message Store](#)
- [Administering Very Large Mailboxes](#)
- [Message Store Message Expiration](#)
- [Configuring Message Expiration \(Tasks\)](#)
- [Configuring POP, IMAP, and HTTP Services](#)
- [Handling Message Store Overload](#)
- [Managing Message Store Partitions and Adding Storage](#)
- [Managing Message Store Quotas](#)
- [Managing Message Types in the Message Store](#)
- [Managing Shared Folders](#)

- Upgrading the Message Store
- Message Store Automatic Recovery On Startup
- Message Store Maintenance Queue
- Message Store Message Type Overview
- Migrating Mailboxes to a New System
- Monitoring Disk Space
- Protecting Mailboxes from Deletion or Renaming
- Reducing Message Store Size Due to Duplicate Storage
- Specifying Administrator Access to the Message Store
- Valid Message Store UIDs and Folder Names
- Message Store Automatic Failover with Database Replication
- Administering Message Store Database Snapshots (Backups)
- Messaging Server and Tiered Storage Overview
- Message Store Command Reference

Additional information about the Message Store can be found in the following chapters:

- Troubleshooting the Message Store
- Best Practices for Messaging Server and ZFS

Managing Mailboxes

This information describes how to list, create, remove, rename, move, and view information about mailboxes. It also describes how to find the directory location for a particular mailbox, restore expunged messages, and see how many users other than the mailbox owner have read messages in a shared IMAP folder. Specifically it describes the following utilities: "mboxutil," "hashdir," "readership" and consists of the following sections:

- [To Manage Mailboxes with mboxutil](#)
- [To Move Mailboxes to a Different Disk Partition](#)
- [To Remove Orphan Accounts](#)
- [To Find a Mailbox's Directory Using hashdir](#)
- [To Find Out How Many Users Have Read Messages in a Shared Folder](#)

To Manage Mailboxes with mboxutil

Use the "mboxutil" command to perform typical maintenance tasks on mailboxes. **mboxutil** tasks include the following:

- List mailboxes
- List and remove orphaned and inactive mailboxes
- Create mailboxes
- Rename mailboxes
- Move mailboxes from one partition to another
- Expunge mailboxes
- Restore expunged messages that have not been purged
- List personal mailbox subscriptions and unsubscribed mailboxes that no longer exist
- You can also use the "mboxutil" command to view information about quotas. For more information, see "[Managing Message Store Quotas](#)."

When an end user deletes a mailbox, all messages are expunged and purged according to the value of **store.cleanupage**. However, expunged messages can be restored by using the **mboxutil -R** command, as long as you have enabled the **store.mailboxpurgedelay** option. Expunged messages are moved to the new location when a mailbox is renamed.

Caution: Do not kill the **mboxutil** process in the middle of execution. If it is killed with **SIGKILL (kill -9)**, it may potentially require that every server get restarted and a recovery be done.

Examples

To list all mailboxes for all users:

```
mboxutil -l
```

To list all mailboxes and also include path and ACL information:

```
mboxutil -l -x
```

To create the default mailbox named **INBOX** for the user **daphne**:

```
mboxutil -c user/daphne/INBOX
```

To delete a mail folder named **projx** for the user **delilah**:

```
mboxutil -d user/delilah/projx
```

To delete the default mailbox named **INBOX** and *all mail folders* for the user **drucilla**:

```
mboxutil -d user/drucilla/INBOX
```

To rename the mail folder **memos** to **memos-april** for the user **desdemona**:

```
mboxutil -r user/desdemona/memos user/desdemona/memos-april
```

To move the mail account for the user **dimitria** to a new partition:

```
mboxutil -r user/dimitria/INBOX user/dimitria/INBOX partition
```

where *partition* specifies the name of the new partition. To move the mail folder named *personal* for the user **dimitria** to a new partition:

```
mboxutil -r user/dimitria/personal user/dimitria/personal partition
```

To Move Mailboxes to a Different Disk Partition

By default, mailboxes are created in the **primary** partition. If the partition gets full, additional messages cannot be stored. There are several ways to address the problem:

- Reduce the size of user mailboxes
- If you are using volume management software, add additional disks
- Create additional partitions ("[To Add a Message Store Partition](#)") and move mailboxes to the new partitions

If possible, we recommend adding additional disk space to a system using volume management software since this procedure is the most transparent for the user. However, you may also move mailboxes to a different partition.

To move mailboxes to a different partition, see the preceding **mboxutil -r** examples.

To Remove Orphan Accounts

To search for orphaned accounts (orphaned accounts are mailboxes that do not have corresponding entries in LDAP) use the following "mboxutil" commands:

```
mboxutil -o
```

Command output follows:

```
mboxutil: Start checking for orphaned mailboxes
user/annie/INBOX
user/oliver/INBOX
mboxutil: Found 2 orphaned mailbox(es)
mboxutil: Done checking for orphaned mailboxes
```

Use the following command to create a file listing orphaned mailboxes that can be turned into a script file that deletes the orphaned mailboxes (example filename is **orphans.cmd**):

```
mboxutil -o -w orphans.cmd
```

The command output is as follows:

```
mboxutil: Start checking for orphaned mailboxes
mboxutil: Found 2 orphaned mailbox(es)
mboxutil: Done checking for orphaned mailboxes
```

Delete the orphan files with the following command:

```
mboxutil -d -f orphans.cmd
```

To Find a Mailbox's Directory Using hashdir

The mailboxes in the message store are stored in a hash structure for fast searching. Consequently, to find the directory that contains a particular user's mailbox, use the "hashdir" utility.

This utility identifies the directory that contains the message store for a particular account. This utility reports the relative path to the message store, such as **d1/a7/**. The path is relative to the directory level just before the one based on the user ID. The utility sends the path information to the standard output.

For example, to find the relative path to the mailbox for user **crowe**:

```
hashdir crowe
```

To Find Out How Many Users Have Read Messages in a Shared Folder

The "readership" utility reports on how many users other than the mailbox owner have read messages in a shared IMAP folder (see "Shared Folders Overview").

An owner of a IMAP folder may grant permission for others to read mail in the folder. A folder that others are allowed to access is called a **shared folder**. Administrators can use the **readership** utility to see how many users other than the owner are accessing a shared folder.

This utility scans all mailboxes and produces one line of output per shared folder, reporting the number of readers followed by a space and the name of the mailbox.

Each reader is a distinct authentication identity that has selected the shared folder within the past specified number of days. Users are not counted as reading their own personal mailboxes. Personal mailboxes are not reported unless there is at least one reader other than the folder's owner.

For example, the following command counts as a reader any identity that has selected the shared IMAP folder within the last 15 days:

```
readership -d 15
```

Backing Up and Restoring the Message Store

This chapter describes how to back up and restore mailboxes. For conceptual information on the message store, see the following:

- [Administering Message Store Database Snapshots \(Backups\)](#)
- [Message Store Disaster Backup and Recovery and Message Store Directory Layout](#)

Topics:

- [Mailbox Backup and Restore Overview](#)
- [To Create a Mailbox Backup Policy](#)
- [To Create Backup Groups](#)
- [To Run the imbackup Utility](#)
- [To Restore Mailboxes and Messages](#)
- [To Use StorageTek Enterprise Backup Software](#)
- [To Use a Third Party Backup Software \(Besides StorageTek Enterprise Backup Software\)](#)
- [Troubleshooting Backup and Restore Problems](#)
- [Message Store Disaster Backup and Recovery](#)

Mailbox Backup and Restore Overview

Mailbox backup and restore is one of the most common and important administrative tasks. You must implement a backup and restore policy for your message store to ensure that data is not lost if the following problems occur:

- System crashes
- Hardware failure
- Accidental deletion of messages or mailboxes
- Problems when reinstalling or upgrading a system
- Natural disasters (for example, earthquakes, fire, hurricanes)
- Migrating users

You can back up and restore mailboxes by using the command-line utilities "imbackup" and "imsrestore," or the integrated backup and restore solution that uses Oracle StorageTek Enterprise Backup Software (EBS).

Oracle Communications Messaging Server provides a single-copy backup procedure. Regardless of how many user folders contain a particular message, during backup, the message file is backed up only once using the first message file found. The second message copy is backed up as a link to the name of the first message file, and so on. **imsbackup** maintains a hash table of all messages using the device and inode of the message files as the index. This method does have implications when restoring data, however. For more information, see "[Considerations for Partial Restore](#)."

Note: You can also back up and restore the message store by backing up all relevant message files and directories. See "[Message Store Disaster Backup and Recovery](#)" for more information.

Backing up mailboxes includes three steps:

1. [To Create a Mailbox Backup Policy](#)
2. [To Create Backup Groups](#)
3. [To Run the imsbackup Utility](#)

To Create a Mailbox Backup Policy

Your backup policy will depend on several factors, such as:

- [Peak Business Loads](#)
- [Full and Incremental Backups](#)
- [Parallel or Serial Backups](#)

Peak Business Loads

Take into account peak business loads when scheduling backups for your system as this can reduce system load during peak hours. For example, backups are probably best scheduled for early morning hours such as 2:00 AM.

Full and Incremental Backups

Incremental backups (see "[Incremental Backup](#)") scan the message store for changed data and back up only what has changed. Full backups back up the entire message store. Determine how often the system should perform full as opposed to incremental backups. For example, you probably want to perform incremental backups as a daily maintenance procedure and full backups once a week.

Parallel or Serial Backups

When user data is stored on multiple disks, you can back up user groups in parallel. Depending on system resources, parallel backups can speed up the overall backup procedure. However, you might want to use serial backups to reduce backup impact on the server's performance. Whether to use parallel or serial backups can depend on many factors, including system load, hardware configuration, how many tape drives are available, and so on.

To Create Backup Groups

A backup group is an arbitrary set of user mailboxes defined by regular expressions. By organizing user mailboxes into backup groups, you can define more flexible backup management.

For example, you could create three backup groups, the first containing user IDs starting with the letters A through L, the second with users whose user IDs begin with M through Z, and the third with users whose user IDs begin with a number. Administrators could use these backup groups to back up mailboxes in parallel, or perhaps only certain groups on one day and other groups on another.

Consider the following points about backup groups:

1. They are arbitrary *virtual* groups of mail users that do not precisely map to the "Message Store Directory Layout," although backup groups could resemble the message store directory.
2. Administrators define backup groups by using UNIX regular expressions. The regular expressions are defined in the *MessagingServer_home/config/backup-groups.conf* file.
3. When backup groups are referenced in **imsbackup** and **imsrestore**, they use the path format: */partition_name/backup_group*
4. When you run the **imsbackup** command, it evaluates the entire **backup-groups.conf**, and if it finds more than one group that matches a user, it uses the first match. For example, the following **backup-groups.conf** contains these definitions:

```
groupA=a.*
...
groupN=. *n$
```

Because both groups match the user ID **admin**, the **imsbackup** command uses the first match, which is **groupA**. Thus, **groupA** includes the **admin** mailbox. Furthermore, the **groupN** backup does not include the **admin** mailbox.

The format of **backup-groups.conf** is as follows:

```
group_name=definition
group_name=definition
.
.
.
```

Using the example described in the previous paragraph, you would use the following definitions to create the three backup groups:

```
groupA=[a-l].*
groupB=[m,-z].*
groupC=[0-9].*
```

You can now scope **imsbackup** and **imsrestore** at several levels. You can backup the whole message store by using the following backup commands:

```
imsbackup -f <device> /
```

To back up all mailboxes for all users in **groupA** use the following command:

```
imsbackup -f <device> /<partition>/groupA
```

The default partition is called **primary**.

Pre-defined Backup Group

Oracle Communications Messaging Server includes one predefined backup group that is available without creating the **backup-groups** configuration file. This group is called **user** and includes all users. For example, the following command backs up all users on the **primary** partition:

```
imbackup -f backupfile /primary/user
```

To Run the imbackup Utility

To back up and restore your mailboxes, Messaging Server provides the "imbackup" and "imsrestore" utilities. The **imbackup** and **imsrestore** utilities do not have the advanced features found in general purpose tools like StorageTek Enterprise Backup Software (EBS). For example, the utilities have only very limited support for tape auto-changers, and they cannot write a single store to multiple concurrent devices. Comprehensive backup is achieved by using plug-ins to generalized tools like EBS. For more information about using EBS, see "[To Use StorageTek Enterprise Backup Software.](#)"

Running the imbackup Utility

With **imbackup**, you can write selected contents of the message store to any serial device, including magnetic tape, a UNIX pipe, or a plain file. The backup or selected parts of the backup can later be recovered by using the **imsrestore** utility. The output of **imbackup** can be piped to **imsrestore**.

The following example backs up the entire message store to **/dev/rmt/0**:

```
imbackup -f /dev/rmt/0 /
```

This example backs up the mailboxes of user ID **joe** to **/dev/rmt/0**:

```
imbackup -f /dev/rmt/0 /primary/user/joe
```

This example backs up all the mailboxes of all the users defined in the backup group **groupA** to **backupfile** (see "[To Create Backup Groups](#)"):

```
imbackup -f /primary/groupA > backupfile
```

Incremental Backup

The following example backs up messages stored from May 1, 2004 at 1:10 pm to the present. The default is to back up all the messages regardless of their dates:

```
imbackup -f /dev/rmt/0 -d 20040501:131000 /
```

This command uses the default blocking factor of 20. For a complete syntax description, see "[imbackup.](#)"

Regarding date-time stamp:

```
20040501:131000  
YYYYMMDD:HHMMSS
```

```
2004 05 01 : 13 10 00  
YYYY MM DD : HH MM SS
```

Excluding Bulk Mail When You Perform Backups

When you perform a backup operation, you can specify mailboxes that will be excluded from being backed up. By excluding bulk or trash mailboxes that can accrue large numbers of unimportant messages, you can streamline the backup session, reduce the time to complete the operation, and minimize the disk space required to store the backup data.

To exclude mailboxes, specify a value for the **store.backupexclude** option.

You can specify a single mailbox or a list of mailboxes separated by the '%' character. ('%' is an illegal character in a mailbox name.) For example, you could specify the following values:

Trash

Trash%Bulk Mail%Third Class Mail

In the first example, the folder **Trash** is excluded. In the second example, the folders **Trash**, **Bulk Mail**, and **Third Class Mail** are excluded.

Example commands:

```
cd /opt/sun/comms/messaging64/bin
msconfig set store.backupexclude "Trash%Bulk Mail%Third Class Mail"
msconfig show store.backupexclude
role.store.backupexclude = Trash%Bulk Mail%Third Class Mail
```

The backup utility backs up all folders in a user mailbox except those folders specified in the **store.backupexclude** option.

This feature works with the Messaging Server backup utility, StorageTek Enterprise Backup Software, and third-party backup software.

You can override the **store.backupexclude** setting and back up an excluded mailbox by specifying its full logical name during the operation. For example, suppose the Trash folder has been excluded. You can still back up Trash by specifying the following:

```
/primary/user/user1/trash
```

However, if you specify

```
/primary/user/user1
```

the Trash folder is excluded.

To Restore Mailboxes and Messages

To restore messages from the backup device, use the **"imsrestore"** command. For example, the following command restores messages for **user1** from the file **backupfile**.

```
imsrestore -f backupfile /primary/user1
```

Considerations for Partial Restore

A partial restore is when only a part of the message store is restored. A full restore is when the entire message store is restored. The message store uses a single-copy message system. That is, only a single copy of any message is saved in the store as a single file. Any other instances of that message (for example, when a message is sent to

multiple mailboxes) are stored as links to that copy. Because of this, there are implications when restoring messages. For example:

- **Full Restore.** During a full restore, linked messages still point to the same inode as the message file to which they are linked.
- **Partial Backup/Restore.** During a partial backup and partial restore, however, the single-copy characteristic of the message store might not be preserved.

The following examples demonstrate what happens to a message that is used by multiple users when a partial restore is performed. Assume there are three messages, all the same, belonging to three users A, B, and C, as follows:

```
A/INBOX/1  
B/INBOX/1  
C/INBOX/1
```

Example 1. In the first example, the system performs a partial backup and full restore procedure as follows:

1. Back up mailboxes for users B and C.
2. Delete mailboxes of users B and C.
3. Restore the backup data from step 1.

In this example, **B/INBOX/1** and **C/INBOX/1** are assigned a new inode number and the message data is written to a new place on the disk. Only one message is restored. The second message is a hard link to the first message.

Example 2. In this example, the system performs a full backup and a partial restore as follows:

1. Perform full backup.
2. Delete mailboxes for user A.
3. Restore mailboxes for user A.

A/INBOX/1 is assigned a new inode number.

Example 3. In this example, partial restore might require more than one attempt:

1. Perform full backup.
B/INBOX/1 and **C/INBOX/1** are backed up as links to **A/INBOX/1**.
2. Delete mailboxes for users A and B.
3. Restore mailboxes for user B.

The restore utilities ask the administrator to restore **A/INBOX** first.

4. Restore mailboxes for users A and B.
5. Delete mailboxes for user A (optional).

Note: If you want to ensure that all messages are restored for a partial restore, you can run the **imsbackup** command with the **-i** option. The **-i** option backs up every message multiple times if necessary.

If the backup device is seekable (for example, a drive or tape), **imsrestore** seeks to the position containing **A/INBOX/1** and restores it as **B/INBOX/1**. If the backup device is non-seekable (for example, a UNIX pipe), **imsrestore** logs the object ID and the ID of the depending (linked) object to a file, and the administrator must invoke **imsrestore** again with the **-r** option to restore the missing message references.

To Restore Messages from a Mailbox that Has Been Incrementally Backed-up

If you are restoring messages from a mailbox that has been incrementally backed-up, and if that mailbox exists on the server on which you want to restore the messages, then restoring the messages requires a straightforward **imesrestore**. However, if you want to restore messages from a mailbox that has been incrementally backed-up, and if that mailbox no longer exists, you must follow different restore procedures.

Use one of the following procedures to restore messages to a mailbox that does not exist on the message store server:

- During the restore operation, disable delivery of messages to the user. Do this by setting the LDAP attribute **mailDeliveryOption** to **hold**.
- Before you use **imesrestore**, create the mailbox with the **mboxutil -c** command.

The reason why these instructions must be followed for restoring an incremental backup is as follows: When a mailbox has been deleted or is being migrated, the **imsrestore** utility recreates the mailbox with the mailbox unique identification validity and message unique identifications (UIDs) stored in the backup archive.

In the past, when **imsrestore** would recreate a deleted or migrated mailbox, it would assign a new UID validity to the mailbox and new UIDs to the messages. In that situation, a client with cached messages would have to resynchronize the mailbox UID validity and message UIDs. The client would have to download the new data again, increasing the workload on the server.

With the new **imesrestore** behavior, the client cache remains synchronized, and the restore process operates transparently with no negative impact on performance.

If a mailbox exists, **imsrestore** assigns new UIDs to the restored messages so that the new UIDs remain consistent with the UIDs already assigned to existing messages. To ensure UID consistency, **imsrestore** locks the mailbox during the restore operation. However, because **imsrestore** now uses the mailbox UID validity and message UIDs from the backup archive instead of assigning new UID values, UIDs could become inconsistent if you perform incremental backups and restores.

If you perform incremental backups with the **-d** date option of the **imsbackup** utility, you might have to invoke **imesrestore** multiple times to complete the restore operation. If incremental backups were performed, you must restore the latest full backup and all subsequent incremental backups.

New messages can be delivered to the mailbox between the restore operations, but in this case, the message UIDs can become inconsistent. To prevent inconsistency in the UIDs, you need to take one of the actions previously described on this page.

To Use StorageTek Enterprise Backup Software

Messaging Server includes a backup API that provides an interface with third-party backup tools, such as EBS. The physical message store structure and data format are encapsulated within the backup API. The backup API interacts directly with the message store. It presents a logical view of the message store to the backup service. The backup service uses the conceptual representation of the message store to store and retrieve the backup objects.

Messaging Server provides an Application Specific Module (ASM) that can be invoked by the EBS's **save** and **recover** commands to back up and restore the message store data. The ASM then invokes the Messaging Server **imsbackup** and **imsrestore** utilities.

Note: This section provides information about how to use EBS with the Messaging Server message store. To understand the EBS interface, see your StorageTek Enterprise Backup Software documentation.

To Back Up Data By Using StorageTek Enterprise Backup Software

1. Create a symbolic link from `/usr/lib/nsr/imsasm` to `MessagingServer_home/lib/msg/imsasm`.
2. From Oracle or EMC, obtain a copy of the `nsrfile` binary and copy it to the following directory: `/usr/bin/nsr`

This is required only if you are using an older version of Networker (5.x). With Networker 6.0 and above, `nsrfile` is automatically installed under `/usr/bin/nsr`.

3. If you want to back up users by groups, perform the following steps:
 - a. Create a backup group file as described in "To Create Backup Groups."
 - b. To verify your configuration, run `mkbackupdir.sh`. Look at the directory structure created by `mkbackupdir.sh`. The structure should look similar to that shown in "Message Store Directory Layout." If you do not specify a `backup-groups.conf` file, the backup process uses the default backup group `ALL` for all users.
4. In the directory `/nsr/res/`, create a `res` file for your save group to invoke the `mkbackupdir.sh` script before the backup. See "Message Store Directory Layout" for an example.

Note: Earlier versions of Networker have a limitation of 64 characters for the save set name. If the name of this directory plus the logical name of the mailbox (for example, `/primary/groupA/fred`) is greater than 64 characters, then you must run `mkbackupdir.sh -p`. Therefore, you should use a short path name for the `-p` option of `mkbackupdir.sh`. For example the following command will create the backup image under the directory `/backup`:

```
mkbackupdir.sh -p /backup
```

Important: The backup directory must be writable by the message store owner (example: `mailsrv`).

The following is a sample backup groups directory structure.

```
/backup/primary/groupA/amy
```

```

/bob
/carly
/groupB/mary
/nancy
/zelda
/groupC/123go
/lbill
/354hut

```

The following example shows a sample **res** file named **IMS.res** in the **/nsr/res** directory:

```

type: savepnpc;
precmd: "echo mkbackupdir started",
"/usr/example/server5/msg-example/bin/mkbackupdir.sh -p /backup";
pstcmd: "echo imsbackup Completed";
timeout: "12:00 pm";

```

You are now ready to run the EBS interface as follows:

5. Create the Messaging Server save group if necessary.
 - a. Run **nwadmin**.
 - b. Select Customize | Group | Create.
6. Create a backup client using **savepnpc** as the backup command:
 1. Set the save set to the directory created by **mkbackupdir**. For a single session backup, use **/backup**.
For parallel backups, use **/backup/server/group**. Be sure you have already created *group* as defined in "To Create Backup Groups." You must also set the parallelism to the number of backup sessions. See "To Back Up Data By Using StorageTek Enterprise Backup Software."
7. Select Group Control | Start to test your backup configuration. Example. Creating A Backup Client in EBS: To create a backup client in EBS. From **nwadmin**, select Client | Client Setup | Create

```

Name: example
Group: IMS
Savesets:/backup/primary/groupA
/backup/secondary/groupB
/backup/tertiary/groupC
.
.
Backup Command:savepnpc
Parallelism: 4

```

Restoring Data Using StorageTek Enterprise Backup Software

To recover data, you can use the EBS **nwrecover** interface or the **recover** command-line utility. The following example recovers user **a1**'s INBOX:

```
recover -a -f -s example /backup/example/groupA/a1/INBOX
```

The next example recovers the entire message store:

```
recover -a -f -s example /backup/example
```

To Use a Third Party Backup Software (Besides StorageTek Enterprise Backup Software)

Messaging Server provides two message store backup solutions, the command line **imsbackup** and the StorageTek Enterprise Backup Software. A large message store running a single **imsbackup** to back up the entire message store can take a significant amount of time. The EBS solution supports concurrent backup sessions on multiple backup devices. Concurrent backup can shorten backup time dramatically (backups of 25GB of data per hour have been achieved).

If you are using another third party concurrent backup software (for example, Netbackup), you can use the following method to integrate your backup software with the Messaging Server.

1. Divide your users into groups (see "[To Create Backup Groups](#)") and create a **backup-groups.conf** file under the directory *MessagingServer_home/config/*.

Note: This backup solution requires additional disk space. To backup all the groups concurrently, the disk space requirement is two times the message store size. If you do not have that much disk space, divide your users into smaller groups, and then backup a set of groups at a time. For example group1 - group5, group6 - group10. Remove the group data files after backup.

2. Run **imsbackup** to back up each group into files under a staging area. The command is **imsbackup -f <device>/<instance>/<group>** You can run multiple **imsbackup** processes simultaneously. For example:

```
imsbackup -f- /primary/groupA > /bkdata/groupA &
imsbackup -f- /primary/groupB > /bkdata/groupB &
. . .
```

3. Use your third party backup software to back up the group data files in the staging area (in our example that is **/bkdata**).
4. To restore a user, identify the group filename of the user, restore that file from tape, and then use **imsrestore** to restore the user from the data file.

Troubleshooting Backup and Restore Problems

This section describes common backup and restore problems and their solutions.

- **Problem:** **msprobe** restarts everything during a long **imsrestore** during message store migration. This can also happen with **imsbackup**, **imsimport**, or any processing intensive utility.
- **Solution:** When **imsrestore** or any processing intensive operation takes significantly more system resources than normal, and continues doing so longer than the **msprobe** interval, there might be a temporary backlog of DB transaction log files to be cleared. If there are more files than specified in **store.maxlog**, then **msprobe** may erroneously restart all the processes during a restore. To prevent this from happening, disable **msprobe** during the **imsrestore**.
- **Problem:** When I do an restore of a folder or INBOX using **imsrestore** or **imsasm**, it appends all the messages in that folder onto the current folder. This results in multiple copies of the messages in that folder.
- **Solution:** Make sure the **-i** flag of **imsrestore** is not set in the **imsasm** script.

- **Problem:** I want to do an incremental backup of just new messages added in a mail folder, but when I try, the entire folder gets backed up. How do I just back up the new messages?
- **Solution:** Set the `-ddatetime` flag on `imsbackup`. This will backup messages stored from the specified date and time to the present. The default is to back up all messages regardless of their dates.

Message Store Disaster Backup and Recovery

A disaster refers to a catastrophic failure of the entire message store as opposed to a mailbox or set of mailboxes. That is, a situation where all data on the message store servers are lost. A complete message store disaster restore will consist of restoring the following lost data:

- All message store data. These can be backed up using the procedures described above. If file system backup method is used, be sure to back up the following data:
 - All message store partitions
 - The message store database files at `MessagingServer_home/data/store/mbxlist`.
 - The message store database snapshots at `MessagingServer_home/data/store/dbdata/snapshots` (Note that the location of message store database snapshot files can be configured with the `store.snapshotpath` option.)
- Configuration data. Including the local configuration file at `MessagingServer_home/data/config`. See also "Message Store Directory Layout."

If you want to back up your message store for disaster recovery, you can use file system snapshot tools to take a snapshot of the file system. The snapshot **must** be a *point-in-time* file system snapshot.

It is best to capture all the data (message store partitions, database files and so on) at the same point-in-time, however, if this cannot be done, then you must backup the data in this order:

1. Database snapshots
2. Database files
3. Message Store indexes if separated from the messages using the `partition:partition-name.messagepath` option (`store.partition.*.messagepath` in legacy configuration). For additional information, see the discussion about the `messagepath` option in *Messaging Server Reference*.
4. Messages
5. Configuration data

If database files and Message Store indexes are not backed up with the same point-in-time snapshot (or database files, Message Store indexes, and Messages, if indexes are not separate) then `reconstruct -m` is required after restore.

Administering Very Large Mailboxes

Oracle Communications Messaging Server enables the use of "very large" mailboxes.

Topics:

- [Very Large Mailboxes Overview](#)
- [The Structure of a Mailbox](#)
- [Mailbox Size Limit](#)
- [Mailbox Migration](#)
- [Pre-Deployment Preparations](#)
- [Checking Mailbox Data](#)


Very Large Mailboxes Overview

To increase the mailbox size limit and improve **expunge** performance, the index records and cache records have been split into separate files with support added for multiple cache files. An index file contains a mailbox header and a 128-byte fixed-length record for every message in the mailbox. The index record contains the location and size of the corresponding cache record. The cache files contain frequently used data in variable length records. The cache file size is configurable (**store.maxcachefiles**). The default cache size is 500 Mbs, with a maximum size 2 Gbs. New cache records are appended to the newest cache file. As a cache file fills up, a new cache file is created, enabling a mailbox to continue to grow.

To optimize expunge performance, only the index file **store.idx** is purged when a mailbox is expunged. Expunge removes the obsoleted index records from the index file. Cache record and message file removal is deferred until the size of the expunged data exceeds a configurable threshold (**store.purge.count** or **store.purge.percentage** for cache records; **store.cleanupsize** for the *number* of messages). When the expunged size exceeds the threshold, expunge enqueues a purge request to the "Message Store Maintenance Queue." impurge dequeues the request, removes the unused message files and purge the cache files when the expunged size exceeds the purge threshold.

The Structure of a Mailbox

The following figure is an example of a large mailbox:

Figure 53–1 Large Mailbox Example

Mailbox Size Limit

IMAP defines the UID as a 32-bit value. Therefore, the maximum number of messages in a mailbox is limited to 4,294,967,295. For a 64 bit messaging server, the maximum number of messages in a folder is 4,294,967,295. For a 32-bit messaging server, the maximum number of messages in a folder is 16,777,215. The **msconfig** option **store.maxmessages** can be used to limit the size of a folder. The default limit is 16,000,000.

Note, the 4 billion limit is a hard limit. Effective limit is much smaller (normally less than 1 million). Mailbox expunge has to rewrite the **store.idx** file. The larger the mailbox, the longer it takes to expunge.

Mailbox Migration

Mailboxes are migrated to the new format automatically when they are opened. The operation is transparent to the end users.

Pre-Deployment Preparations

The high-level steps for preparing to use very large mailboxes include:

1. Configuring the maximum cache file size
2. Configuring the mailbox expunge size
3. Configuring **store.maxmessages**, quotas or expire rules to prevent mailboxes from getting too big

The new **msconfig** options for managing large mailboxes are:

Table 53–1 msconfig Options for Managing Large Mailboxes

msconfig Option	Description
store.maxcachefilesize	Sets the maximum cache file size in bytes, maximum of 2 Gbs.

Table 53-1 (Cont.) msconfig Options for Managing Large Mailboxes

msconfig Option	Description
store.cleanupsize	Cleans the mailbox when the number of expunged messages exceeds this value. Default is 100.

Checking Mailbox Data

The **imcheck** command line utility can be used to dump the data of mailboxes in readable format. For example, **imcheck -m mailbox** dumps the content of a **store.idx** file:

```
$ imcheck -m user/dumbo/INBOX
-----
Name: user/dumbo
Version: 103
Exists: 4
Flags: 0
Largest Msg: 1521 bytes
Last Append: 20080131104858
Last Repair: -
Last UID: 4
Oldest Msg: 20080131104843
Oldest Uid: 1
Quota Used: 2394
Bytes Expunged: 0
UID Validity: 1201805323
Last CacheId: 1
Start Offset: 256
Append CacheId: 1
ACL: dumbo lrswipcdan
Subscribed: 0
Partition: primary
Path: /var/opt/SUNWmsgsr/store/partition/primary/=user/94/60/=dumbo
Msg Path: /var/opt/SUNWmsgsr/store/partition/primary/=user/94/60/=dumbo

MsgNo Uid Internal-Date Sent-Date Size HSize Cache-Id C-Offset C-Len Last-Updated
Save-Date MT SFlags UFlags Original-Uid Message-id
-----
1 1 20080131104843 20080131104843 257 244 1 16 864 20080131104843 20080131104843 2
R 0.0.0 1201805323-1 -
2 2 20080131104851 20020301191039 338 229 1 880 816 20080131104851 20080131104851
2 R 0.0.0 1201805323-2 <001@red.iplanet.com>
3 3 20080131104855 20020301191039 1521 241 1 1696 840 20080131104855
20080131104855 2 R 0.0.0 1201805323-3 <002@red.iplanet.com>
4 4 20080131104858 20050919110532 278 252 1 2536 860 20080131104858 20080131104858
1 R 0.0.0 1201805323-4 <003@red.iplanet.com>
```

imcheck -m mailbox -c msgno displays the cache records of a message:

```
$ imcheck -m user/dumbo/INBOX -c 1
Cache items of user/dumbo message number 1:

ENVELOPE {300}
("Thu, 31 Jan 2008 10:48:43 -0800" "welcome" (("Mail Administrator" NIL
"Postmaster" "puzzle.red.iplanet.com")) ("Mail Administrator" NIL "Postmaster"
"puzzle.red.iplanet.com")) ("Mail Administrator" NIL "Postmaster"
"puzzle.red.iplanet.com")) ((NIL NIL "dumbo" "red.iplanet.com")) NIL NIL NIL NIL)
```

BODYSTRUCTURE {75}
("TEXT" "PLAIN" ("CHARSET" "us-ascii") NIL NIL "7BIT" 13 1 NIL NIL NIL NIL)

BODY {59}
("TEXT" "PLAIN" ("CHARSET" "us-ascii") NIL NIL "7BIT" 13 1)

SECTION {48}
Header offset: 0 len: 244
Body offset: 244 len: 13
1 subparts
(1) offset: 244 len: 13 charset: 0 encoding: 0

CACHEHEADERS {244}
Subject: welcome
To: dumbo@red.iplanet.com
Date: Thu, 31 Jan 2008 10:48:43 -0800
From: Mail Administrator <Postmaster@puzzle.red.iplanet.com>
MIME-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit

FROM {53}
mailadministrator <postmaster@puzzle.red.iplanet.com>

TO {23}
<dumbo@red.iplanet.com>

CC {0}

BCC {0}

SUBJECT {9}
"welcome"

XSENDER {0}

Message Store Message Expiration

This chapter describes message expiration concepts. See "[Configuring Message Expiration \(Tasks\)](#)" for information on message removal tasks.

Message expiration automatically removes messages from the message store based on criteria that you set. For example, you can remove old messages, overly large messages, seen or deleted messages, messages with specific Subject: lines, messages of a certain type, and so on.

Note: Oracle Communications Messaging Server removes messages without giving a warning, so it is important to inform users about message expiration policies. Unexpected message removal can be a source of consternation for users and administrators.

The topics in this chapter include:

- [imexpire Theory of Operation](#)
- [To Deploy the Message Expiration Feature](#)
- [Localized Mailbox Names in imexpire](#)

imexpire Theory of Operation

Message expiration is performed by the **imexpire** utility, which performs a specific action to the expired messages. (See "[Deleting, Expunging, Purging, and Cleaning Up Messages](#)" for details on the message removal process.) You can launch the **imexpire** utility from the command-line or schedule it to launch through the **imsched** daemon. You specify a set of expiration rules in the **store.expirerule** file. You can have multiple rules files, each located the directory that pertains to the scope of the rules. That is, rules that apply globally to the entire message store are put in one directory, rules that apply to a partition in another, rules that apply to users in yet another, and so on.

Note: Although global expiration rules can be specified with **store.expire.attribute configutil** options, use **store.expirerule** files to specify these rules. If too many rules are created by using **configutil**, performance problems can result.

imexpire loads all of the expire rules at start up. By default, **imexpire** creates one thread per partition. Each thread goes through the list of user folders under its assigned partition and loads the local expire rule files as it goes. The expire function

checks each folder against the expire rules applicable to this folder and expunges messages as needed.

It is also possible to exclude specified users from the expire rules by adding their user ID, one per line, in a file called **expire_exclude_list** in the *MessagingServer_home/config/* directory.

To Deploy the Message Expiration Feature

Message expiration requires the following three steps:

1. Define message expiration policy: Which messages will be expired? What users, folders, domains, and partitions will have messages expired? What size, message age, and headers will define the removal criteria? Define the scope of messages to be removed. See ["To Define Message Expiration Policy."](#)
2. Specify the **imexpire** rules to implement this policy. See ["To Set Rules Implementing Message Expiration Policy."](#)
3. Specify the **imexpire** scheduling. See ["To Schedule Message Expiration and Logging Level."](#)

To Define Message Expiration Policy

You can define message expiration based on criteria such as:

- **Age of Message** – Expire messages older than X days. Attribute: **messagedays**.
- **Message Count** – Expire messages in a folder exceeding X messages. Attribute: **messagecount**.
- **Age of Oversized Message** – Expire messages that exceed X bytes after Y days grace period. Attributes: **messagesize** and **messagesizedays**.
- **Seen and Deleted Message Flag** – Expire messages with the Seen or Deleted flag set. These criteria can be set to **"and"** or **"or."** If set to **or**, the message's Seen/Delete flag will cause expiration regardless of other criteria. If set to **and**, the message's Seen/Delete flag must be set along with passing all other specified criteria. Attributes: **seen** and **deleted**.
- **Header Field of Message** – Allows you to specify a header and string as criteria for expiring a message, for example, removing all messages with the header "Subject: Work from Home." Note that this feature also allows you to use message type as a criteria too. See ["Expiring Messages by Message Type."](#)
- **Folder of Messages** – Allows you to specify the folder on which to expire messages. Attribute: **folderpattern**. Note that this attribute only uses the modified UTF-7 character set.

Note: **imexpire** does not allow you to delete or preserve messages based on how long it has been since that message was read. For example, you cannot specify that messages that have not been read for 200 days will be removed.

Examples of Message Expiration Policy

Example 1: Remove all messages 365 days old in a folder exceeding 1,000 messages.

Example 2: Remove messages in domain **example.org** that are older than 180 days.

Example 3: Remove all messages that have been marked as deleted.

Example 4: Remove messages in example.com that have been marked as seen, are older than 30 days, are larger than 100 kilobytes, from folders exceeding 1,000 messages, with the header **X-spam**.

To Set Rules Implementing Message Expiration Policy

Rules are set by putting them into a **store.expirerule** file. An example of two global **store.expirerule** rules is shown below:

```
Rule1.regex: 1
Rule1.folderpattern: user/.*/trash
Rule1.messagedays: 2
Rule2.regex: 1
Rule2.folderpattern: user/.*
Rule2.messagedays: 14
```

In this example, Rule 1 specifies that all messages in all users' trash folders are removed after two days. Rule 2 specifies that all messages in any folder in the message store are removed after 14 days.

Expiration Rules Guidelines

This section sets the guidelines for the **store.expirerule** file rules.

Note: In earlier Messaging Server releases, expiration rules could be set with **configutil** options **store.expirerule.attribute** (see the *Messaging Server Reference*.) This is still true, but expire rules using header constraints (example: expiring a message with a specific subject line) are not supported. Also, regular expressions in the expire rules created with **configutil** need to be POSIX compliant rules. If you want to use UNIX compliant regular expressions you must use the **store.expire** file. In addition, using both **configutil** options and the global **store.expirerule** configuration file is not supported. If the configuration file is present, **configutil** options are not used. In any case, it is best to use **store.expirerule** to specify all expiration rules.

- Rules are specified in a file called **store.expirerule**.
- Multiple expiration criteria can be specified with the same rule. (See preceding example.)
- Rules can apply to the entire message store (global rules), a partition, a user, or a folder.
 - The global rules are stored in the *MessagingServer*_home/**config/store.expirerule** file.

Note: Each global rule will be checked against every mailbox, which can cause some processing overhead depending on the number of global rules you specify. For this reason, you should not specify partition, mailbox or user rules in the global rules file. In general, you should try not to put any more expiration rules than necessary in this file.

- Partition rules are stored in *store_root/partition/partition_name/store.expirerule* (more accurately, the location specified by the **store.partition.*.path configutil** option).
- User rules are specified in *store_root/partition/partition_name/userid/store.expirerule* or by specifying the **folderpattern** rule to be **user/userid.***.
- Folder rules are specified in *store_root/partition/partition_name/userid/folder/store.expirerule* or by specifying the **folderpattern** rule to be **user/userid/folder**.

Multiple non-global rules (user, folder, partition) using *rule_name* was only implemented starting in Messaging Server 6.2p4.

- Multiple expire rules can be applied to a mailbox at the same time. An expire policy for a mailbox consists of global rules and local rules. Local rules apply to the mailbox under the same directory and all of its sub-folders.
- **imexpire** unifies all of the expiration rules applying to a mailbox, unless there is an exclusive rule specified for this mailbox (see Table 54–1). The resulting rule set represents the most restrictive expiration policy based on all applicable rules. For example, if rule X expires messages such that the maximum message life is 10 days, and rule Y specifies 5 days, the union will be 5 days.

Table 54–1 imexpire Attributes

Attribute	Description (Attribute Value)
action	Specifies an action to perform on the messages caught by the expire rules. The possible values are: <ul style="list-style-type: none"> ■ discard – discards the message. This is the default. ■ report – prints the mailbox name, uid-validity and uid to stdout. ■ archive – archives the message with the Compliance and Content Management System and then discards the message. ■ fileinto: folder – files the message into the specified folder. The shared folder prefix can be used to file messages to folders owned by another user.
exclusive	Specifies whether or not this is an exclusive rule. If specified as exclusive , only this rule applies to the specified mailbox(es) and all other rules are ignored. If more than one exclusive rule exists, the last exclusive rule loaded will be used. For example, if a global and a local exclusive rule are specified, the local rule will be used. If there is more than one global exclusive rule, the last global rule listed by configutil is used. (1/0)
expires	imexpire will select the message if the date value specified with these header fields is older than the expiration date based on the messagedays attribute. If multiple expiration header fields are specified, the earliest expiration date will be used. (string)
expiry-date	imexpire will select the message if the date value specified with these header fields is older than the expiration date based on the messagedays attribute. If multiple expiration header fields are specified, the earliest expiration date will be used. (string)
foldpattern	Specifies the folders affected by this rule. The format must start with a user/ , which represents the directory <i>store_root/partition/*</i> . See Table 54–2, "imexpire Folder Patterns Using Regular Expressions". (POSIX regular expression)

Table 54–1 (Cont.) *imexpire* Attributes

Attribute	Description (Attribute Value)
messagecount	Maximum number of messages in a folder. Oldest messages are expunged as additional messages are delivered. (integer)
foldersize	Maximum size of folder before the oldest messages are expunged when additional messages are delivered. (integer in bytes)
messagedays	Number of days in the message store before being expunged. (integer)
messagesize	Maximum size of message in bytes before it is marked to be expunged. (integer)
messagesizedays	Grace period. Days an over-sized message should remain in a folder. (integer)
messageheader . <i>header</i>	Specifies a header field and string. Values are not case-sensitive and regular expressions are not recognized. Example: Rule1.messageheader.Subject: Get Rich Now! Headers other than Subject: can be used.
regexp	Enable UNIX regular expressions in rules creation. (1 or 0). If not specified, IMAP expressions will be used.
savedays	Number of days the messages are saved in a folder until they are expunged.
seen	seen is a message status flag set by the system when the user opens a message. If the attribute seen is set to and , then the message must be seen and other criteria must be met before the rule is fulfilled. If the attribute seen is set to or , then the message only needs to be seen or another criteria be met before the rule is fulfilled. (and/or)
sieve	A Sieve test specifying message selection criteria. Example: Rule17.sieve: header :contains "Subject" "Vigara"
deleted	deleted is a message status flag set by the system when the user deletes a message. If the attribute deleted is set to and , then the message must be deleted and another criteria must be met before the rule is fulfilled. If the attribute deleted is set to or , then the message only needs to be deleted or another criteria be met before the rule is fulfilled. (and/or)

Localized Mailbox Names in imexpire

The IMAP protocol specifies that mailbox names use modified UTF-7 encoding. Messaging Server supports localized character sets on external interfaces so that mailbox names can be localized. Internally, however, the system converts the localized name to MUTF-7. Thus, a folder that has a localized mailbox name on a client will have a corresponding mailbox file name in MUTF-7. (Note that IMAP error messages will output mailbox names in MUTF-7 and not the localized character set.)

In general, most message store utilities that require mailbox names expect the names in the localized character set, although they may have an option flag that allows a different character set to be used. These utilities include **reconstruct**, **mboxutil**, **imsbackup**, **imsrestore**, and **hashdir**. However, **imexpire** requires that the mailbox name, specified as the attribute **folderpattern**, be in MUTF-7. Using a localized name will not work.

To obtain the appropriate **folderpattern** for **imexpire** it may be necessary to convert a localized mailbox name to the modified UTF-7 equivalent. This can be done using the **mboxutil -E** command as follows:

```
$ mboxutil -l -p user/han/*

  msgs  Kbytes last msg          partition  quotaroot mailbox
-----
    57   100 2010/04/29 11:18 primary      5242880 user/han/INBOX
    1     1 2010/04/30 12:56 primary          - user/han/<multibyte_mailbox>

$ mboxutil -l -E MUTF-7 -p user/han/*

  msgs  Kbytes last msg          partition  quotaroot mailbox
-----
    57   100 2010/04/29 11:18 primary      5242880 user/han/INBOX
    1     1 2010/04/30 12:56 primary          -
user/han/&kAFP4W4IMH8wojCkMMYw4A-
```

The first **mboxutil** shows the localized mailbox name. The second **mboxutil** shows the mailbox name in MUTF-7. The MUTF-7 mailbox name is identical to the IMAP list command:

```
x list "" *
* LIST (\NoInferiors) "/" INBOX
* LIST (\HasNoChildren) "/" &kAFP4W4IMH8wojCkMMYw4A-
```

To convert the local **charset** to modified UTF-7 encoding, use the **mboxutil** command with the **-E** option:

```
$ mboxutil -l -E MUTF-7 -P user/han/<multibyte_mailbox>

  msgs  Kbytes last msg          partition  quotaroot mailbox
-----
    1     1 2010/04/30 12:56 primary          -
user/han/&kAFP4W4IMH8wojCkMMYw4A-
```

Note that **mboxutil -E** can be used for any command that requires the use of a MUTF-7 mailbox name including **imexpire**.

Setting imexpire Rules Textually

Message expiration rules are set by specifying expire criteria in a **store.expirerule** file. The **store.expirerule** file contains one expire criteria per line. An expire criteria of the global rule configuration file (*MessagingServer_home/data/store/store.expirerule*) has the following format:

rule_name.attribute: value

"[Example imexpire Rules](#)" shows a set of global expiration rules in *MessagingServer_home/config/store.expirerule*.

Rule 1 sets the global expiration policy (that is, policy that applies to all messages), as follows:

- Enable UNIX regular expressions in rules creation.
- Removes messages larger than 100,000 bytes after 3 days.
- Removes messages deleted by the user.
- Removes any message with the strings "Viagra Now!" or "XXX Porn!" in the Subject: header.

- Limits all folders to 1,000 messages, the system removes the oldest messages on a folder to keep the total to 1,000.
- Removes all messages older than 365 days.

Rule 2 sets the message expiration policy for users at the hosted domain **example.org**. It limits mailbox sizes to 1 megabyte, removes messages that have been deleted, and removes messages older than 14 days.

Rule 3 sets the message expiration policy for messages in the inbox folder of user **f.dostoevski**. It removes messages with a subject line having the expression "On-line Casino."

Example imexpire Rules

```
Rule1.regex: 1
Rule1.folderpattern: user/. *
Rule1.messagesize: 100000
Rule1.messagesizedays: 3
Rule1.deleted: or
Rule1.Subject: Vigara Now!
Rule1.Subject: XXX Porn!
Rule1.messagecount: 1000
Rule1.messagedays: 365
Rule2.regex: 1
Rule2.folderpattern: user/. *@example.org/. *
Rule2.exclusive: 1
Rule2.deleted: or
Rule2.messagedays: 14
Rule2.messagecount: 1000
Rule3.folderpattern: user/f.dostoevski/inbox
Rule3.Subject: On-line Casino
```

Setting imexpire Folder Patterns

Folder patterns can be specified using POSIX regular expressions by setting the **imexpire** attribute **regex** to 1. If not specified, IMAP expressions will be used. The format must start with a **user/** followed by a pattern. [Table 54–2](#) shows the folder pattern for various folders.

Table 54–2 *imexpire Folder Patterns Using Regular Expressions*

Scope	Folder Patterns (regex=0)	Folder Pattern (regex=1)
Apply rule to all messages in all folders of <i>userid</i> .	user/<i>userid</i>/*	user/<i>userid</i>/.*
Apply rule to messages of <i>userid</i> in folder Sent .	user/<i>userid</i>/Sent	user/<i>userid</i>/Sent
Apply rule to entire message store.	user/*	user/.*
Apply rule to any folder called Trash anywhere in any user's hierarchy.	user/*/Trash	user/.*/Trash
Apply rule to folders in hosted domain example.org .	user/*@example.org/*	user/.*@example.org/.*
Apply rule to folders in default domain.	Not applicable	user/[^@]*/.*

Configuring Message Expiration (Tasks)

This information describes the tasks you use to expire messages. See ["Message Store Message Expiration"](#) for overview and conceptual information.

Topics:

- [To Set imexpire Rules Textually](#)
- [To Set imexpire Folder Patterns](#)
- [To Schedule Message Expiration and Logging Level](#)
- [To Exclude Specified Users from Message Expiration](#)

To Set imexpire Rules Textually

Note: In Unified Configuration, you can continue to configure messaging expiration as explained in this information, or use the **msconfig** command to individually set the appropriate configuration options. Modifying the **store.expirerule** file enables access to more functionality than using the **msconfig** command.

You expire messages by specifying rules in a **store.expirerule** file. The **store.expirerule** file contains one expire criteria per line. An expire criteria of the global rule configuration file *MessagingServer_home/config/store.expirerule* has the following format:

```
<rule_name>.<attribute>: <value>
```

An expiration rule for a user or mailbox rule configuration file has the following format:

```
<attribute>: <value>
```

The following example shows a set of global expiration rules in the *MessagingServer_home/config/store.expirerule* file.

Example imexpire Rules

```
Rule1.regex: 1
Rule1.folderpattern: user/. *
Rule1.messagesize: 100000
Rule1.messagesizedays: 3
Rule1.deleted: or
Rule1.messageheader.Subject: Vigara Now!
```

```
Rule1.messageheader.Subject: XXX Porn!  
Rule1.messagecount: 1000  
Rule1.messagedays: 365  
Rule2.regex: 1  
Rule2.folderpattern: user/.*@example.org/. *  
Rule2.exclusive: 1  
Rule2.deleted: or  
Rule2.messagedays: 14  
Rule2.messagecount: 1000  
Rule3.folderpattern: user/f.dostoevski/inbox  
Rule3.messageheader.subject: On-line Casino
```

Rule 1 sets the global expiration policy (that is, policy that applies to all messages) to:

- Enable UNIX regular expressions in rules creation.
- Remove messages larger than 100,000 bytes after 3 days.
- Remove messages deleted by the user.
- Remove any message with the strings "Vigara Now!" or "XXX Porn!" in the Subject: header.
- Limit all folders to 1,000 messages. After 1,000 messages, the system removes the oldest messages on a folder to keep the total to 1,000.
- Remove all messages older than 365 days.

Rule 2 sets the message expiration policy for users at the hosted domain **example.org**. It limits mailbox sizes to 1 megabyte, removes messages that have been deleted, and removes messages older than 14 days.

Rule 3 sets the message expiration policy for messages in the **inbox** folder of user **f.dostoevski**. It removes messages with a subject line having the expression "On-line Casino."

Note that headers other than Subject: can be used.

To Set Expiration Rules by Using the msconfig Command

- In Unified Configuration, you can set expiration rules by using the following **msconfig** command:

```
msconfig set expirerule:<name>.<option> <value>
```

For example:

```
msconfig set expirerule:Rule1.folderpattern user/. *  
./msconfig set expirerule:Rule1.messagesize 100000  
./msconfig set expirerule:Rule1.messagesizedays 3
```

You can set the following options in this way:

- **deleted**
- **exclusive**
- **folderpattern**
- **foldersizebytes**
- **messagecount**
- **messagedays**
- **messagesize**

- `messagesizedays`
- `seen`

To Set `imexpire` Folder Patterns

Folder patterns can be specified by using POSIX regular expressions by setting the `imexpire` attribute `regex` to 1. If not specified, IMAP expressions are used. The format must start with a `user/` followed by a pattern. Table 55–1 shows the folder pattern for various folders.

Table 55–1 *imexpire Folder Patterns Using Regular Expressions*

Scope	Folder Pattern (regex=0)	Folder Pattern (regex=1)
Apply rule to all messages in all folders of <code>userid</code> .	<code>user/userid/*</code>	<code>user/userid/*</code>
Apply rule to messages of <code>userid</code> in folder <code>Sent</code> .	<code>user/userid/Sent</code>	<code>user/userid/Sent</code>
Apply rule to entire message store.	<code>user/*</code>	<code>user/*</code>
Apply rule to any folder called <code>Trash</code> anywhere in any user's hierarchy.	<code>user/*/Trash</code>	<code>user/*/Trash</code>
Apply rule to folders in hosted domain <code>example.org</code> .	<code>user/*@example.org/*</code>	<code>user/*@example.org/*</code>
Apply rule to folders in default domain.	Not applicable	<code>user/[Configuring Message Expiration in Unified Configuration (Tasks)^@]/*</code>

To Schedule Message Expiration and Logging Level

You activate message expiration by using the `imsched` scheduling daemon. By default, `imsched` invokes `imexpire` at 23:00 every day. Use the `"impurge"` command for the purge function. The `imexpire` schedule can be customized by setting the `schedule.task:expire.crontab` option as described in "Expire and Purge Log and Scheduling Options."

Expire and purge can take a long time to complete on a large message store. You should experiment and decide how often to run these processes. For example, if an expire and purge cycle takes 10 hours, you might not want the default schedule of running expire and purge once a day. Schedule expire and purge by using the `imexpire` command and the automatic task scheduling option (see "Scheduling Automatic Tasks"). For example:

```
msconfig
msconfig> set schedule.task:expire.crontab "0 1 * * 6 bin/imexpire -e"
msconfig# write
msconfig> exit
```

In this example, messages are expired at 1 AM Saturdays and purged every night at 11 PM. If no purge schedule is set, `imexpire` performs purge after an expire.

Expire and Purge Log and Scheduling Options

For Purge options, see "Maintenance Queue Configuration Options." Clicking on the specific option will often provide both the legacy (`configutil`) parameter and the Unified Configuration (`msconfig`) option.

Table 55–2 Expire and Purge Log and Scheduling Options

Option	Description
<code>schedule.task:expire.enable</code>	Whether the expire task should be scheduled. Default: 1
<code>schedule.task:expire.crontab</code>	<p>Interval for running imexpire. Uses UNIX crontab format: <i>minute hour day-of-month month-of-year day-of-week</i></p> <p>The values are separated by a space or tab and can be 0-59, 0-23, 1-31, 1-12 and 0-6 (with 0=Sunday) respectively. Each time field can be either an asterisk (meaning all legal values), a list of comma-separated values, or a range of two values separated by a hyphen. Note that days can be specified by both day of the month and day of the week, however, it is not typical to use them both since the number of such occurrences are very small. If they are both specified, then both will be required. For example, setting the 17th day of the month and Tuesday will require both values to be true.</p> <p>You can also use the -e and -c flags with imexpire to and expire only or purge only respectively. See "imexpire."</p> <p>Interval Examples:</p> <ol style="list-style-type: none"> 1. Run imexpire at 12:30am, 8:30am, and 4:30pm:30 0,8,16 * * bin/imexpire 2. Run imexpire at weekday morning at 3:15 am:15 3 * * 1-5 bin/imexpire 3. Run imexpire only on Mondays:0 0 * * 1 bin/imexpire Default: 0 23 * * * bin/imexpire <p>To disable: Run msconfig set schedule.task:expire.enable 0</p>
<code>store.expire.exploglevel</code>	<p>Specify a log level:</p> <p>0 = No log.</p> <p>1 = Log summary for the entire expire session.</p> <p>2 = Log one message per mailbox expired.</p> <p>3 = Log one message per message expired.</p> <p>Default: 0</p>

To Set imexpire Logging Levels

imexpire logs a summary to the default log file upon completion. If expire is invoked from the command line, the **-v** (verbose) and **-d** (debug) options can be used to instruct **imexpire** to log detail status and debug messages to **stderr**. If **imexpire** is invoked by **imsched**, the configuration option **store.expire.exploglevel** can be set to 0, 1, 2, or 3 for different levels of logging. Loglevel 0 is the default, and no logging is performed. Loglevel 1 logs a summary for the entire expire session. Loglevel 2 logs one message per mailbox expired. Loglevel 3 logs one message per message expired.

The following example invokes expire from the command line to set verbose logging and shows the resulting messages in the default log file, *MessagingServer_home/log/default*.

```
cd /opt/sun/comms/messaging64/bin
imexpire -n -v 1
tail ../log/default
[25/Nov/2010:11:51:51 +1100] server imexpire[20730]: General Notice: imexpire
started
[25/Nov/2010:11:51:51 +1100] server imexpire[20730]: General Notice: iBiff plugin
loaded: ms-internal
```

```
[25/Nov/2010:11:51:51 +1100] server imexpire[20730]: General Notice: primary
partition: expired 0 messages
[25/Nov/2010:11:51:51 +1100] server imexpire[20730]: General Notice: Expired 0
messages
[25/Nov/2010:11:51:51 +1100] server imexpire[20730]: General Notice: Expire
finished
```

To Exclude Specified Users from Message Expiration

Exclude specified users from the expire rules by adding their user ID, one per line, in a file called **expire_exclude_list** in the *MessagingServer_home/config* directory. Or, configure a "dummy" exclusive expire rule under the user's mailbox.

Configuring POP, IMAP, and HTTP Services

Oracle Communications Messaging Server supports the Post Office Protocol 3 (POP3), the Internet Mail Access Protocol 4 (IMAP4), and the HyperText Transfer Protocol (HTTP) for client access to mailboxes. IMAP and POP are both Internet-standard mailbox protocols. Convergence, a web-enabled electronic mail program, enables end users to access their mailboxes by using a browser running on an Internet-connected computer system using HTTP.

This information describes how to configure your server to support one or more of these services by using command-line utilities.

Topics:

- [General Configuration](#)
- [Login Requirements](#)
- [Performance Options](#)
- [Client Access Controls](#)
- [To Configure POP Services](#)
- [To Configure IMAP Services](#)
- [To Configure the mshttpd Process for Use by Convergence](#)

General Configuration

Configuring the general features of the Messaging Server POP, IMAP, and HTTP services includes enabling or disabling the services, assigning port numbers, and optionally modifying service banners sent to connecting clients. This section provides background information. For the steps you follow to make these settings, see "[To Configure POP Services](#)," "[To Configure IMAP Services](#)," and "[To Configure the mshttpd Process for Use by Convergence](#)."

This section consists of the following subsections:

- [Enabling and Disabling Services](#)
- [Specifying Port Numbers](#)
- [Ports for Encrypted Communications](#)
- [Service Banner](#)

Enabling and Disabling Services

You can control whether any particular instance of Messaging Server makes its POP, IMAP, or HTTP service available for use. This is not the same as starting and stopping services (see "[Starting and Stopping Services](#)"). To function, POP, IMAP, or HTTP must be both enabled and started.

Enabling a service is a more "global" process than starting or stopping a service. For example, the Enable setting persists across system reboots, whereas you must restart a previously "stopped" service after a reboot.

There is no need to enable services that you do not plan to use. For example, if a Messaging Server instance is used only as a Mail Transfer Agent (MTA), you should disable POP, IMAP, and HTTP services. If a Messaging Server instance is used only for POP services, you should disable IMAP and HTTP. If a Messaging Server instance is used only for web-based email, you should disable both POP and IMAP.

You can enable or disable services at the server level, described later in this information. "[To Specify What Services Can Be Started](#)" also describes this process. In addition, you can enable or disable services at the user level by setting the LDAP attribute `mailAllowedServiceAccess` seen in the *Schema Reference*.

Specifying Port Numbers

For each service, you can specify the port number that the server is to use for service connections:

- If you enable the POP service, you can specify the port number that the server is to use for POP connections. The default is 110.
- If you enable the IMAP service, you can specify the port number that the server is to use for IMAP connections. The default is 143.
- If you enable the HTTP service, you can specify the port number that the server is to use for HTTP connections. The default is 8990.

You might need to specify a port number other than the default if you have, for example, two or more IMAP server instances on a single host machine, or if you are using the same host machine as both an IMAP server and a Messaging Multiplexor server. (For information about the Multiplexor, see "[Configuring and Administering Multiplexor Services](#)".)

Keep the following in mind when you specify a port:

- Port numbers can be any number from 1 to 65535.
- Make sure the port you choose isn't already in use or reserved for another service.

Ports for Encrypted Communications

Messaging Server supports encrypted communications with IMAP, POP, and HTTP clients by using the Secure Sockets Layer (SSL) protocol. For general information on support for SSL in Messaging Server, see the discussion on configuring encryption and certificate-based authentication in the *Messaging Server Security Guide*.

IMAP Over SSL

You can accept the default (recommended) IMAP over SSL port number (993) or you can specify a different port for IMAP over SSL.

Messaging Server provides the option of using separate ports for IMAP and IMAP over SSL because most current IMAP clients require separate ports for them.

Same-port communication with both IMAP and IMAP over SSL is an emerging standard. As long as your Messaging Server has an installed SSL certificate (see the discussion on obtaining certificates in the *Messaging Server Security Guide*), it can support same-port IMAP over SSL.

POP Over SSL

The default separate SSL port for POP is 995. You can also initiate SSL over normal POP port with the command "STLS" (see "[To Configure POP Services](#)").

HTTP Over SSL

You can accept the default HTTP over SSL port number (8991) or you can specify a different port for HTTPS.

Service Banner

When a client first connects to the Messaging Server POP or IMAP port, the server sends an identifying text string to the client. This service banner (not normally displayed to the client's user) identifies the server as Messaging Server, and gives the server's version number. The banner is most typically used for client debugging or problem-isolation purposes.

You can replace the default banner for the POP or IMAP service if you want a different message sent to connecting clients.

Use the `msconfig` utility and the (`pop.banner`) option to set service banners.

Login Requirements

You can control how users are permitted to log in to the POP, IMAP, or HTTP service to retrieve mail. You can allow password-based login (for all services), and certificate-based login (for IMAP or HTTP services). This section provides background information. For the steps you follow to make these settings, see "[To Configure POP Services](#)," "[To Configure IMAP Services](#)," or "[To Configure the mshttpd Process for Use by Convergence](#)." In addition, you can specify the valid login separator for POP logins. This section consists of the following subsections:

- [To Set the Separator for POP Clients](#)
- [To Allow Log In without Using the Domain Name](#)
- [Password-Based Login](#)
- [Certificate-Based Login](#)

To Set the Separator for POP Clients

Some older mail clients do not accept @ as the login separator (that is, the @ in an address like `uid@domain`). If you are using one of these older mail clients, the workaround is as follows:

1. Make + a valid separator with the following command:

```
msconfig set base.loginseparator "@+"
```

2. Inform POP client users that they should log in with + as the login separator, not @.

To Allow Log In without Using the Domain Name

A typical login involves the user entering a user ID followed by a separator and the domain name and then the password. Users in the default domain specified during installation, however, can log in without entering a domain name or separator.

To allow users of other domains to log in with just the user ID (that is, without having to use the domain name and separator) set the **auth.searchfordomain** option to **0**. The user ID must be unique to the entire directory tree. If it is not unique, logging in without the domain name will not work.

You might want to modify the attribute that user must enter to log in. For example, to allow the user to log in with a phone number (**telephoneNumber**) or employee number (**employeeID**), change the LDAP search defined by the **auth.searchfilter** option. This option is a global default setting for the **inetDomainSearchFilter** per-domain attribute and follows the same syntax.

Refer to the *Messaging Server Reference* for further information on these options.

Password-Based Login

In typical messaging installations, users access their mailboxes by entering a password into their POP, IMAP, or HTTP mail client. The client sends the password to the server, which uses it to authenticate the user. If the user is authenticated, the server decides, based on access-control rules, whether or not to grant the user access to certain mailboxes stored on that server.

If you allow password login, users can access POP, IMAP, or HTTP by entering a password. (Password- or SSL-based login is the only authentication method for POP services.) Passwords are stored in an LDAP directory. Directory policies determine what password policies, such as minimum length, are in effect.

If you disallow password login for IMAP or HTTP services, password-based authentication is not permitted. Users are then required to use certificate-based login, as described in the next section.

To increase the security of password transmission for IMAP and HTTP services, you can require that passwords be encrypted before they are sent to your server. You do this by selecting a minimum cipher-length requirement for login.

- If you choose **0**, you do not require encryption. Passwords are sent in the clear or they are encrypted, depending on client policy.
- If you choose a nonzero value, the client must establish an SSL session with the server by using a cipher whose key length is at least the value you specify, thus encrypting any IMAP or HTTP user passwords the client sends.

If the client is configured to require encryption with key lengths greater than the maximum your server supports, or if your server is configured to require encryption with key lengths greater than what the client supports, password-based login cannot occur. For information on setting up your server to support various ciphers and key lengths, see the discussion on enabling SSL and selecting ciphers in the *Messaging Server Security Guide*.

Certificate-Based Login

In addition to password-based authentication, Oracle servers support the authentication of users through examination of their digital certificates. Instead of presenting a password, the client presents the user's certificate when it establishes an

SSL session with the server. If the certificate is validated, the user is considered authenticated.

For instructions on setting up Messaging Server to accept certificate-based user login to the IMAP or HTTP service, see the discussion on setting up certificate-based login in the *Messaging Server Security Guide*.

If you have performed the tasks required to set up certificate-based login, both password-based and certificate-based login are supported. Then, if the client establishes an SSL session and supplies a certificate, certificate-based login is used. If the client does not use SSL or does not present a client certificate, it sends a password instead.

Performance Options

You can set some of the basic performance options for the POP, IMAP, and HTTP services of Messaging Server. Based on your hardware capacity and your user base, you can adjust these options for maximum efficiency of service. This section provides background information. For the steps you follow to make these settings, see "[To Configure POP Services](#)," "[To Configure IMAP Services](#)," or "[To Configure the mshttpd Process for Use by Convergence](#)."

This section consists of the following subsections:

- [Number of Processes](#)
- [Number of Connections per Process](#)
- [Number of Threads per Process](#)
- [Dropping Idle Connections](#)
- [Logging Out HTTP Clients](#)

Number of Processes

Messaging Server can divide its work among several executing processes, which in some cases can increase efficiency. This capability is especially useful with multiprocessor server machines, in which adjusting the number of server processes can allow more efficient distribution of multiple tasks among the hardware processors.

There is a performance overhead, however, in allocating tasks among multiple processes and in switching from one process to another. The advantage of having multiple processes diminishes with each new one added. A simple rule of thumb for most configurations is to have one **IMAPD** and one **POPD** process per hardware processor on your server machine, up to a maximum of perhaps four processes. Your optimum configuration might be different. This rule of thumb is meant only as a starting point for your own analysis.

Note: On some platforms you might also want to increase the number of processes to get around certain per-process limits (such as the maximum number of file descriptors), specific to that platform, that might affect performance. The default number of processes is one each for the POP, IMAP, or HTTP service.

Number of Connections per Process

The more simultaneous client connections your POP, IMAP, or HTTP service can maintain, the better it is for clients. If clients are denied service because no connections are available, they must then wait until another client disconnects.

On the other hand, each open connection consumes memory resources and makes demands on the I/O subsystem of your server machine, so there is a practical limit to the number of simultaneous sessions you can expect the server to support. (You might be able to increase that limit by increasing server memory or I/O capacity.)

IMAP, HTTP, and POP have different needs in this regard:

- IMAP connections are generally long-lived compared to POP and HTTP connections. When a user connects to IMAP to download messages, the connection is usually maintained until the user quits or the connection times out. In contrast, a POP or HTTP connection is usually closed as soon as the POP or HTTP request has been serviced.
- IMAP and HTTP connections are generally very efficient compared to POP connections. Each POP re-connection requires re-authentication of the user. In contrast, an IMAP connection requires only a single authentication because the connection remains open for the duration of the IMAP session (login to logout). An HTTP connection is short, but the user need not re-authenticate for each connection because multiple connections are allowed for each HTTP session (login to logout). POP connections, therefore, involve much greater performance overhead than IMAP or HTTP connections. Messaging Server, in particular, has been designed to require very low overhead by open but idle IMAP connections and by multiple HTTP connections.

Note: For more information about HTTP session security, see the discussion about HTTP security in the *Messaging Server Security Guide*.

Thus, at a given moment for a given user demand, Messaging Server may be able to support many more open IMAP or HTTP connections than POP connections.

The default value for IMAP is 4000. The default value for HTTP is 6000 connections per process. The default value for POP is 600. These values represent roughly equivalent demands that can be handled by a typically configured server machine. Your optimum configuration might be different. These defaults are meant only as general guidelines.

Typically, active POP connections are much more demanding on server resources and bandwidth than active IMAP connections since IMAP connections are idle most of the time while POP connections are constantly downloading messages. Having a lower number of sessions for POP is correct. Conversely, POP connections only last as long as it takes to download email, so an active POP user is only connected a small percentage of the time, while IMAP connections stay connected between successive mail checks.

Number of Threads per Process

Besides supporting multiple processes, Messaging Server further improves performance by subdividing its work among multiple threads. The server's use of threads greatly increases execution efficiency, because commands in progress are not holding up the execution of other commands. Threads are created and destroyed, as needed during execution, up to the maximum number you have set.

Having more simultaneously executing threads means that more client requests can be handled without delay, so that a greater number of clients can be serviced quickly. However, there is a performance overhead to dispatching among threads, so there is a practical limit to the number of threads the server can make use of.

For POP, IMAP, and HTTP, the default maximum value is 250 threads per process. The numbers are equal despite the fact that the default number of connections for IMAP and HTTP is greater than for POP. It is assumed that the more numerous IMAP and HTTP connections can be handled efficiently with the same maximum number of threads as the fewer, but busier, POP connections. Your optimum configuration might be different, but these defaults are high enough that it is unlikely you would ever need to increase them; the defaults should provide reasonable performance for most installations.

Dropping Idle Connections

To reclaim system resources used by connections from unresponsive clients, the IMAP4, POP3, and HTTP protocols permit the server to unilaterally drop connections that have been idle for a certain amount of time.

The respective protocol specifications require the server to keep an idle connection open for a minimum amount of time. The default times are 10 minutes for POP, 30 minutes for IMAP, 3 minutes for HTTP. You can increase the idle times beyond the default values, but you cannot make them less.

If a POP or IMAP connection is dropped, the user must re-authenticate to establish a new connection. In contrast, if an HTTP connection is dropped, the user need not re-authenticate because the HTTP session remains open. For more information about HTTP session security, see the discussion about HTTP security in the *Messaging Server Security Guide*.

Idle POP connections are usually caused by some problem (such as a crash or hang) that makes the client unresponsive. Idle IMAP connections, on the other hand, are a normal occurrence. To keep IMAP users from being disconnected unilaterally, IMAP clients typically send a command to the IMAP server at some regular interval that is less than 30 minutes.

Logging Out HTTP Clients

An HTTP session can persist across multiple connections. HTTP clients are not logged out when a connection is dropped. However, if an HTTP session remains idle for a specified time period, the server will automatically drop the HTTP session and the client is logged out (the default time period is 2 hours). When the session is dropped, the client's session ID becomes invalid and the client must re-authenticate to establish another session. For more information about HTTP security and session ID's, see the discussion about HTTP security in the *Messaging Server Security Guide*.

Client Access Controls

Messaging Server includes access-control features that enable you to determine which clients can gain access to its POP, IMAP, or HTTP messaging services (and SMTP as well). You can create flexible access filters that allow or deny access to clients based on a variety of criteria.

Client access control is an important security feature of Messaging Server. For information on creating client access-control filters and examples of their use, see the

discussion on configuring client access to POP, IMAP, and HTTP services in the *Messaging Server Security Guide*.

To Configure POP Services

You configure the Messaging Server POP service by using the **msconfig** command. This section lists the more common POP services options. The *Messaging Server Reference* provides a complete listing of options.

Note: For the POP service, password-based login is automatically enabled.

For more information, see also:

- [Enabling and Disabling Services](#)
- [To Set the Separator for POP Clients](#)
- [Specifying Port Numbers](#)
- [Number of Connections per Process](#)
- [Dropping Idle Connections](#)
- [Number of Threads per Process](#)
- [Number of Processes](#)

- To enable the POP service:

```
msconfig set pop.enable 1
```
- To disable the POP service:

```
msconfig set pop.enable 0
```
- To specify the port number:

```
msconfig set pop.port <port number>
```
- To set the maximum number of network connections per process (see "[Number of Connections per Process](#)" for details):

```
msconfig set pop.maxsessions <number>
```
- To set the maximum idle time for connections (see "[Dropping Idle Connections](#)" for details):

```
msconfig set pop.idletimeout <number>
```
- To set the maximum number of threads per process (see "[Number of Threads per Process](#)" for more information):

```
msconfig set pop.maxthreads <number>
```
- To set the maximum number of processes (see "[Number of Processes](#)" for additional information):

```
msconfig set pop.numprocesses <number>
```
- To enable POP over SSL on port 995:

```

msconfig
msconfig> set pop.enablesslport 1
msconfig# set pop.sslusessl 1
msconfig# set pop.sslport 995
msconfig# write
msconfig> exit
stop-msg pop
start-msg pop

```

TLS is also supported if SSL is configured correctly.

- To specify a protocol welcome banner:

```
msconfig set pop.banner <banner>
```

To Configure IMAP Services

You configure the Messaging Server IMAP service by using the **msconfig** command. This section lists the common IMAP services options. The *Messaging Server Reference* provides a complete listing of options. For more information, see also:

- [Enabling and Disabling Services](#)
 - [Specifying Port Numbers](#)
 - [Password-Based Login](#)
 - [Number of Connections per Process](#)
 - [Dropping Idle Connections](#)
 - [Number of Threads per Process](#)
 - [Number of Processes](#)
 - [Configuring IMAP IDLE](#)
- To enable the IMAP service:

```
msconfig set imap.enable 1
```
 - To disable the IMAP service:

```
msconfig set imap.enable 0
```
 - To specify the port number:

```
msconfig set imap.port <number>
```
 - To enable a separate port for IMAP over SSL:

```
msconfig set imap.enablesslport 1
```
 - To specify a port number for IMAP over SSL:

```
msconfig set imap.sslport <number>
```
 - To enable or disable password login to the IMAP service:

```
msconfig set imap.plaintextmncipher <value>
```

If *value* is greater than 0, disable use of plaintext passwords unless a security layer (SSL or TLS) is activated. This forces users to enable SSL or TLS on their client to log in, which prevents exposure of their passwords on the network. Default is 0.

- To set the maximum number of network connections per process (see "[Number of Connections per Process](#)" for additional information):

```
msconfig set imap.maxsessions <number>
```

- To set the maximum idle time for connections (see "[Dropping Idle Connections](#)" for additional information):

```
msconfig set imap.idletimeout <number>
```

- To set the maximum number of threads per process (see "[Number of Threads per Process](#)"):

```
msconfig set imap.maxthreads <number>
```

- To set the maximum number of processes (see "[Number of Processes](#)"):

```
msconfig set imap.numprocesses <number>
```

- To specify a protocol welcome banner:

```
msconfig set imap.banner <banner>
```

- To enable IMAP over SSL on port 993:

```
msconfig
msconfig> set imap.enablesslport 1
msconfig# set imap.sslusessl 1
msconfig# set imap.sslport 993
msconfig# write
msconfig> exit
stop-msg imap
start-msg imap
```

Configuring IMAP IDLE

The IMAP IDLE extension to the IMAP specification, defined in RFC 2177, enables an IMAP server to notify the mail client when new messages arrive and other updates take place in a user's mailbox. See "[Configuring IMAP IDLE](#)" for conceptual and task information on enabling IMAP IDLE in Messaging Server.

To Configure the mshttpd Process for Use by Convergence

Messaging Server supports the mail client Convergence.

While POP and IMAP clients send mail directly to a Messaging Server MTA for routing or delivery, HTTP clients send mail to a specialized web server called the Webmail Server (also called **mshttpd** or Messaging Server HTTP daemon). Depending on where the message is addressed, the Webmail Server directs the mail to an outbound MTA for routing or to one of the back-end message stores using IMAP. Convergence simply routes requests to and from the Webmail Server.

The Webmail Server accesses the message store through the IMAP server. This provides several advantages:

- Convergence clients are able to access shared folders that are located on different back-end message stores.
- The Webmail Server does not need to be installed on each back-end server.
- The Webmail Server can serve as a front-end server performing multiplexing capabilities.
- Users can access shared folders that are not on their message store.

The Webmail Server operates as a front-end server receiving HTTP client email requests. It translates these requests to SMTP or IMAP calls and forwards the calls to either the MTA or the appropriate IMAP server on the back-end message store. If Messaging Server is used only for web-based email, make sure that IMAP is enabled.

Configuring Your HTTP Service

Many of the HTTP configuration options are similar to the options available for the POP and IMAP services, including options for connection settings and process settings. This section lists common HTTP service options. The *Messaging Server Reference* provides a complete listing of options. For more information, see also:

- [Enabling and Disabling Services](#)
- [Specifying Port Numbers](#)
- [Password-Based Login](#)
- [Number of Connections per Process](#)
- [Dropping Idle Connections](#)
- [Logging Out HTTP Clients](#)
- [Number of Threads per Process](#)
- [Number of Processes](#)

For each IMAP server that users access, the Webmail Server needs to know the IMAP port, whether to use SSL, and the administrative credentials for user log-in. The configuration options to do this are as follows:

- **base.proxyimapport**: IMAP port on which to connect (default 143).
- **base.proxyimapssl**: Enable SSL (default no).
- **base.proxyadmin**: Specifies the store Admin ID.
- **base.proxyadminpass**: Specifies the store Admin password.

You can set these options globally in Unified Configuration to apply to every IMAP back-end server by using **base.proxyadmin**. Alternatively, you can set these options for each individual IMAP back-end server by using **proxy:storeaffinitygroup.imapadmin**.

To use IMAP over SSL, you must configure **mshttpd** as an SSL HTTP server, and the **mshttpd** certificate database must trust the IMAP back end's CA. You must enable **http.sslusessl**. If the back-end message store running IMAP is using a self-signed certificate (for example, as created by **generate-certDB**), then this certificate needs to be added to the front-end **mshttpd** daemon server.

If **base.proxyadmin** and **base.proxyadminpass** are not configured, logins are rejected. The system provides the error message, "Mail server unavailable. Administrator, check server log for details" and the HTTP log lists the missing configuration options.

Additional values for HTTP attributes can be set at the command line as follows:

- To enable the HTTP service:

```
msconfig set http.enable 1
```

- To disable the HTTP service:

```
msconfig set http.enable 0
```

By default, the HTTP service sends outgoing web mail to the local MTA for routing or delivery. You might want to configure the HTTP service to send mail to a remote MTA, for example, if your site is a hosting service and most recipients are not in the same domain as the local host machine. To send web mail to a remote MTA, you need to specify the remote host name and the SMTP port number for the remote host.

- To specify the port number:

```
msconfig set http.port <number>
```

- To enable a separate port for HTTP over SSL:

```
msconfig set http.enablesslport 1
```

- To specify a port number for HTTP over SSL:

```
msconfig set http.sslport <number>
```

- To enable or disable password login:

```
msconfig set http.plaintextmncipher <value>
```

If *value* is greater than 0, then disable use of plaintext passwords unless a security layer (SSL or TLS) is activated. This forces users to enable SSL or TLS on their client to log in, which prevents exposure of their passwords on the network. Default is 0.

- To set the maximum number of network connections per process (for more information, see "[Number of Connections per Process](#)):

```
msconfig set http.maxsessions <number>
```

For more information, see "[Dropping Idle Connections](#)."

- To set the maximum idle time for client sessions (for more information, see "[Logging Out HTTP Clients](#)):

```
msconfig set http.sessiontimeout <number>
```

- To set the maximum number of threads per process:

```
msconfig set http.maxthreads <number>
```

- To set the maximum number of processes:

```
msconfig set http.numprocesses <number>
```

When an HTTP client constructs a message with attachments, the attachments are uploaded to the server and stored in a file. The HTTP service retrieves the attachments and constructs the message before sending the message to an MTA for routing or delivery. You can accept the default attachment spool directory or specify an alternate directory. You can also specify a maximum size allowed for attachments. To specify the attachment spool directory for client outgoing mail use the following command. This includes all the attachments encoded in base64, and that base64 encoding requires an

extra 33 percent more space. Thus, a 5 Mb limit in the option results in the maximum size of one message and attachments being about 3.75 Mb.

- To set the spool directory:

```
msconfig set http.spooldir <dirpath>
```

- To specify the maximum message size:

```
msconfig set http.maxmessagesize <size>
```

where *size* is a number in bytes. This includes all the attachments encoded in base64, and that base64 encoding requires an extra 33 percent more space. Thus, a 5 Mb limit in the option results in the maximum size of one message and attachments being about 3.75 Mbs.

- To specify an alternate MTA host name:

```
msconfig set http.smtphost <hostname>
```

- To specify the port number for the alternate MTA host name:

```
msconfig set http.smtpport <portnum>
```

To enable HTTP access over SSL on port 8991:

```
msconfig
msconfig> set http.enablesslport -1
msconfig# set http.sslusessl 1
msconfig# set http.sslport 8991
msconfig# write
msconfig> exit
stop-msg http
start-msg http
```

Handling Message Store Overload

This chapter contains the following topics:

- [Overview of Managing Message Store Load](#)
- [Message Store Load Throttling](#)
- [Job Controller Stress Handling](#)
- [Default Job Controller Configuration](#)

Overview of Managing Message Store Load

An overloaded message store can suffer from degraded performance. The **mboxlist** database is particularly sensitive to overload conditions. When the database detects deadlocks, all database operations that cannot acquire the locks they need must abort the transactions and retry, thereby decreasing the throughput. If this situation continues, the message store can become very inefficient. In extreme cases, you need to restart the message store to recover.

Therefore, having the ability to control the message store load is crucial to prevent performance degradation. The message store uses transaction checkpoint time as the stress indicator. The **stored** daemon measures the transaction checkpoint duration (the time it takes to sync the database pages from the memory pool to disks). When the transaction checkpoint exceeds one minute, it raises an alarm.

Message Store Load Throttling

Message store throttling is used to regulate short spikes of activities. When the **ims_master** program detects the stressed status from the message store, it informs the Job Controller. The Job Controller responds by temporarily decreasing the number of **ims_master** processes for the **ims-ms** channel. Similarly, when the LMTP server detects the stressed status, it tells the LMTP client, which informs the Job Controller, to back off. By decreasing the number of delivery threads, the Job Controller enables the message store to recover before performance begins to degrade.

Job Controller Stress Handling

Channel programs can now tell the Job Controller if they are being overwhelmed. If this occurs, then the job controller sees if it has happened recently. The job controller ignores stressed channel messages that are received within **job_controller.stressblackout** seconds of a previous stressed message for the same channel. If the message is processed, then the job controller multiplies the effective **threaddepth** option for the channel by **job_controller.stressfactor**, and subtracts **job_**

controller.stressjob from the job limit for the channel. **threaddepth** never goes over 134,217,727, and job limit never goes below 1. In addition, then the Job Controller asks all current master programs for the channel to exit, and, if the queue is not empty, starts an appropriate number of processes.

When **job_controller.stresstime** seconds has passed after the last stress change, the Job Controller divides **threaddepth** by **job_controller.unstressfactor** (never allowing thread depth to drop below the original configured **threaddepth**), and adds **UnstressJob** to the job limit (never allowing the job limit to rise above the original configured limit. a "stress change" is either an increase in stress or a decrease in stress.

Default Job Controller Configuration

These configuration options have the following default values:

job_controller.stressblackout=60

job_controller.stresstime=120

job_controller.stressfactor=5

job_controller.stressjobs=2

job_controller.unstressfactor=stressfactor

job_controller.unstressjobs=stressjobs

Managing Message Store Partitions and Adding Storage

This information describes message store partitions and adding storage.

Topics:

- [Message Store Partition Overview](#)
- [To Add a Message Store Partition](#)
- [To Change the Default Message Store Partition](#)
- [Adding More Physical Disks to the Message Store](#)

See also:

- [To Move Mailboxes to a Different Disk Partition](#)

Message Store Partition Overview

Mailboxes are stored in message store partitions, an area on a disk partition specifically devoted to storing the message store. Message store partitions are not the same as disk partitions, though for ease of maintenance, it is recommended that you have one disk partition and one file system for each message store partition. Message store partitions are directories specifically designated as a message store.

User mailboxes are stored by default in the `store_root/partition/` directory (see "Message Store Directory Layout"). The `partition` directory is a logical directory that might contain a single partition or multiple partitions. At start-up time, the `partition` directory contains one subpartition called the **primary** partition.

You can add partitions to the `partition` directory as necessary. For example, you might want to partition a single disk to organize your users as follows:

```
store_root/partition/mkting/  
store_root/partition/eng/  
store_root/partition/sales/
```

As disk storage requirements increase, you might want to map these partitions to different physical disk drives.

You should limit the number of mailboxes on any one disk. Distributing mailboxes across disks improves message delivery time (although it does not necessarily change the SMTP accept rate). The number of mailboxes you allocate per disk depends on the disk capacity and the amount of disk space allocated to each user. For example, you can allocate more mailboxes per disk if you allocate less disk space per user.

If your message store requires multiple disks, you can use RAID (Redundant Array of Inexpensive Disks) technology to ease management of multiple disks. With RAID technology, you can spread data across a series of disks but the disks appear as one logical volume so disk management is simplified. You might also want to use RAID technology for redundancy purposes; that is, to duplicate the store for failure recovery purposes.

Note: To improve disk access, the message store and the message queue should reside on separate disks.

To Add a Message Store Partition

When adding a partition, you specify both an absolute physical path where the partition is stored on disk, and a logical name (called the partition nickname).

The partition nickname enables you to map users to a logical partition name regardless of the physical path. When setting up user accounts and specifying the message store for a user, you can use the partition nickname. The name you enter must be an alphanumeric name and must use lowercase letters.

To create and manage the partition, the user ID used to run the server must have permission to write to the location specified in the physical path.

Note: After adding a partition, you must stop then restart Oracle Communications Messaging Server to refresh the configuration information.

- **Command Line,** To add a partition to the store at the command line: The location of message files is controlled by setting the **partition:partition_name.messagepath** option, where *partition_name* is the logical name of the partition.

```
msconfig set partition:<partition_name>.messagepath <path>
```

path indicates the absolute path name where the partition is stored. To specify the path to the primary partition:

```
msconfig set partition:primary.path <path>
```

To Change the Default Message Store Partition

The default partition is the partition used when a user is created and the **mailMessageStore** LDAP attribute is not specified in the user entry. The **mailMessageStore** LDAP attribute, which specifies a user's message store partition, should be specified in all user entries so that a default partition is not necessary. In addition, the default partition should **not** be changed for load balancing or any other reason. It is invalid and dangerous to change the default partition while there are still users depending on the default partition definition.

If it is absolutely necessary to change the default partition, make sure that all users on the old default partition (the one being left behind) have their **mailMessageStore** attribute set to their current partition (which will no longer be the default), before changing the definition of default with the **store.defaultpartition** option.

Adding More Physical Disks to the Message Store

The Messaging Server message store contains the user mailboxes for a particular Messaging Server instance. The size of the message store increases as the number of mailboxes, folders, and log files increase.

As you add more users to your system, your disk storage requirements increase. Depending on the number of users your server supports, the message store might require one physical disk or multiple physical disks. Messaging Server enables you to add more stores as needed.

Managing Message Store Quotas

This information describes managing Message Store quotas and the quota tasks.

Topics:

- [Message Store Quota Overview](#)
- [To Specify a Default User Quota](#)
- [To Specify Individual User Quotas](#)
- [To Specify Domain Quotas](#)
- [To Set Up Quota Notification](#)
- [To Enable or Disable Quota Enforcement](#)
- [To Set a Grace Period](#)
- [Netscape Messaging Server Quota Compatibility Mode](#)

See also "[Monitoring the Message Store.](#)"

Message Store Quota Overview

This chapter describes quota concepts. See "[Managing Message Store Quotas](#)" for information on how to use quotas in your system.

Message store quotas limit or reduce message store usage. They enable you to set *quotas* for how much disk space or how many messages can be used by a user or domain.

Topics:

- [Quota Overview](#)
- [Quota Theory of Operations](#)
- [Message Store Quota Attributes and Options](#)

Quota Overview

Quotas can be set, in terms of number of messages or number of bytes or both, for specific users or domains. Quotas can also be set for specific folders and message types. For example, you can set different quotas based on whether a message is a voice mail or an email. Folder quotas set limits to the size of a user's folder in bytes or number of messages. For example, a quota can be set on the Trash folder. Oracle Communications Messaging Server enables you to set default quotas for domains and users as well as customized quotas.

You can also configure how the system responds to users or domains that are either over quota or approaching the quota. One response is to send users an *over quota notification*. Another response is to halt delivery of messages into the message store when quota is exceeded. This is called *quota enforcement* and usually occurs after a specified *grace period*. A grace period is how long the mailbox can be over the quota before enforcement occurs. If message delivery is halted due to over quota, incoming messages can either remain in the MTA queue until one of the following occurs or be rejected by the MTA immediately, if **local.store.overquotastatus** is enabled:

- The size or number of the user's messages no longer exceeds the quota, at which time the MTA delivers the messages.
- The undelivered message remains in the MTA queue longer than the specified *grace period*, at which time messages are returned to sender. (See ["To Set a Grace Period."](#))
- The message has remained in the message queue longer than the maximum message queue time. This is controlled by the **notices** MTA channel keyword (see the discussion on setting notification message delivery intervals in the *Messaging Server Reference*).

For example, if your grace period is set for two days, and you exceed quota for one day, new messages continue to be received and held in the message queue, and delivery attempts continue. After the second day, the messages bounce back to the sender.

Disk space becomes available when a user deletes and expunges messages or when the server deletes messages according to expiration policies established (see ["Message Store Message Expiration"](#)).

Exceptions for Telephony Application Servers

To support unified messaging requirements, Messaging Server provides the ability to override quota limitations imposed by the message store. This guarantees the delivery of messages that have been accepted by certain agents, namely telephony application servers (TAS). Messages accepted by a TAS can be routed through a special MTA channel that ensures the message is delivered to the store regardless of quota limits. This is a fairly esoteric usage, but can be useful to telephony applications. For more information about configuring a TAS channel, contact your Oracle messaging representative.

Quota by message type is useful for telephony applications that use unified messaging. For example, if a mix of messages, say text and voice mail, is stored in a user's mailbox, then the administrator can set different quotas for different types of messages. One quota can be set for email and another can be set for voice mail.

Quota Theory of Operations

Customized user and domain quotas are specified by adding quota attributes to LDAP user and domain entries. Quota defaults, notification policy, enforcement, and grace period are specified in **msconfig** options or by using the **"imquotacheck"** command.

To determine if a user is over quota, Messaging Server first determines if a quota has been set for the individual user. If no quota has been set, Messaging Server looks at the default quota set for all users. For a user, the quota is for all the cumulative bytes or messages in all of the user's folders. For a domain, the quota is for all the cumulative bytes or messages of all the users in a particular domain. For a message type, the quota is for all the cumulative bytes or messages for that message type. For a folder, the quota is for all the cumulative bytes or messages for user's folder.

You can specify the following quota values for a user's mailbox tree:

- Quota values for specific folders in the user's mailbox.
- Quota values for specific message types such as voice mail or text messages. (A message type quota applies to messages of that type in all folders in the user's mailbox.)
- A default quota value that applies to all folders and message types in the user's mailbox that are not explicitly assigned quotas.

The following guidelines apply when you assign multiple quota values for a user:

- Quotas do not overlap. For example, when there is a quota for a particular message type or folder, messages of that type or messages in that folder are not counted toward the default quota. Each message counts toward one and only one quota.
- The total quota for the whole user mailbox equals the sum of the values of all the quotas specified by default, type, and folder.
- Message type quotas take precedence over folder quotas. For example, suppose one quota is specified for a user's **memos** folder and another quota is specified for voice messages. Now suppose the user stores eight voice messages in the **memos** folder. The eight messages are counted toward the voice mail quota and excluded from the **memos** folder quota.

Changes made to the quota attributes and **msconfig** options will take effect automatically, but not immediately as information is stored in caches and it may take a little time before the changes fully take effect. Messaging Server provides a command, **"iminitquota"** that updates the changes immediately.

The **"imquotacheck"** utility enables you to check message store usage against assigned quotas.

Methods of Notification

If **store.quotanotification** is enabled, when users approach or exceed their quota limit (depending on **store.quotawarn**), the message defined by **store.quotaexceededmsg** notifies them immediately. Otherwise, you must run **imquotacheck -n** to notify the users. Dynamic user notification by enabling **store.quotanotification** and running **imquotacheck -n** are mutually exclusive. If **store.quotanotification** is enabled, you should not use **imquotacheck -n**. The preferred method is dynamic notification.

Domain level quota enforcement and reporting is done by running **imquotacheck -f**.

For more information on **imquotacheck** options, see [Table 75–15, "imquotacheck Options"](#).

Message Store Quota Attributes and Options

This section lists the major the message store quota attributes and **msconfig** options. The intention is to provide you with an overview of the functionality interface. For detailed information on these attributes and options, refer to the appropriate reference documentation.

[Table 59–1](#) lists the quota attributes. Refer to the *Schema Reference* for more information about the following attributes.

Table 59–1 Message Store Quota Attributes

Attribute	Description
mailQuota	Bytes of disk space allowed for the user's mailbox.
mailMsgQuota	Maximum number of messages permitted for a user. This is a cumulative count for all folders in the store.
mailUserStatus	Status of the mail user. Some of the possible values are active , inactive , deleted , hold , and overquota .
mailDomainDiskQuota	Bytes of disk space allowed for the cumulative count of all the mailboxes in a domain.
mailDomainMsgQuota	Maximum number of messages permitted for a domain, that is, the total count for all mailboxes in the store.
mailDomainStatus	Status of the mail domain. Values and default are the same as mailUserStatus .

To have the preceding attributes take effect, run "iminitquota" to make quota and usage up-to-date. The changes would take effect without running this, but not immediately, as information is stored in caches and it takes a little time before the changes take effect.

Table 59–2 lists the quota options. Refer to "Overview of Messaging Server Unified Configuration" for the latest and most detailed information.

Table 59–2 Message Store msconfig Options

Option	Description
store.quotaenforcement	Enable quota enforcement. When off, the quota database is still updated, but messages are always delivered. Default: 1 (bool).
store.quotanotification	Enable quota notification. Default: 0 (bool).
store.defaultmailboxquota	Store default quota by number of bytes. Default: "-1" (string) (unlimited).
store.defaultmessagequota	Store default quota by number of messages. Numeric. Default: "-1" (string) (unlimited).
store.quotaexceededmsg	Message to be sent to user when quota exceeds store.quotawarn . If none, notification is not sent. Default: No default (non-empty string). The message must contain a header (with at least a subject line), followed by \$\$, then the message body. The \$ represents a new line. There is support for the following variables: [ID] - userid, [DISKUSAGE] - disk usage, [NUMMSG] - number of messages, [PERCENT] - store.quotawarn percentage, [QUOTA] - mailquota attribute, [MSGQUOTA] - mailmsgquota attribute.
store.quotaexceededmsginterval	Interval, in days, for sending overquota notification. Default: 7 (int32).
store.quotagraceperiod	Time, in hours, a mailbox has been overquota before messages to the mailbox will bounce back to the sender. Number of hours. Default: 120 (uint32).
store.quotawarn	Quota warning threshold. Percentage of quota exceeded before clients are sent an over quota warning. Default: 90 (int32).

Table 59–2 (Cont.) Message Store msconfig Options

Option	Description
<code>store.quotaoverdraft</code>	Used to provide compatibility with systems that migrated from the Netscape Messaging Server. When ON, allow delivery of one message that puts disk usage over quota. After the user is over quota, messages are deferred or bounced, the quota warning message is sent, and the quota grace period timer starts. (The default is that the quota warning messages are sent when the message store reaches the threshold.) Default: 0 (bool), but is treated as on if <code>store.overquotastatus</code> is set, otherwise the user can never go over quota and the <code>overquotastatus</code> is never used.
<code>store.overquotastatus</code>	Enable quota enforcement before messages are enqueued in the MTA. This prevents the MTA queues from filling up. When set, and a user is not yet over quota, but an incoming message pushes the user over quota, then the message is delivered, but the <code>mailuserstatus</code> LDAP attribute is set to overquota so no more messages are accepted by the MTA. Default: 0 (bool).

To have the preceding `msconfig` options take effect, restart Messaging Server.

The `"imquotacheck"` utility enables you to check message store usage against assigned quotas.

Also see `"iminitquota"` to update the information in the quota database if a user's quota-related LDAP attributes (or the system defaults) have been changed recently and the changes have not yet been propagated automatically to the quota database.

To Specify a Default User Quota

A default quota applies to users who do not have individual quotas set in their LDAP entries. The process consists of two steps:

1. Specifying a user default quota
2. Specifying which users are bound to the default quota

The following examples show how to set default user quotas. Refer to *Messaging Server Reference* for detailed option information.

- To specify a default user disk quota for message size in bytes:

```
msconfig set store.defaultmailboxquota [ -1 | <number> ]
```

where `-1` indicates no quota (unlimited message usage) and `number` indicates a number of bytes.

- To specify a default user quota for total number of messages:

```
msconfig set store.defaultmessagequota [ -1 | <number> ]
```

where `-1` indicates no quota (unlimited messages) and `number` indicates the number of messages.

- To specify the default quota for specific users: Set the `mailQuota` attribute to `-2` in the user entries that use the default message store quota. Note that if `mailQuota` is not specified, the system default quota is used.

To Specify Individual User Quotas

Each user can have individualized quotas. To set user-specific quotas, set the **mailQuota** or **mailMsgQuota** attributes in the user's LDAP entry. For more information on the previous attributes see the *Schema Reference*.

In addition, a number of configuration options of the form **store.quota*** can be used to implement more finely grained quota policies. See the *Messaging Server Reference*. The following examples show how to set user quotas.

- To specify the system default quota, do not add **mailQuota** to the LDAP entry, or set it to **-2**.
- To set the quota to 1,000 messages, set **mailMsgQuota** to **1000**.
- To set the quota to two megabytes, set **mailQuota** to **2M**.
- To set the quota to two gigabytes, set **mailQuota** to **2G** or **2000M**.

The following LDAP entry specifies a 2 Gigabyte quota; a 20 Megabyte voice mail quota; and a 100 Megabyte quota for the Archive folder:

```
mailQuota: 2G;#voice%20M;Archive%100M
```

The two gigabyte quota represents all folders in the user's mailbox that are not explicitly assigned quotas. In this example, that excludes messages in the **Archive** folder, and messages of type **voice**. The 100 megabyte quota includes messages in any folders within the **Archive** folder.

Refer to the *Schema Reference* for more information about **mailQuota** values. Also see "[Managing Message Types in the Message Store](#)."

To Specify Domain Quotas

You can set disk space or message quotas for domains. These quotas are for the cumulative bytes or messages of all users in a particular domain. To specify domain quotas, set the **mailDomainDiskQuota** or **mailDomainMsgQuota** attributes in the desired LDAP domain entry. For more information about the previous attributes see the *Schema Reference*.

- To set the quota to 1,000 messages, set **mailDomainMsgQuota** to **1000**.
- To set the quota to two megabytes, set **mailDomainDiskQuota** to **2M** or **2000000**.
- To set the quota to two gigabytes, set **mailDomainDiskQuota** to **2G** or **2000000000** or **2000M**.

To Set Up Quota Notification

Quota notification is the process of sending users a warning message when they are getting close to their quota.

1. To enable quota notification, run the following command:

```
msconfig set store.quotanotification 1
```

If the message is not set, no quota warning message is sent to the user.

2. To define a quota warning message, run the following command:

```
msconfig set store.quotaexceededmsg '<message>'
```

For example:

```
msconfig set store.quotaexceededmsg 'Subject: WARNING: User quota
exceeded$$User quota threshold exceeded - reduce space used.'
```

The Warning Message is the message that is sent to users who are close to exceeding their disk quota. The warning message must:

- Be in RFC 822 format
- Contain a header with at least a subject line, followed by \$\$, then the message body
- Use "\$" to represent a new line. Depending on the shell that you are using, it might be necessary to append a \ before \$ to escape the special meaning of \$. (\$ is often the escape character for the shell.) You can also use the following variables in the message:
 - [ID] - userid
 - [DISKUSAGE] - disk usage
 - [NUMMSG] - number of messages
 - [PERCENT] - `store.quotawarn` percentage
 - [QUOTA] - `mailquota` attribute
 - [MSGQUOTA] - `mailmsgquota` attribute. The following example shows a warning message that uses these variables:

```
msconfig set store.quotaexceededmsg 'Subject: Overquota
Warning$$[ID],$$Your mailbox size has exceeded [PERCENT] of its allotted
quota.$Disk Usage: [DISKUSAGE]$Number of Messages: [NUMMSG]$Mailquota:
[QUOTA]$Message Quota: [MSGQUOTA]$$-Postmaster'
```

3. To specify how often the warning message is sent, run the following command:

```
msconfig set store.quotaexceededmsginterval <number>
```

where **number** indicates a number of days. For example, 3 would mean the message is sent every 3 days.

4. To specify a quota threshold, run the following command:

```
msconfig set store.quotawarn <number>
```

where **number** indicates a percentage of the allowed quota. A quota threshold is a percentage of a quota that is exceeded before clients are sent a warning. When a user's disk usage exceeds the specified threshold, the server sends a warning message to the user.

Note: When the `store.quotaoverdraft` is enabled, email notifications are not triggered until the user's disk usage exceeds 100 percent of the quota, regardless of the threshold set with `store.quotawarn`.

For IMAP users whose clients support the IMAP ALERT mechanism, the message is displayed on the user's screen each time the user selects a mailbox and a message is also written to the IMAP log.

To Disable Quota Notification

- To disable quota notification, run the following command:

```
msconfig set store.quotanotification 0
```

To Enable or Disable Quota Enforcement

By default, users or domains can exceed their quotas with no effect except for receiving an over quota notification (if set). Quota enforcement locks the mailboxes from receiving further messages until the disk usage is reduced below the quota level.

To Enable Quota Enforcement at the User level

- To enable quota enforcement at the user level, run the following command:

```
msconfig set store.quotaenforcement 1
```

The MTA saves over-quota messages in its queues and notifies users that their messages were not delivered but that a redelivery attempt is to be made later. Delivery retries continue until either the grace period expires and all messages are sent back to the senders, or the disk usage falls below the quota and messages can be dequeued from the MTA and delivered to the message store. If you want to return messages that are over quota before they get to the message queues, use the following command:

```
msconfig set store.overquotastatus 1
```

To Perform Quota Enforcement at the Domain Level

Unlike user-level quotas, domain-level quotas are not maintained dynamically.

- To enforce quotas for a particular domain, use the following command:

```
imquotacheck -f -d <domain>
```

To enforce quotas for all domains, exclude the **-d** option.

When **imquotacheck -f** finds a domain with **maildomainstatus=active** that has exceeded its quota, the **maildomainstatus** attribute is set to **overquota**, which halts all delivery to this domain. When **imquotacheck -f** is run again and the domain is back under quota, the value is set to **active**.

Disabling Quota Enforcement

If it appears that user quotas are being enforced, even when you have disabled them, check that the following options are disabled or not set:

- **store.quotaenforcement**
- **store.overquotastatus**
- **store.quotaoverdraft**

When **store.overquotastatus** is enabled (set to 1), it always treats **store.quotaoverdraft** as enabled, otherwise users never go over quota to trigger the rejection. Also, when **store.quotaoverdraft** is enabled, users are allowed one message that is smaller than the quota only. That is, it never accepts a message that is greater than the user's quota.

After changing these options, restart Messaging Server.

These Message Store attributes should be active:

- **maildomainstatus**

- **mailuserstatus**

Messages bounce if they are larger than the mailbox quota, regardless of quota enforcement configuration.

To Set a Grace Period

The grace period specifies how long the mailbox can be over the quota (disk space or number of messages) before messages are bounced back to sender. The grace period is not how long the message is held in the message queue, it's how long the mailbox is over quota before all incoming messages, including those in the message queue, are bounced. (see "[Message Store Quota Overview](#)" for more details.) The grace period starts when the user has reached the quota threshold and been warned. See "[To Set Up Quota Notification](#)."

- To specify a quota grace period at the command line:

```
msconfig set store.quotagraceperiod <number>
```

where **number** indicates number of hours.

Netscape Messaging Server Quota Compatibility Mode

After disk usage exceeded the quota in the Netscape Messaging Server, the server deferred or bounced message delivery, sent an over quota notification, and started the grace period. Messaging Server provides an option, **store.quotaoverdraft**, which retains this behavior.

When **store.quotaoverdraft** is enabled (set to 1), messages are delivered until disk usage is over quota. At that time, messages are deferred (messages stay in the MTA message queue but are not delivered to the message store), an over quota warning message is sent to the user, and a grace period starts. The grace period determines how long a mailbox is overquota before the overquota messages bounce. (The default is that the quota warning messages are sent when the message store reaches the threshold.) The default for this option is disabled (set to 0).

Managing Message Types in the Message Store

This information describes how to work with message types. For conceptual information, see "[Message Store Message Type Overview](#)."

Topics:

- [To Configure Message Types](#)
- [Sending Notification Messages for Message Types](#)
- [Administering Quotas by Message Type](#)
- [Expiring Messages by Message Type](#)

To Configure Message Types

To configure a message type, use the **msconfig** utility to set the **store.messagestype** option values that define and identify the message type.

1. Enable message types by setting the **store.messagestype.enable** option to 1. This option enables the message store to identify and manipulate message types. You must set this option before you can configure an individual message type. For example, type the following command:

```
msconfig set store.messagestype.enable 1
```

2. Define and identify the message type by setting **store.messagestype.mtindexx** options. The variable *x* identifies this particular message type in the message store. The variable *x* must be an integer greater than zero and less than 64. You can define up to 63 message types by iteratively configuring this option with unique integers. You define the value of the message type with a text string that describes the type.

- For example, to define a text message type, type the following command:

```
msconfig set store.messagestype.mtindex:1.contenttype text/plain
```

- To define a voice message type, type the following command:

```
msconfig set store.messagestype.mtindex:2.contenttype  
multipart/voice-message
```

3. Provide a flag name for the message type by setting the **store.messagestype.mtindex:x.flagname** option. This option creates a unique flag that identifies the message type. The flag is automatically set whenever a message of this type first arrives in the message store and remains associated with the message until it is purged. The flag name value is a text string that describes the message type. It does not have to be the same as the value set with the

store.message.type.mtindex:x option. The variable *x* is the integer ID of the message type defined with the **store.message.type.mtindex:x** option. For example, to define flag names for the message types configured in the preceding step, type the following commands:

```
msconfig set store.message.type.mtindex:1.flagname text
msconfig set store.message.type.mtindex:2.flagname voice_message
```

4. Configure a quota root name for the message type by setting the **store.message.type.mtindex:x.quotaroot** option. This option enables the quota function to identify and manage a quota root for this message type. The option value is a name-a text string that describes the message type. It does not have to be the same as the value set with the **store.message.type.x** option. The variable *x* is the integer ID of the message type defined with the **store.message.type.x** option. When this option is configured, you can set a quota that applies to the specified message type. For more information, see ["Administering Quotas by Message Type."](#) For example, to enable the use of quota roots for the message types configured in the preceding steps, type the following commands:

```
msconfig set store.message.type.mtindex:1.quotaroot text
msconfig set store.message.type.mtindex:2.quotaroot voice
```

5. To configure an alternate header field for identifying the message type, set the **store.message.type.header** option. By default, the message store reads the Content-Type header field to determine the message type. Configure the **store.message.type.header** option only if you want to use a different header field for identifying the message type. The value of this option is a text string. For example, to use a field called **X-Message-Type**, enter the following command:

```
msconfig set store.message.type.header X-Message-Type
```

Sending Notification Messages for Message Types

Notifications can deliver status information about messages of different types, such as text messages, voice mail, and image data. Oracle Communications Messaging Server uses Message Queue to send notification information for message types. For information about configuring the JMQ notification plug-in for Message Queue, see ["Configuring a JMQ Notification Service \(Tasks and Examples\)."](#)

To enable the JMQ notification plug-in to recognize a particular message type, you must configure the **store.message.type** options, including the **store.message.type.mtindex:x.flagname** option. For details, see ["To Configure Message Types."](#)

Once the message types have been configured, JMQ notification messages can identify the particular message types. You can write a Message Queue client to interpret notification messages by message type and deliver status information about each type to the mail client.

The JMQ notification function counts the number of messages currently in the mailbox, by message type. Instead of sending one count, an array specifying the count for each message type is sent with the notification message. For example, a **NewMsg** notification message can carry data to users informing them that their inbox has new voice mail and text messages.

For more information about sending notifications by message type, see ["Notifications for Particular Message Types."](#)

Administering Quotas by Message Type

When you set a quota for a message type, you include that value in a *quota root*. A quota root specifies quotas for a user. It can specify different quotas for particular message types and mailbox folders, and it can specify a default quota that applies to all remaining message types, folders, and messages not defined by type.

For complete information about setting and managing quotas, see "[Quota Theory of Operations](#)."

Before You Set Message-Type Quotas

Before you can set quotas for message types, you must configure the following options:

- Set the **store.message_type.mtindex:x.quotaroot** option for each message type. For details, see "[To Configure Message Types](#)."
- Set the **store.typequota.enable** option to **1**. For example, type the following command:

```
msconfig set store.typequota.enable 1
```

Methods of Setting Message-Type Quotas

Use one of the following methods to set quotas for message types:

- Set message-type quotas for a user with the LDAP attributes **mailQuota** or **mailMsgQuota** (or both).

For information about how to set quota roots with these attributes, see the **mailQuota** and **mailMsgQuota** entries in the *Schema Reference*.

- Set default message-type quotas that apply to all individual users when the **mailQuota** and **mailMsgQuota** attributes are not set.

To set default quotas, use the **store.defaultmessagequota** or **store.defaultmailboxquota** option (or both).

For information about how to set quota roots with these options, see "[Managing Message Store Quotas](#)."

When you set a quota for the message type with a **msconfig** option or LDAP attribute, you must use the quota root specified with the *store.message_type.mtindex:x.quotaroot* option.

Example of a Message-Type Quota Root

The example described in this section sets the following quotas for the user **joe**:

- The default mailbox storage quota is 40 MBytes.
- The default mailbox message quota is 5000.
- The storage quota for the Archive folder is 100 MBytes.
- The storage quota for text message types is 10 MBytes.
- The message quota for text message types is 2000.
- The storage quota for voice message types is 10 MBytes.
- The message quota for voice message types is 200.

This quota root permits greater storage in the Archive folder (100 MBytes) than in all the other folders and message types combined (60 MBytes). Also, no message limit is set for the Archive folder. In this example, only storage limits matter for archiving.

The message types have both storage and number-of-message quotas. The message-type quotas apply to the sum of all messages of those types, whether they are stored in the Archive folder or in any other folder.

The default mailbox quotas apply to all messages that are not text or voice message types and are not stored in the Archive folder. That is, the message-type quotas and Archive quota are not counted as part of the default mailbox quotas.

To set the quota root in this example:

1. Configure the **store.message.type.mtindexx.quotaroot** option as follows:

```
msconfig set store.message.type.mtindex:1.quotaroot text
msconfig set store.message.type.mtindex:2.quotaroot voice
```

2. Configure the **mailQuota** attribute for the user **joe** as follows:

```
mailQuota: 20M;#text%10M;#voice%10M;Archive%100M
```

3. Configure the **mailMsgQuota** attribute for the user **joe** as follows:

```
mailMsgQuota: 5000;#text%2000;#voice%200
```

When you run the **getquotaroot** IMAP command, the resulting IMAP session displays all quota roots for the user **joe**'s mailbox, as shown here:

```
1 getquotaroot INBOX
* QUOTAROOT INBOX user/joe user/joe/#text user/joe/#voice
* QUOTA user/joe (STORAGE 12340 20480 MESSAGE 148 5000)
* QUOTA user/joe/#text (STORAGE 1966 10240 MESSAGE 92 2000)
* QUOTA user/joe/#voice (STORAGE 7050 10240 MESSAGE 24 200)
2 getquotaroot Archive
* QUOTAROOT user/joe/Archive user/joe/#text user/joe/#voice
* QUOTA user/joe/Archive (STORAGE 35424 102400)
* QUOTA user/joe/#text (STORAGE 1966 10240 MESSAGE 92 2000)
* QUOTA user/joe/#voice (STORAGE 7050 10240 MESSAGE 24 200)
```

Expiring Messages by Message Type

The expire and purge feature enables you to move messages from one folder to another, archive messages, and remove messages from the message store, according to criteria you define in expire rules. You perform these tasks with the **imexpire** utility. Because the **imexpire** utility is run by the administrator, it bypasses quota enforcement. For information about how to write expire rules and use the **imexpire** utility, see "[Configuring Message Expiration \(Tasks\)](#)."

You can write expire rules so that messages of different types are expired according to different criteria. The expire feature is extremely flexible, offering many choices for setting expire criteria. This section describes one example in which text and voice messages are expired according to different criteria.

The following example assumes you have configured text and voice message types as follows:

```
store.message.type.mtindex:1 text/plain
store.message.type.mtindex:2 multipart/voice-message
```

Assume also that the message store is configured to read the Content-Type header field to determine the message type.

Example: Sample Rules for Expiring Different Message Types

```
TextInbox.folderpattern: user/*/INBOX
TextInbox.messageheader.Content-Type: text/plain
TextInbox.messagedays: 365
TextInbox.action: fileinto:Archive
VoiceInbox.folderpattern: user/*/INBOX
VoiceInbox.messageheader.Content-Type: multipart/voice-message
VoiceInbox.savedays: 14
VoiceInbox.action: fileinto:OldMail
VoiceOldMail.folderpattern: user/*/OldMail
VoiceOldMail.messageheader.Content-Type: multipart/voice-message
VoiceOldMail.savedays: 30
VoiceOldMail.action: fileinto:Trash
Trash.folderpattern: user/*/Trash
Trash.savedays: 7
Trash.action: discard
```

In this example, text messages and voice mail are expired in different ways, and they follow different schedules, as follows:

- Text messages are moved from a user's inbox to the user's Archive folder one year after they arrive in the message store.
- Voice mail is moved from the inbox to the OldMail folder after two weeks. If the user saves a voice message, the saved date is reset, and the message is moved two weeks after the new date.
- Voice mail is moved from the OldMail folder to the Trash folder after 30 days. The user also can save a voice message in the OldMail folder, which postpones the removal of the message for another 30 days after the new saved date.
- Messages of all types are discarded seven days after they are moved to the Trash folder. The expire rules move voice mail to Trash automatically. Text messages are moved to Trash when a user deletes them.

Note: The **savedays** rule causes a message to be expired the specified number of days after the message is saved. In a typical voice mail system, a user can save voice mail on the voice mail menu. For text messages, a message is saved when it is moved to a folder. The **messagedays** rule causes a message to be expired the specified number of days after it first arrives in the message store, no matter which folder it is stored in or how often it is moved.

Managing Shared Folders

This information describes the tasks that you use to administer shared folders. See "Shared Folders Overview" for conceptual information.

Topics:

- [Specifying Sharing Attributes for Private Shared Folders](#)
- [To Create a Public Shared Folder](#)
- [To Grant Folder Access Rights Based on Group Membership](#)
- [To Set or Change a Shared Folder's Access Control Rights](#)
- [Enabling or Disabling Listing of Shared Folders](#)
- [Setting Up Distributed Shared Folders](#)
- [Monitoring and Maintaining Shared Folder Data](#)

Specifying Sharing Attributes for Private Shared Folders

A private shared folder is a normal folder, created by users in the same way that they create other folders. A folder becomes "shared" when its owner grants access rights to other users or groups. Methods to manage folder access include:

- Many IMAP clients
- Convergence web client
- Oracle Communications Messaging Server "readership" command, for mail administrators

Table 61-1 explains the `msconfig` options that pertain to private shared folders:

Table 61-1 *Disabling Quota Enforcement*

msconfig Option	Description	Default
<code>store.privatesharedfolders.restrictanyone</code>	If enabled (1), disallows regular users from setting the permission on private shared folders to anyone .	0
<code>store.privatesharedfolders.restrictdomain</code>	If enabled (1), disallows regular users sharing private folders to users outside of their domain.	0
<code>store.privatesharedfolders.shareflags</code>	If disabled (0), users of a shared folder have their own set of flags (for example, seen , deleted , and so on) for messages in that folder. If enabled (1), a single set of flags is shared between all users of each shared folder.	0

To Create a Public Shared Folder

Public shared folders must be created by the mail administrator because they require access to the LDAP database as well as the "readership" and "mboxutil" commands.

1. Set the userid for Public shared folders. The `store.publicsharedfolders.user` option specifies the userid to act as a container for all public shared folders (see "Shared Folders Overview"). Typically, this is simply *public*. The default is NULL (unset).

```
msconfig set store.publicsharedfolders.user public
```

2. Create an LDAP entry for that user. The `uid` must match that specified by `store.publicsharedfolders.user`, for example:

```
dn: cn=public,ou=people,o=example.com,o=ISP
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: inetUser
objectClass: ipUser
objectClass: inetMailUser
objectClass: inetLocalMailRecipient
objectClass: nsManagedPerson
objectClass: userPresenceProfile
cn: public
mail: public@example.com
mailDeliveryOption: mailbox
mailHost: manatee.example.org
uid: public
inetUserStatus: active
mailUserStatus: active
mailQuota: -1
mailMsgQuota: 100
```

3. Create folders within the **public** account by using the "mboxutil" command, for example:

```
mboxutil -c user/public/gardening
```

4. Use the "readership" command to grant rights to allow users to access the folder. For example, the following command gives everyone in the **example.com** domain lookup, read, and posting access to the public folder **gardening**:

```
readership -s user/public/gardening anyone@example.com lrp
```

The name *anyone@domain* is a special case to designate all users in the specified domain. It does not correspond to any user or group definition in LDAP. The name **anyone** without specifying a domain indicates anyone in any domain. The following command grants the user whose `uid` is **kelly** the same access rights as the owner of the folder:

```
readership -s user/public/gardening kelly@example.com lrswipcdan
```

For individual users, you only need to supply a domain name with hosted domains. Do not use a domain name if the user to whom access is being granted is in the default domain.

See the "readership" command for a list of the "ACL Rights Characters" and their meanings.

To Grant Folder Access Rights Based on Group Membership

In the previous examples, "ACL Rights Characters" have been granted to individual users or to the special case names **anyone** or **anyone@<domain>**. You can also grant rights based on group membership. Members of such a group are identified by having the **aclGroupAddr** attribute. For more information on the **aclGroupAddr** see the *Schema Reference*.

For example, a group called **tennis@example.com** has 25 members and the members have decided that they would like to create a shared folder to store all email going to this group address and to allow members of the group to access that shared folder.

The mail administrator uses the "readership" command to grant group access rights. A group name is distinguished from individual user names by the prefix "group=".

1. Create the folder. In this example, the team decided to use a private shared folder. The user **gregk** could have created the folder by using a mail client, or the mail administrator could have created it by using the "mboxutil" command, for example:

```
mboxutil -c user/gregk/gardening
```

If the team were using a public shared folder, the mail administrator would have had to create it:

```
mboxutil -c user/public/gardening
```

2. Use the "readership" command to grant **lookup**, **read**, and **posting** access privileges to the group:

```
readership -s user/gregk/gardening group=tennis@example.com lrp
```

3. Assign group membership to the individual users. For the purpose of folder access control, group membership is determined by the **aclGroupAddr** attribute on the LDAP entry of the individual users. Add the attribute-value pair **aclGroupAddr=<group-name>** to the user entry of every member of the group, for example:

```
aclGroupAddr: tennis@example.com
```

To create group objects in LDAP, you could use the **aclGroupAddr** attribute as the basis for a dynamic group, for example:

```
memberURL:
ldap:///o=example.com??sub?(&(aclGroupAddr=tennis@example.com)(objectclass=inetmailuser))
```

However, note that the LDAP group object with mail address **tennis@example.com** is not used for determining group membership for the purpose of shared folder access. What matters is that the "xxx" value in **group=xxx** on the **readership** command matches the value of the **aclGroupAddr** attribute on the user's LDAP object.

Also note that if you use the **aclGroupAddr** attribute as the criteria for a dynamic group, you should check to make sure that attribute is indexed properly for such lookups.

To Set or Change a Shared Folder's Access Control Rights

Users can set or change the access control for a shared folder by using Convergence. Administrators can set or change the access control for a shared folder using the "readership" command line utility. The command has the following form:

```
readership -s <foldername> <identifier> <rights_chars>
```

where **foldername** is the name of the folder for which you are setting rights, **identifier** is the person or group to whom you are assigning the rights, and **rights_chars** are the rights you are assigning. For the meaning of each character, see [Table 61-3](#), "readership Options" for more information.

Note: **anyone** is a special identifier. The access rights for **anyone** apply to all users. Similarly, the access rights for **anyone@domain** apply to all users in the same domain. For the identifier, only supply a domain name with hosted domains. Do not use a domain name if the folder is in the default domain.

Shared Folder Examples

- To assign everyone in the **example** domain to have lookup, read, and email marking (but not posting) access to the public folder called **golftournament**, type the following command:

```
readership -s user/public/golftournament anyone@example lwr
```

- To assign the same access to everyone on the message store type the following command:

```
readership -s user/public/golftournament anyone lwr
```

- To assign lookup, read, email marking, and posting rights to a group, type the following command:

```
readership -s user/public/golftournament group=golf@example.com lwrp
```

- If you want to assign administrator and posting rights for this folder to an individual, **jdoe**, type the following command:

```
readership -s user/public/golftournament jdoe@example.com lwrpa
```

- To deny an individual or group access to a public folder, prefix the **userid** with a dash. For example, to deny lookup, read, and write rights to **jsmith**, type the following command:

```
readership -s user/public/golftournament -jsmith@example.com lwr
```

- To deny an individual or group an access right, prefix the ACL rights character with a dash. For example, to deny posting rights to **jsmith**, type the following command:

```
readership -s user/public/golftournament jsmith@example.com -p
```

- To remove an individual or group access right setting from a folder, set it to an empty set. This is different from an ACL to deny access:

```
readership -s user/public/golftournament jsmith@example.com ""
```

Note: Posting messages to a shared folder by using the *uid+folder@domain* address requires that the **p** (post) access right be used with the readership command. See "To Set or Change a Shared Folder's Access Control Rights."

Enabling or Disabling Listing of Shared Folders

Use the `store.sharedfolders` option to enable to disable listing of shared folders when responding to an IMAP `LIST` command. Setting the option to `0` disables it. The setting is enabled by default (set to `1`). `SELECT` and `LSUB` commands are not affected by this option. The `LSUB` command returns every subscribed folder, including shared folders. Users can `SELECT` the shared folders they own or are subscribed to.

Setting Up Distributed Shared Folders

Normally, shared folders are only available to users on a particular message store. Messaging Server, however, enables you to create **distributed shared folders** that can be accessed across multiple message stores. That is, access rights to distributed shared folders can be granted to any users within the group of message stores. However, web mail clients do not support remote shared folders access. Users can list and subscribe to the folders, but they cannot view or alter the contents.

Distributed shared folders require the following:

- Every message store `userid` must be unique across the group of message stores.
- The directory data across the deployment must be identical.

The remote message stores (that is the message stores that do not hold the shared folder) must be configured as proxy servers by setting the configuration variables listed in [Table 61-2](#).


Table 61-2 Variables for Configuring Distributed Shared Folders

Name	Value	Data Format
<code>base.proxyserverlist</code>	Message store server list to list shared folders from	space-separated strings
<code>base.proxyadmin</code>	Default store admin login name	string
<code>base.proxyadminpass</code>	Default store admin password	string
<code>proxy:hostname.imapadmin</code>	Store admin login name for a specific host if different from <code>base.proxyadmin</code>	string
<code>proxy:hostname.imapadminpass</code>	Store admin password for a specific host if different from <code>base.proxyadminpass</code>	string

Setting Up Distributed Shared Folders-Example

The following figure shows a distributed folder example of three message store servers called `StoreServer1`, `StoreServer2`, and `StoreServer3`.

Figure 61-1 Distributed Shared Folders-Example



These servers are connected to each other as peer proxy message stores by setting the appropriate `msconf` options. Each server has a private shared folder: `golf` (owned by

Han), *tennis* (owned by Kat), and *hurling* (owned by Luke). In addition, there are two public shared folders called *press_releases* and *Announcements*. Users on any of the three servers can access any of these three shared folders.

The following example shows the ACLs for each server in this configuration.

```
$ StoreServer1 :> imcheck -d lright.db
Ed: user/Han/golf
Ian: user/Han/golf
anyone: user/public/press_releases

$ StoreServer2 :> imcheck -d lright.db
Jan: user/Kat/tennis
Ann: user/Kat/tennis
anyone: user/public+Announcements user/public+press_releases

$ StoreServer3 :> imcheck -d lright.db
Tuck: user/Ian/hurling
Ed: user/Ian/hurling
Jac: user/Ian/hurling
anyone: user/public/Announcements
```

Monitoring and Maintaining Shared Folder Data

The **readership** command-line utility enables you to monitor and maintain shared folder data which is held in the **folder.db**, **peruser.db**, and **lright.db** files. **folder.db** has a record for each folder that holds a copy of the ACLs. The **peruser.db** has an entry per user and mailbox that lists the various flags settings and the last date the user accessed any folders. The **lright.db** has a list of all the users and the shared folders for which they have lookup rights.

The **readership** command-line utility takes the following options:

Table 61–3 *readership* Options

Options	Description
-d days	Returns a report, per shared folder, of the number of users who have selected the folder within the specified days.
-p months	Removes data from the peruser.db for those users who have not selected their shared folders within the specified months.
-l	List the data in lright.db .
-sfolder_identifier_rights	Sets access rights for the specified folder. This updates the lright.db as well as the folder.db .

Using the various options, you can perform the following functions:

- [To Monitor Shared Folder Usage](#)
- [To List Users and Their Shared Folders](#)
- [To Remove Inactive Users](#)
- [To Set Access Rights](#)

To Monitor Shared Folder Usage

To find out how many users are actively accessing shared folders, use the following command:

```
readership -d <days>
```

where **days** is the number of days to check. Note that this option returns the number of active users, not a list of the active users.

Example: To find out the number of users who have selected shared folders within the last 30 days:

```
readership -d 30
```

To List Users and Their Shared Folders

To list users and the shared folders to which they have access, use the following command:

```
imcheck -d lright.db
```

Example output:

```
$ imcheck -d lright.db
group=lee-staff@example.org: user/user2/lee-staff
richb: user/golf user/user10/Drafts user/user2/lee-staff user/user10/Trash
han1: user/public+hurling@example.org user/golf
gregk: user/public+hurling@example.org user/heaving user/tennis
```

To Remove Inactive Users

If you want to remove inactive users (those who have not accessed shared and other folders in a specified time period), use the following commands:

1. This command writes the inactive mailboxes to a file:

```
mboxutil -o [-w <file>] [-t <number of days>]
```

2. This command removes the mailboxes in a given file:

```
mboxutil -d -f <file>
```

Example: Remove users who have not accessed folders for the past six months (180 days) using a file named `inactive_users`:

```
mboxutil -o -w inactive_users -t 180
mboxutil -d -f inactive_users
```

To Set Access Rights

You can assign access rights to a new public folder, or change access rights on a current public folder.

For an example of how to set access rights with this command, see ["To Set or Change a Shared Folder's Access Control Rights."](#)

Upgrading the Message Store

This chapter describes the data components and tools of message store, and how they affect data upgrades. It provides the technical background for upgrade planning. If you are using the upgrade procedure described in the *Messaging Server Installation and Configuration Guide*, you do not need to concern yourself with this level of message store technical detail. However, if you need to customize your message store upgrade process, use this information as a guideline.

For specific information about significant design changes in the message store between versions that can impact disk I/O and performance, see "[Significant Changes in the Message Store Between Versions](#)."


Topics:

- [Architecture and Components](#)
- [Message Store Component Version Compatibilities](#)
- [IMAPD, MSHTTPD and Convergence](#)
- [Upgrading From Messaging Server 32-bit to 64-bit](#)
- [Migrating from x86 to SPARC](#)
- [stored -r](#)
- [ims_db_upgrade](#)
- [Significant Changes in the Message Store Between Versions](#)

Architecture and Components

The following figure shows the message store architecture and its components. See "[Message Store Directory Layout](#)" for a view of the message store file structure.

Figure 62-1 Message Store and Components



The message store uses the *Berkeley Database* provided by Sleepycat Software and, since 2006, Oracle Corporation. The message store database files are stored in a directory called **mboxlist** (see "Message Store Directory Layout," and so it is often called the **mboxlist** database.

The message store has used a number of versions of the Berkeley Database over the course of its existence. Thus, when you upgrade your message store, the Berkeley Database may also be upgraded. The database engine upgrade has complexities and implications for the data upgrade. (These complexities and implications are handled in the Message Store upgrade tools and instructions described in the *Messaging Server Installation and Configuration Guide*, but these details are described here for custom upgrade plans.)

The **mboxlist** database is stored in the *BTREE database files*. The BTREE database files store information about the message store users and mailboxes (see "[Message Store Directory Layout](#)"). The number and types of files, as well as the structure of files themselves have varied over the life of the Berkeley Database.

The Berkeley Database is transactional, so each transaction is logged in a *database log file*. Database log files are used for recovery. All database changes are recorded in the database log file. When the server is crashed and restarted, it uses the log file to bring the database back to a consistent state. Before you do an upgrade, make sure you have a clean shutdown and recovery by running **stored -r**.

Message Store Component Version Compatibilities

Table 62–1 contains the version number of various database components in the message store with respect to upgrade.

Table 62–1 Message Store Database Components

Messaging Server	Mailbox	Berkeley Database	Database BTREE	Database Log Files
4.15	1_1	2.6	6	2
5.x	1_2	2.6	6	2
6.0	1_2	3.2.9	8	3
6.1	1_2	4.2	9	8
6.2	1_2	4.2	9	8
6.3	1_2	4.4	9	11
7.0	1_3	4.4	9	11
7 Update 1	1_3	4.7.25	9	14
7 Update 2	1_3	4.7.25	9	14
7 Update 3	1_3	4.7.25	9	14
7 Update 4	1_3	4.7.25	9	14
7 Update 5	1_4	5.3.21	9	19
8.0	1_5	6.1.19	10	22
8.0.1	1_5	6.1.26	10	22

In general, when you upgrade Oracle Communications Messaging Server from one version to another, you also have to upgrade the components that have a different version. Certain message components are not compatible with other components. The following section describes the various compatibilities and incompatibilities.

Upgrading the Mailboxes

The message store upgrades mailboxes from version 1_1-1_2 and 1_2-1_4 to the current version automatically. Usually, a mailbox is upgraded when the end users select their mailboxes, or when messages are delivered to the mailboxes after the message store software is upgraded. The message store checks the mailbox version in the mailbox header and upgrades the mailbox if needed.

Mailbox upgrade increases the load on the server. Disable the non-essential tasks during the transition period, such as:

- Disable peruser.db archive:
`msconfig -o store.seenckpinterval -v 0` (Unified Configuration)
or
`configutil -o local.store.seenckpinterval -v 0` (legacy configuration)
- Disable last access time tracking:
`msconfig -o base.enablelastaccess -v 0` (Unified Configuration)
or
`configutil -o local.enablelastaccess -v no` (legacy configuration)
- Decrease DB snapshot frequency
- Decrease log level
- Temporarily disable **delivery**
- Temporarily disable **impurge**

You can upgrade the mailboxes manually with **imcheck -H** during off-peak hours. The **imcheck** command opens every mailbox, which triggers the upgrade. You can run this command to make sure all the mailboxes are upgraded.

To upgrade from mail version 1_1 to 1_4, you have to use a migration tool such as **imsbackup** and **imsrestore**.

Downgrading mailbox versions is not automatic. You have to use a migration tool such as **imsbackup** and **imsrestore** to downgrade a mailbox.

Upgrading and Downgrading the Berkeley Database (BDB)

The Message Store uses the Berkeley Database to store various data. The email messages themselves are not part of the data stored in the BDB. If the database is upgraded to a version that is incompatible with the previous version, the database files will become incompatible with older versions of the Message Store. We recommend, therefore, that you make a backup copy of the database before the upgrade in case you want to back out of the upgrade.

The primary database which needs to be backed up is located by default in the **mboxlist/** directory, and consists of **.db** files, and **log** files. The **__*** files are temp files for the database and do not need to be copied. A correct copy of the database ensures data is in a consistent state. The message store and utilities must be shut down. Running **stored -r** will make sure the cache files are synced to the database files. Both database and log files are required.

If you need to back out of a patch or update which upgrades the Berkeley Database to a version that is incompatible with previous versions and you did not make a backup, you may be able to rely on a previous database snapshot. Database snapshots are located by default in the **dbdata/** directory. A valid database snapshot directory will

have a **.verified** file. The **.verified** file indicates that the snapshot has been recovered, verified, and is ready to be used.

Normally when the store is brought up, the **stored** process replaces the **mboxlist** directory with any snapshots needed. In the case where you have backed out of a database upgrade, **stored** may not take into account that a downgrade has occurred. It may therefore be necessary to replace the valid snapshot manually. To do this, move the current database files in **mboxlist/** out of the directory and move all the files from the chosen snapshot directory into the **mboxlist/** directory. Be sure to remove the **__*** tmp files as well. Note that if the store is configured with **store.dbtmpdir**, the tmp files will be in a different location.

If you have no database backup and no valid snapshots, it may be necessary to move the upgraded database out of the way, and rebuild it from scratch. Under normal circumstances, the Message Store rebuilds the database while allowing users to access their mail. Since doing this puts a heavier load on the system, you should create proper database backups instead.

Normal reparation of the database should be done after putting an older version in place by running the **reconstruct -m** and **reconstruct -r** commands.

Some of these manual requirements might be addressed in future releases.

Note that the Berkeley database consist of a number of BTREE files, LOG files and temporary files (tmp file location is configurable with **store.dbtmpdir**). In order to upgrade the database, run "**stored -r**" before replacing the new libraries and binaries. Note that "**stored -r**" runs automatically during the proper upgrade process.

In the unlikely event that **stored -r** fails, check the store log files to determine the cause, run **stored** alone to perform a database recovery, and run **stored -r** again. In normal circumstances, this should not be a problem unless there is some underlying system problem which you should resolve before upgrading.

Database BTREE File

BTREE version 8, 9, and 10 are compatible. Upgrade is not needed.

To upgrade the BTREE files from version 6 (BDB 2.6) to a higher version, copy the database files from the old location (example: **/usr/iplanet/server5/msg-store/store/mboxlist** for 5.2) to the new **mboxlist** location (example: **/var/opt/SUNWmsgsr/store/mboxlist**), then run **ims_db_upgrade**.

Database Log Files

Database log files cannot be upgraded. When the log version changes, run the legacy version of "**stored -r**" to process the log files (recover the database) before upgrading. Do not remove the old log files.

The following message is logged when the server restarts.

'Skipping log file...historic log version'

The store daemon creates a new log file with the new version.

IMAPD, MSHTTPD and Convergence

The webmail server (**mshttpd**) uses **imapd** to access the Message Store. **imapd** and **mshttpd** in 6.3 and 7.x are fully compatible, so simultaneous upgrade is not required. To prevent memory problems due to redundant IMAP sessions from **mshttpd**, you should run with only one **mshttpd** process. If you serve a lot of concurrent webmail

users, this might require an upgrade to 64-bit so that you have enough virtual address space for the process.

Convergence requires webmail server. To use Convergence, you must upgrade **mshttpd** to the current version.

Upgrading From Messaging Server 32-bit to 64-bit

Run the legacy version of "stored -r" before upgrading the Messaging Server software from 32-bit to 64-bit.

If you run the 64-bit version, then it is more efficient to run with only one **mshttpd** process.

You might also want to reduce the number of **imapd** process.

Migrating from x86 to SPARC

The message store data formats are architecture dependent. You cannot transfer any data components in the matrix from one architecture to another directly. To migrate the message store from an x86 machine to a SPARC machine (or from SPARC to x86), use **imsbackup** and **imsrestore** or **rehostuser**.

stored -r

stored -r performs a final recovery and cleanly removes the database temp files to prepare the database for an upgrade.

For example:

```
/opt/SUNWmsgsr/lib/stored -r
removing mboxlist environment ... done
removing lock environment ... done
removing session db ... done
```

Make sure **stored -r** completes successfully.

The **stored -r** command requires the watcher process. Make sure that the watcher is running before you perform **stored -r**.

ims_db_upgrade

ims_db_upgrade copies the database files to a backup directory, upgrades the database files to the current version and validates the new database. If upgrade is not successful, the backup files are restored. If the database is validated, the backup files are removed.

Downgrading

Mailbox and BDB downgrade are not supported. To downgrade a message store to an older version with incompatible mailboxes or databases, you must restore the mailboxes and mboxlist database from backup.

Significant Changes in the Message Store Between Versions

This document describes the significant design changes in the message store that can impact disk I/O and performance. Administrators and deployment designers may

need to be aware of these changes for deployments where mboxlist, index, and message data are on different file systems or pools.

Topics:

- [Changes from Messaging Server 6.3 to Messaging Server 7.0](#)
- [Changes from Messaging Server 7 to Messaging Server 7 Update 1](#)
- [Changes from Messaging Server 7 Update 1 to Messaging Server 7 Update 5](#)

Changes from Messaging Server 6.3 to Messaging Server 7.0

The following changes were made between Messaging Server 6.3 and Messaging Server 7.0:

- [Changes to store.idx](#)
- [Message Store Maintenance Queue and impurge](#)
- [Mailbox Self-Healing \(Auto-Repair\)](#)

Changes to store.idx

The **store.idx** file is separated into two or more files. The **store.idx** file contains the mailbox's meta data and a 128 bytes fixed length record for every message in the folder. The **store.cN** files contain the variable length cache records. The average cache record size is approximately 2 kilobytes.

Benefits

- Mailbox expunge is much faster
- Supports very large mailboxes (up to 33554431 messages)

I/O Impact

- More files (inodes) on the index partitions
- More I/O for very small mailboxes
- Initial mailbox access triggers mailbox upgrade automatically

Message Store Maintenance Queue and impurge

A BDB queue is added to the message store to manage maintenance tasks. Cache file purge and message file cleanup tasks are queued and executed by **impurge** which runs continuously.

Benefits

- Less I/O overall (especially write)
- Event driven message store maintenance
- Higher performance
- Eliminates sudden increase in load during nightly cleanup

I/O Impact

- Less write access
- Higher disk space usage (approximately 10%)

Mailbox Self-Healing (Auto-Repair)

Mailbox corruption is detected and scheduled for repair automatically by **impurge**.

Benefits

- Increases availability
- Reduces mailbox administration cost

I/O Impact

- Minor I/O increase due to error detection and repair.

Changes from Messaging Server 7 to Messaging Server 7 Update 1

The following change was made between Messaging Server 7 and Messaging Server 7 Update 1:

- [Berkeley Database Upgrade](#)

Berkeley Database Upgrade

The Berkeley Database was upgraded from version 4.4 to version 4.7.25.

Benefits

- Higher throughput

I/O Impact

- None

Changes from Messaging Server 7 Update 1 to Messaging Server 7 Update 5

The following changes were made between Messaging Server 7 Update 1 and Messaging Server 7 Update 5:

- [Changes to the Owner's Seen and Deleted Flags](#)
- [Immediate flag update and state sharing](#)
- [Change to the service.imap.capability.condstore option](#)
- [Changes to the Berkeley Database](#)
- [Changes to mboxlist and lockdir BDB environments](#)

Changes to the Owner's Seen and Deleted Flags

The owner's seen and deleted flags have been moved from the Berkeley Database to the **store.idx** file.

Benefits

- Higher throughput by reducing Berkeley Database lock contentions
- Performance bottlenecks can be fixed by adding more spindles

I/O Impact

- Less I/O on the mboxlist file system
- More I/O on the index file system
- Initial mailbox access triggers mailbox upgrade automatically

Immediate flag update and state sharing

Flags and ACL changes are flushed to the disk and shared immediately across IMAP sessions. New message arrival are notified immediately.

Benefits

- IMAP sessions are more up-to-date

I/O Impact

- More I/O on the index file system

Change to the `service.imap.capability.condstore` option

The `service.imap.capability.condstore` option is enabled by default.

Benefits

- IMAP clients can utilize CONDSTORE
- Improve overall performance (when CONDSTORE is utilized)

I/O Impact

- Minimal (because flags and MODSEQ are updated together)

Changes to the Berkeley Database

The Berkeley Database was upgrade from version 4.7.25 to version 5.3.21. The default cache size increased from 16 MB to 64 MB. Moved the maintenance queue to the mboxlist environment.

Benefits

- Increase performance and scalability
- Simplify store maintenance tasks

I/O Impact

- A small increase in database environment (`tmp` file) size.

Changes to mboxlist and lockdir BDB environments

Mboxlist and lockdir DB environments (`store.dbtmpdir` and `local.lockdir`) default to `tmpfs`.

Benefits

- Better performance

I/O Impact

- Less I/O on the data root file system.

Message Store Automatic Recovery On Startup

This information provides a conceptual overview of the startup and automatic recovery process of **stored**, and Message Store Database Snapshot. See "[Administering Message Store Database Snapshots \(Backups\)](#)" for task information.

Topics:

- [Overview of Automatic Recovery on Startup](#)
- [Automatic Startup and Recovery Theory of Operations](#)

Overview of Automatic Recovery on Startup

Message store data consists of the messages, index data, and the message store database. While this data is fairly robust, on rare occasions there may be message store data problems in the system. These problems are indicated in the default log file, and almost always are fixed transparently. In rare cases an error message in the log file may indicate that you need to run the **reconstruct** utility. In addition, as a last resort, messages are protected by the backup and restore processes described in "[Backing Up and Restoring the Message Store.](#)"

The message store automates many recovery operations which were previously the responsibility of the administrator. These operations are performed by message store daemon **stored** during startup and include database snapshots and automatic fast recovery as necessary. **stored** thoroughly checks the message store's database and automatically initiates repairs if it detects a problem.

stored also provides a comprehensive analysis of the state of the database by writing status messages to the default log, reporting on repairs done to the message store and automatic attempts to bring it into operation.

Automatic Startup and Recovery Theory of Operations

The **stored** daemon starts before the other message store processes. It initializes and, if necessary, recovers the message store database. The message store database keeps folder, quota, subscription, and message flag information. The database is logging and transactional, so recovery is already built in. In addition, some database information is copied redundantly in the message index area for each folder.

Although the database is fairly robust, on the rare occasions that it breaks, in most cases **stored** recovers and repairs it transparently. However, whenever **stored** is restarted, you should check the default log files to make sure that additional administrative intervention is not required. Status messages in the log file indicate that **reconstruct** should be run if the database requires further rebuilding.

Before opening the message store database, **stored** analyzes its integrity and sends status messages to the default log under the category of **warning**. Some messages are useful to administrators and some messages consist of coded data to be used for internal analysis. If **stored** detects any problems, it attempts to fix the database and try starting it again.

When the database is opened, **stored** signals that the rest of the services may start. If the automatic fixes failed, messages in the default log specify what actions to take. See ["Error Messages Signifying that reconstruct is Needed."](#)

After most recoveries, the database is usually be up-to-date and no further action is required. However, some recoveries require a **reconstruct -m** to synchronize redundant data in the message store. Again, this is stated in the default log, so it is important to monitor the default log after a startup. Even though the message store seems to be up and running normally, it is important to run any requested operations such as **reconstruct**.

Another reason for reading the log file is to determine what caused damage to the database in the first place. Although **stored** is designed to bring up the message store regardless of any problem on the system, you should ascertain cause of the database damage as this may be a sign of a larger hidden problem.

Error Messages Signifying that reconstruct is Needed

This section describes the type of error messages that require **reconstruct** to be run.

When the error message indicates mailbox error, run **reconstruct <mailbox>**. Example:

```
Invalid cache data for msg 102 in mailbox user/joe/INBOX. Needs reconstruct
```

```
Mailbox corrupted, missing fixed headers: user/joe/INBOX
```

```
Mailbox corrupted, start_offset beyond EOF: user/joe/INBOX
```

When the error message indicates a database error, run **reconstruct -m**. Example:

```
Removing extra database logs. Run reconstruct -m soon after startup to resync  
redundant data
```

```
Recovering database from snapshot. Run reconstruct -m soon after startup to resync  
redundant data
```

Message Store Database Snapshot Theory of Operations

This section describes concepts about the message store database snapshot. See ["Administering Message Store Database Snapshots \(Backups\)"](#) for a description of related tasks.

A snapshot is a hot backup of the database and is used by **stored** to restore a broken database transparently. This is much quicker than using **reconstruct** to rebuild the entire database from scratch with the information stored in the message and index partitions.

Snapshots of the database are taken automatically by the scheduler. The default snapshot schedule consists of a full snapshot every day and incremental snapshots every 10 minutes. (Note that older versions of Oracle Communications Messaging Server have a more frequent default incremental snapshot schedule).

If the recovery process decides to remove the current database because it is determined to be bad, **stored** will move it into the **removed** directory if it can. This allows the database to be analyzed if desired.

Message Store Database Snapshot Interval and Location

There should be five times as much space for the database and snapshots combined. It is highly recommended that the administrator reconfigure snapshots to run on a separate disk, and that it is tuned to the system's needs.

If **stored** detects a problem with the **mboxlist** database at startup, the most recent verified snapshot is automatically restored. Two snapshot options can be configured: the location of the snapshot file and number of snapshots saved. These options are shown in "[Message Store Database Snapshot Options](#)."

Having a snapshot interval which is too small results in a frequent burden to the system and a greater chance that a problem in the database is copied as a snapshot. Having a snapshot interval too large can create a situation where the database holds the state it had back when the snapshot was taken.

A snapshot updates all the snapshots with the current data in the **mboxlist** database.

The ultimate role the snapshot plays is to get the system as close to up-to-date and ease the burden of the rest of the system trying to rebuild the data on the fly.

Message Store Database Snapshot Options

[Table 63-1](#) shows the snapshot options that you set with the **msconfig** command.

Table 63-1 *Message Store Database Snapshot Options*

Option	Description
store.snapshotpath	Location of message store database snapshot files. Either existing absolute path or path relative to the store directory. Default: dbdata/snapshots
store.snapshotdirs	Number of snapshots to maintain. Do not set this to more than 3. Default: 3

Message Store Maintenance Queue

Only the index file is purged when a mailbox is expunged. The purging of cache records is deferred until the amount of expunged data has exceeded a configurable limit. In addition, the message store uses a maintenance queue to schedule mailbox purge and repair tasks. Mailbox corruptions detected by the message store are also queued for repair automatically.


Topics:

- [Message Store Maintenance Queue Overview](#)
- [Displaying the Maintenance Queue](#)
- [Deleting, Expunging, Purging, and Cleaning Up Messages](#)
- [Mailbox Self Healing \(Auto Repair\)](#)
- [Maintenance Queue Configuration Options](#)
- [The `impurge` Command](#)

Message Store Maintenance Queue Overview

The following diagram summarizes the major components and their interactions with the queue.

Figure 64–1 Message Store Maintenance Queue



When a mailbox is expunged, the message store processes update the number of expunged messages and enqueue the mailbox name for purging when the number of expunged messages has exceeded the configured limit (set by the `store.cleanupsize` option). Similarly, when store corruptions are detected by the store processes, they enqueue the mailbox name for repair. A maintenance queue entry contains a mailbox

name, a task ID, and a timestamp. The **impurge** process dequeues the entries, checks the timestamp and mailbox size or last repair time and performs the maintenance task if required. The mailbox size and last repair time are stored in the **mboxlist** database. Mailboxes that have already been repaired after the enqueue time are ignored. Obsolete cache and message files are enqueued for removal after the cleanup age has expired.

You can set the **impurge** process to execute as a daemon or through the scheduler. When you use the scheduler, you can configure **impurge** to exit when the queue is empty or when the end time expires. In this way, you can configure **impurge** to run during off peak hours.

impurge executes as a daemon by default.

Displaying the Maintenance Queue

Use the **imcheck -q** command to display the contents of the maintenance queue, for example:

```
imcheck -q
RecNo      Mailbox                                     Timestamp      Action
-----
2          user/username                               20080225095558 Clean
```

The maintenance queue database is stored under the *MessagingServer_home/data/store/mboxlist* directory and persists across Oracle Communications Messaging Server restarts.

Deleting, Expunging, Purging, and Cleaning Up Messages

Message removal is a four steps process:

1. Delete
 - A user deletes a message. This results in the per-user **\Deleted** flag being set on the message. If there is a second client, the deleted flag may not be recognized immediately by that second client. You can set the **imap.immediateflagupdate** option to enable immediate flag update.
2. Expunge and Expire
 - When the mailbox is expunged, messages with the **\Deleted** flag are removed from the **store.idx** file. The message itself is still on disk, but once messages are expunged, clients can no longer restore them. The number of the expunged messages is also recorded in the "Expunged" mailbox meta-data field. This field can be reviewed using the **./imcheck -mmailbox** command.
 - When a message matches an expiration rule during an **imexpire** run, the **imexpire** command removes the message from the **store.idx** file and increments the "Expunged" mailbox meta-data field.
3. Cleanup
 - Cleanup of a folder is scheduled when the number of expunged in the folder messages exceeds **store.cleanupsize**, default is 100. The **store.exp** file is renamed to **store.exp.timestamp**.
 - Obsolete message files and cache files are removed from the store partitions when the **store.cleanupage** period has elapsed. **store.cleanupage** controls the cleanup grace period.

4. Purge

- **impurge** purges the cache files when the number of expunged cache records exceeds **store.purge.count** (default is 500) and the percentage of expunged cache records exceeds **store.purge.percentage** (default is 5%). This is different from the **imsimta purge** command, which purges older versions of the MTA log files.

Mailbox Self Healing (Auto Repair)

The message store performs sanity checks when a mailbox is opened and when index and cache data are accessed. When the message store encounters a mailbox format error, an error is returned to the client and the mailbox is enqueued for repair. **impurge** dequeues the mailbox and repairs the mailbox automatically. The mailbox size and last repair time are stored in the mboxlist database. Mailboxes that have already been repaired after the enqueue time are ignored.

The following is an example of the auto repair process. In this example, the **store.idx** file for a user's INBOX has been removed. When the user accessed their INBOX the following error was reported in the **imap** log file:

```
[25/Feb/2012:09:55:57 +1100] hostname imapd[7264]: Store Critical: Unable to open mailbox user/username: Mailbox does not exist
[25/Feb/2012:09:55:57 +1100] hostname imapd[7264]: Store Error: user/username: Mailbox has an invalid format; repair requested
```

In the **default** log file an **impurge** record indicates that the requested mailbox repair has been processed and the **store.idx** restored:

```
[25/Feb/2012:09:55:58 +1100] hostname impurge[7263]: General Notice: repairing user/username
```

Maintenance Queue Configuration Options

The **store.purge.enable** option enables the purge server daemon process on **start-msg** startup. Various other **store.*** options control the maintenance queue. For more information on these options, see the *Messaging Server Reference*.

The impurge Command

You can use the **impurge** command to manually purge unused cache records and message files in mailboxes when the purge server daemon process is not running, that is, when the **store.purge.enable** option is disabled.

Attempting to run the **impurge** command while the purge server daemon process is running results in the following error message in the Messaging Server **default** log:

```
[24/Feb/2012:14:47:15 +1100] hostname impurge[17986]: General Error: Could not get purge session lock.
Possibly another impurge is running
```

For details, see "[impurge](#)."

Message Store Message Type Overview

This information provides the conceptual background for using message types in the message store. See ["Managing Message Types in the Message Store"](#) for task information.

This chapter includes the following topics:

- [About Message Type](#)
- [Planning the Message-Type Configuration](#)
- [Defining and Using Message Types](#)

About Message Type

A unified messaging application can receive, send, store, and administer messages of many types, including text messages, voice mail, fax mail, image data, and other data formats. The message store allows you to define up to 63 different message types.

One method of managing messages by type is to organize the messages by their types into individual folders.

With the introduction of the message type feature, you do not have to maintain different message types in individual mailbox folders. Once you configure a message type, the message store can identify it, no matter where it is stored. Thus, you can store heterogeneous message types in the same folder. You also can perform the following tasks:

- Track the usage of message types
- Send notifications grouped by message type
- Set and administer different quotas for different message types, whether they are stored in the same folder or different folders
- Move messages from one folder to another, according to criteria configured uniquely for each message type
- Expire messages according to criteria configured for each message type

Planning the Message-Type Configuration

In a unified messaging application, data of heterogeneous formats are given standard internet message headers so that Oracle Communications Messaging Server can store and manage the data. For example, when voice mail is sent to an end-user's phone, a telephone front-end system adds a message header to the incoming voice mail and delivers it to the message store.

To recognize and administer messages of different types, all components of the unified messaging system must use the same message-type definitions and the same header fields to identify the messages.

Before you configure the message store to support message types, you must:

- Plan which message types you intend to use
- Decide on the definition for each message type
- Decide which header field to use

For example, if your application includes phone messages, you can define this message type as "multipart/voice-message" and use the Content-Type header field to identify message types.

You would then configure the telephone front-end system to add the following header information to each phone message to be delivered to the message store:

Content-Type: multipart/voice-message

Next, you would configure the message store to recognize the **multipart/voice-message** message type, as described in the sections that follow.

Defining and Using Message Types

You define a message type by giving it a unique definition such as **multipart/voice-message**. By default, the message store reads the Content-Type header field to determine the message-type. If you prefer, you can configure a different header field to identify the message types.

The message store reads the Content-Type (or other specified) header field, ignoring case. That is, the message store accepts the header field as valid even if the header's combination of uppercase and lowercase letters differs from the expected combination.

The message store reads only the message-type name in the header field. It ignores additional arguments or options.

To define a message type, use the **msconfig** command to set values for the **store.messagestype** option. For instructions, see "[To Configure Message Types](#)."

Configuring a message type allows the message store to identify and manipulate messages of the specified type. It is the first, essential step in administering message types in a unified messaging application.

To take full advantage of the message-type features provided by the message store, you also should perform some or all of the following tasks:

- Configure a JMQ notification plug-in and write Message Queue clients for retrieving notifications that track the status of the message types
- Configure quota roots that apply to each message type
- Write expire rules and set LDAP attribute values to expire and purge messages according to message type

Message Types in IMAP Commands

When you configure the **store.messagestype.mtindex:n.flagname** option for a message type, you create a unique flag that identifies the message type. This flag cannot be modified by end users.

Messaging Server presents the message-type flag as a user flag to IMAP clients. Mapping the message type to a user flag allows mail clients to use simple IMAP commands to manipulate messages by message type.

For example, you can perform the following operations:

- Use the **IMAP FETCH FLAGS** command to display a message-type flag name as a user-defined flag to the client. For a sample use of the **IMAP FETCH FLAGS** command, see ["Example 1: IMAP FETCH Session Based on the Message-Type msconfig Configurations."](#)
- Use a message-type flag as a keyword in an **IMAP SEARCH** command. For a sample use of the **IMAP SEARCH** command, see ["Example 2: IMAP SEARCH Session Based on the Message-Type msconfig Configurations."](#)

The message-type user flag is read only. It cannot be modified by IMAP commands.

The following examples assume that you configure the message-type **msconfig** options with the values shown here:

```
store.message-type.enable 1
store.message-type.mtindex:1.content-type text/plain
store.message-type.mtindex:1.flagname text
store.message-type.mtindex:1.quotaroot text
store.message-type.mtindex:2.content-type multipart/voice-message
store.message-type.mtindex:2.flagname voice_message
store.message-type.mtindex:2.quotaroot voice
```

Example 1: IMAP FETCH Session Based on the Message-Type msconfig Configurations

The following IMAP session fetches messages for the currently selected mailbox:

```
2 fetch 1:2 (flags rfc822)
* 1 FETCH (FLAGS (\Seen text) RFC822 {164}
Date: Wed, 8 July 2006 03:39:57 -0700 (PDT)
From: bob.smith@example.com
To: john.doe@example.com
Subject: Hello
Content-Type: TEXT/plain; charset=us-ascii
* 2 FETCH (FLAGS (\Seen voice_message) RFC822 {164}
Date: Wed, 8 July 2006 04:17:22 -0700 (PDT)
From: sally.lee@example.com
To: john.doe@example.com
Subject: Our Meeting
Content-Type: MULTIPART/voice-message; ver=2.0
2 OK COMPLETED
```

In the preceding example, two messages are fetched, one text message and one voice mail.

The message-type flags are displayed in the format configured with the **store.message-type.mtindex:n.flagname** option.

The Content-Type header fields identify the message types. The message-type names are displayed as they were received in the incoming messages. They use mixed uppercase and lowercase letters and include the message-type arguments such as **charset=us-ascii**.

Example 2: IMAP SEARCH Session Based on the Message-Type msconfig Configurations

The following IMAP session searches for voice messages for the currently selected mailbox:

```
3 search keyword voice_message
* SEARCH 2 4 6
3 OK COMPLETED
```

In the preceding example, messages 2, 4, and 6 are voice messages. The keyword used in the search is **voice_message**, the value of the **store.message.type.mtindex:2.flagname** option.

Migrating Mailboxes to a New System

Oracle Communications Messaging Server provides several ways to move or relocate mailboxes from one Messaging Server host (mailhost) to another.

Topics:

- [Tools Summary for Relocating Messaging Server Users to a New Mailhost](#)
- [Migrating Mailboxes from an x86 Host to a SPARC Host](#)
- [Moving Mailboxes to Another Messaging Server While Online](#)
- [To Move Mailboxes Using an IMAP Client](#)
- [To Move Mailboxes by Using the `imsimport` Command](#)

For More Information:

- [rehostuser](#)

Tools Summary for Relocating Messaging Server Users to a New Mailhost

The following tools are available for relocating users from one mailhost to another:

- **rehostuser**, Enables you to move a Messaging Server user's mail store from one mailhost to another. The **rehostuser** utility also disconnects any active session, locks the store to ensure atomicity of the move from the user's perspective (no loss of data, flag change, and so on), changes the user's LDAP entry, flushes LDAP caches as necessary, and causes any queued mail to be rerouted to the new store. See "[rehostuser](#)" for more information.
- **imsbackup** | **imsrestore**: Manually backs up the account from one host and restores on another. Can be combined with **ssh** and "piped" across the network. You can also back up to a file and then access that file from the other host through NFS, or move the file in between.
- **imsimport**: Imports messages from a UNIX `/var/mail` format file into the Messaging Server message store.

Notes:

- If both the source and destination mailhosts are installed with at least Messaging Server 7, use "[rehostuser](#)." The **rehostuser** utility solves issues present in the other solutions, such as preventing users from accessing mail while it is being moved, making sure mail is held in the MTA during the move and redirected to the new message store, and so on.
- You can manually run "[imsbackup](#)" on one host and "[imsrestore](#)" on the other host. You can combine multiple users in one step. You can combine "[imsbackup](#)" with

ssh or NFS, or you could use **gzip** and transfer the backup file from one host to the other by using **ftp**. This method is actually used within the **rehostuser** utility. However, **rehostuser** shields you from having to set the users' status, disconnect them, and so on. If you use **imsbackup** and **imsrestore** manually, you also need to deal with those details manually. But if you have a large amount of data to move and can afford for the MTA and IMAP user access to be down while it is being moved, this method might be more efficient.

- Similar to **imsrestore** and **imsbackup**, you can use "**imsimport**" to import UNIX **/var/mail** format files into the message store. If you are moving from a non-Oracle server and it does not have IMAP access capability, perhaps you can get it to export the folders in UNIX **/var/mail** format and then import them this way.

Migrating Mailboxes from an x86 Host to a SPARC Host

The message store data formats are architecture dependent. You cannot transfer any data components in "[Message Store Component Version Compatibilities](#)" from one architecture to another directly. To migrate the message store from an x86 host to a SPARC host (or from a SPARC host to an x86 host), use **imsbackup** and **imsrestore**, or **rehostuser**.

Moving Mailboxes to Another Messaging Server While Online

You can migrate the message store from an older version of Messaging Server to a newer version, or move mailboxes from one Messaging Server message store to another, while remaining online. This procedure works for iPlanet Messaging Server 5.0 and later. You cannot move messages from prior versions of Messaging Server or a non-Oracle Communications Suite message store. Moving mailboxes while online have the following advantages and disadvantages.

Advantages

- You can move the mailboxes from the old source system to the new destination system without user involvement.
- This process is faster than any of the other processes.
- Re-linking is not required if you are moving an entire partition.
- Both Messaging Server systems remain active and online.
- You can migrate all the mailboxes on a messages store or a subset of those messages. This procedure allows for incremental migrations.

Disadvantages

- This method does not work with non-Oracle Communications Suite messaging servers.
- The users being migrated do not have access to their mailboxes until the migration of their own mailbox is complete.
- This method can be complex and time consuming.

Incremental Mailbox Migration While Online

Incremental migration provides numerous advantages for safely and effectively moving your message store to a different system or upgrading to a new system,

incremental migration allows you to build a new back-end message store system alongside the old back-end message store. You can then test the new system, migrate a few friendly users, then test the new system again. Once you are comfortable with the new system and configuration, and you are comfortable with the migration procedure, you can start migrating real commercial users. These users can be split into discrete backup groups so that during migration, only members of this group are offline, and only for a short time.

Another advantage of on-line incremental migration is that you do not have to plan for a system-wide back out in case your upgrade fails. A back out is a procedure for reverting changes you have made to a system to return the system to the original working state. When doing a migration, you have to plan for failure, which means that for every step in the migration requires a plan to return your system back to its previous operational state.

The problem with offline migrations is that you can't be sure your migration is successful until you've completed all the migration steps and switched the service back on. If the system doesn't work and cannot be quickly fixed, you'll need a back out procedure for all the steps performed. This can be stressful and take some time, during which your users will remain offline.

With an online incremental migration you perform the following basic steps:

1. Build the new system alongside the old one so that both can operate independently.
2. Configure the old system for coexistence with the new.
3. Migrate a group of "friendly" users and test the new system and its coexistence with the old system.
4. Divide the users on the old system into groups and migrate group by group to the new one as desired.
5. Disassemble the old system.

Because both systems will coexist, you have time to test and get comfortable with the new system before migrating to it. If you do have to perform a back-out procedure, which should be very unlikely, you only have to plan for steps 2 and 4. Step 2 is easy to revert because you do not ever touch user's data. In step 4, the backout is to revert the user's state to active and their mailhost attribute back to the old host. No system-wide back out is required.

Online Migration Overview

Migrating mailboxes while remaining online is a straightforward process. Complications arise when you try to ensure that messages in transit to the mailbox (sitting in an MTA channel queue waiting for delivery) are not lost in the migration process. One solution is to hold messages sent during the migration process in a *held* state and wait for the messages in the various channel queues to be delivered. However, messages can get stuck in queues because of system problems or because a particular user is over quota. In this case, you must address this situation before migrating the mailboxes.

You can take various measures to reduce the likelihood of lost messages and to verify that messages are not stuck in a channel queue, but at a cost of increased complexity of the procedure.

The order and necessity of steps in the procedure vary depending upon your deployment and whether every message addressed to every mailbox must not be lost. This section describes the theory and concepts behind the steps. It is incumbent on you

to understand each step and decide which to take and in which order, given your specific deployment. Following is an overview of the process of moving mailboxes. This process might vary depending upon your deployment.

1. Block user access to the mailboxes being moved.
2. Temporarily hold messages addressed to the mailbox being moved.
3. Verify that messages are not stuck in the channel queues.
4. Change the user's mailhost attribute to the new mailbox location.
5. Move the mailboxes to the new location.
6. Release held mail to be delivered to the new mailbox and enable incoming messages to be delivered to the migrated mailboxes.
7. Examine the old message store to see if any messages were delivered after the migration.
8. Unblock user access to mailbox.

To Migrate User Mailboxes from One Messaging Server to Another While Online

The requirements for this type of migration are as follows:

- **stored** should be running on both the source (old) and destination (new) messaging servers.
- The source system and destination system must be able to route messages to each other if both systems will operate in co-existence. This is needed, for instance, so that delivery status notification messages can be generated on the destination system and get delivered to the source system.

Note: Some steps apply only if you are upgrading the messaging server from an earlier version to a later version. These steps might not apply if you are only migrating mailboxes from one message store to another. The steps that apply to migrating entire systems are noted.

1. On the source system, split your user entries to be moved into equal backup groups by using the **backup-groups.conf** file. This step is in preparation for the mailbox migration, Step 8, that occurs later in this procedure. See "[To Create Backup Groups](#)" for detailed instructions. You can also place the user names into files and use the **-u** option in the **imsbackup** command.
2. Notify users to be moved that they will not have access to their mailboxes until the move is completed. Ensure that users to be moved are logged out of their mail systems before the data move occurs. (See "[Monitoring User Access to the Message Store](#).")
3. Set the authentication cache timeout to 0 on the back-end message store and MMP systems, and **alias_entry_cache_timeout** option to 0 on the MTAs.

```
msconfig set mta.alias_entry_cache_timeout 0
```

- a. On the back-end message stores containing the mailboxes to be moved, set the authentication cache timeout to 0.

```
msconfig set base.authcachettl -0
```

This step and Step 7 (changing **mailUserStatus** to **hold**) immediately prevents users from accessing their mailboxes during migration.

- b. On all MMPs, set the LDAP and authentication cache timeout to 0. For the IMAP proxy and POP proxy, set both **ldapcachettl** and **authcachettl** to 0. For example:

```
msconfig
msconfig> set imaproxy.authcachettl 0
msconfig# set imaproxy.ldapcachettl 0
msconfig# set popproxy.authcachettl 0
msconfig# set popproxy.ldapcachettl 0
msconfig# write
```

- c. On any Messaging Server host that contains an MTA that inserts messages into mailboxes that are to be migrated, set the **alias_entry_cache_timeout** option to 0. Messaging Server hosts that run an MTA that inserts messages into the migrating mailboxes are typically the back-end message store. However, if the system is using LMTP, then that system is the inbound MTA. Check your configuration to make sure. Resetting the **alias_entry_cache_timeout** option forces the MTA to bypass the cache and look directly at the LDAP entry so that intermediate channel queues (for example, the **conversion** or **reprocess** channels) see the new **mailUserStatus (hold)** of the users being moved rather than the out-of-date cached information.
- d. Restart the systems on which the caches were reset. You must restart the system for these changes to take place. See ["Starting and Stopping Services"](#) for instructions.
4. Ensure that both your source Messaging Server and destination Messaging Server are up and running. The source Messaging Server must be able to route incoming messages to the new destination server.
5. Change the LDAP attribute **mailUserStatus** on all user entries whose mailboxes will be moved from **active** to **hold**. Changing the attribute holds incoming messages in the **hold** queue and prevents access to the mailboxes over IMAP, POP, and HTTP. Typically, users are moved in groups of users. If you are moving all the mailboxes of a single domain, you can use the **mailDomainStatus** attribute. For more information on **mailUserStatus**, see the **mailUserStatus attribute** in the *Schema Reference*.
6. Make sure that messages addressed to mailboxes being migrated are not stuck in the **ims-ms** or **tcp_lmtp*** channel queues (if LMTP has been deployed). Use the following commands to see if messages exist in the channel queue directory tree and in the *held* state (to see **.HELD** files) addressed to a user to be migrated:

```
imsimta qm directory -to=<user_address_to_be_migrated> -directory_tree
imsimta qm directory -to=<user_address_to_be_migrated> -held -directory_tree
```

If there are messages in the queue, run these same commands later to see if the MTA has dequeued them. If there are messages that are not being dequeued, then you must address this problem before migrating. This should be a rare occurrence, but possible causes are recipient mailboxes being over quota, mailboxes being locked perhaps because users are logged in and moving messages, the LMTP backend server is not responding, network or name server problems, and so on).

7. Change the LDAP attribute **mailHost** in the user entries to be moved as well as in any mail group entries. Use the **ldapmodify** command to change the entries to the new mail server. Use the **ldapmodify** that comes with Messaging or Directory Server. Do not use the Oracle Solaris **ldapmodify** command.
- You only need to change the **mailHost** attribute in the mail group entry if the old mail host is being shut down. You can either change this attribute to the

new mail host name or just eliminate the attribute altogether. It is optional for mail groups to have a **mailHost**. Having a **mailHost** means that only that host can do the group expansion. Omitting a **mailHost** (which is the more common case) means all MTAs can do the group expansion. Mail group entries do not have mailboxes to be migrated and typically do not even have the **mailhost** attribute. For more information on **mailhost**, see the **mailHost attribute** in the *Schema Reference*.

8. Move the mailbox data from the source Messaging Server message store to the destination Messaging Server message store and record the time when started. Back up the mailboxes with the **imsbackup** utility and restore them to the new Messaging Server with the **imsrestore** utility. For example, to migrate mailboxes from a Messaging Server system called **oldmail.example.com** to **newmail.example.com**, run the following command on **oldmail.example.com**:

```
<server-root>/bin/msg/store/bin/imsbackup -f- instance/group | rsh
newmail.example.com /opt/SUNWmsgsr/lib/msg/imsrestore.sh -f- -c y -v 1
```

You can run multiple concurrent "**imsbackup**" and "**imsrestore**" sessions (one per group) to maximize the transfer rate into the new message store. See also "[Backing Up and Restoring the Message Store.](#)"

Note: When **imsrestore** or any processing intensive operation takes significantly more system resources than normal, and continues doing so longer than the **msprobe** interval, there may be a temporary backlog of DB transaction log files to be cleared. If there are more files than specified in **store.maxlog**, then **msprobe** may erroneously restart all the processes during a restore. To prevent this from happening, disable **msprobe** during the **imsrestore**.

Note: Record the timestamp of when **imsbackup** is run for later delivery validation.

9. (*Conditional Step for System Upgrades*) If your mailbox migration is part of the process of upgrading from an earlier version of Messaging Server to the current version, set this current version of Messaging Server to be the new default Messaging Server for the system. Change the DNS **A** record of **oldmail.example.com** to point to **newmail.example.com** (the server responsible for domain(s) previously hosted on **oldmail.example.com**).
10. Enable user access to the new message store. Set the LDAP attribute **mailUserStatus** or **mailDomainStatus**, if applicable, to whatever value it had been before it was changed to **hold** (for example, **active**).
11. Release the messages in the *held* state on all source Messaging Servers. Any system that may be holding incoming messages needs to run the following command to release all the user messages:

```
imsimta qm release -channel=hold -scope
```

where *scope* can be **all**, which releases all messages; **user**, which is the user ID; or **domain** which is the domain where the user resides.

12. Reset the authentication cache timeout and the **alias_entry_cache_timeout** option to the default or desired values and restart the system. At this point, you've migrated all the user mailboxes that need to be migrated. Before proceeding, make

sure that no new entries in LDAP have been created with the old system as the **mailhost**, and if some have, migrate them. Also, make sure that no such entries can be created by modifying the provisioning systems. You also want to change the **preferredmailhost** attribute to the name of the new mail host. For back-end messages stores, set authentication cache timeout to 900 as follows:

```
msconfig set base.authcachettl 900
```

For the IMAP proxy and POP proxy, set both **ldapcachettl** and **authcachettl** to 900. For example:

```
./msconfig
msconfig> set imaproxy.authcachettl 900
msconfig# set imaproxy.ldapcachettl 900
msconfig# set popproxy.authcachettl 900
msconfig# set popproxy.ldapcachettl 900
msconfig# write
```

For MTAs, set the **alias_entry_cache_timeout** option to 600.

```
msconfig set mta.alias_entry_cache_timeout 600
```

You must restart the system for these changes to take place. See "[Starting and Stopping Services](#)" for instructions.

13. Ensure that the user clients are pointing to the new mail server. After the upgrade finishes, have the users point to the new server through their mail client program (in this example, users would point to **newmail.example.com** from **oldmail.example.com**). An alternative is to use a Messaging Multiplexor (MMP), which obviates the need to have users point their clients directly to the new mail server. The MMP gets that information from the **mailHost** attribute that is stored in the LDAP user entries and automatically redirects the client to the new server.
14. After everything works, verify that no messages were delivered to the old message store after the migration. Go to the old message store and run **mboxutil -l** to list the mailboxes. Check the last message delivery timestamp. If a message was delivered after the migration timestamp (the date stamp when you ran the **imsbackup** command), then migrate those messages with a backup and restore command. Because of the preparatory steps provided, it would be exceedingly rare to see a message delivered after migration. Theoretically, a message could be stuck in a queue for the number of days or hours specified by the **notices** channel options. See the discussion on setting notification message delivery intervals in the *Messaging Server Reference*).
15. Remove duplicate messages on the new message store, run the **relinker** command. This command might free disk space on the new message store.
16. Remove the old messages from the store you migrated from and delete users from the database on the old store. Run the **mboxutil -d** command. (See "[mboxutil](#).")

To Move Mailboxes Using an IMAP Client

This procedure can be used anytime messages need to be migrated from one messaging server to a different messaging server. Consider the advantages and disadvantages before moving mailboxes using this method.

Advantages

- This method can be used to migrate from a non-Oracle Communications host to the Messaging Server host. It can also be used to move mailboxes from one physical server to a different physical server.
- After you set up the new mail server or message store, responsibility for moving mailboxes to the new system is left to users.
- The process for moving mailboxes is relatively simple.
- User access to mailboxes does not have to be disabled.

Disadvantages

- Requires that both the old and new systems be simultaneously running and accessible to users.
 - Cumulatively, this method takes longer to move mailboxes than other methods.
 - Responsibility for moving mailboxes to the new system is left to users.
 - The size of the new message store will be significantly larger than the old message store until the re-linking operation is performed.
1. Install and configure the new Messaging Server.
 2. Set `store.relinker.enable` to 1. This reduces the message store size on the new system caused by duplicate storage of identical messages.
 3. Provision users on the new Messaging Server. You can use Delegated Administrator to do this. As soon as users are provisioned on the new system, newly arriving mail is delivered to the new INBOX.
 4. Have users configure their mail client to view both new and old Messaging Server mailboxes. This may involve setting up a new email account on the client. See mail client documentation for details.
 5. Instruct users to drag folders from their old Messaging Server to their new Messaging Server.
 6. Verify with users that all mailboxes are migrated to the new system, then shut down the user account on the old system.

To Move Mailboxes by Using the `imsimport` Command

This procedure is specifically used to move mailboxes from UNIX `/var/mail` format folders into a Messaging Server message store. However, if the Messaging Server host from which you are migrating can convert the IMAP message stores to UNIX `/var/mail` format, then you can use the `imsimport` command to migrate messages to Messaging Server. Consider the advantages and disadvantages before moving mailboxes using this method.

Advantages

- You have complete responsibility for moving mailboxes from the old system to the new system. Users do not have to do anything.

Disadvantages

- This method takes longer to move mailboxes than the other non-IMAP methods.
- Users access to mailboxes must be disabled while mailboxes are being moved.
- The size of the new message store will be significantly larger than the old message store until the re-linking operation is performed.

1. Install and configure the new Messaging Server.
2. Set **`store.relinker.enabled`** to **1**. This reduces the message store size on the new system caused by duplicate storage of identical messages.
3. Provision users on the new Messaging Server if needed. You can use Delegated Administrator to do this. Do not switch over to the new system yet.
4. Disable user access to both the new and old messaging store. Set the **`mailUserStatus`** LDAP attribute to **hold**. User's mail is sent to the hold queue and access to the mailbox over IMAP, POP, and HTTP is disallowed. MTA and Message Access Servers on the store server must comply with this requirement. This setting overrides any other **`mailDeliveryOption`** settings.
5. If the mail store from the existing mail server is not already in the **`/var/mail`** format, convert the mail store to **`/var/mail`** files. Refer to the third-party mail server documentation.
6. Run the **`imsimport`** command. For example:

```
imsimport -s /var/mail/joe -d INBOX -u joe
```

See the "**`imsimport`**" for details.
7. Enable user access to the message store. Set the **`mailUserStatus`** LDAP attribute to **active**.
8. Enable user access to the new messaging store.
9. Shut down the old system.

Monitoring Disk Space

This chapter describes configuration options for monitoring disk and partition usage and for generating warnings about disk space availability.

Topics:

- [Overview](#)
- [Symptoms of Insufficient Disk Space](#)
- [Monitoring Disk Space](#)
- [Monitoring the Message Store](#)

Overview

Inadequate disk space or inadequate space within a disk partition are among the most common causes of mail server problems and failure. Typical causes of inadequate space are:

- Message store quotas are not enforced and the message store outgrows the disk space available for a partition.
- Over-long MTA message queues.
- Log files that are not adequately monitored and kept within defined limits. (Note that there are a number of log files such as LDAP, MTA, and Message Access, and that each of these log files can be stored on different disks.)

Symptoms of Insufficient Disk Space

Symptoms of insufficient disk-space are:

- MTA queues overflow and reject SMTP connections.
- Messages remain in the **ims-master** queue and are not delivered to the message store.
- Log files overflow.

If a message store partition fills up, message access daemons can fail and message store data can be corrupted. Message store maintenance utilities such as **imexpire** and **reconstruct** can repair the damage and reduce disk usage. However, these utilities require additional disk space, and repairing a partition that has filled an entire disk can cause down time.

Monitoring Disk Space

Depending upon the system configuration you may need to monitor various disks and partitions. For example, MTA queues may reside on one disk/partition, message stores may reside on another, and log files may reside on yet another. Each of these spaces will require monitoring and the methods to monitor these spaces may differ.

Oracle Communications Messaging Server provides specific methods for monitoring message store disk usage and preventing partitions from filling up all available disk space.

You can take the following steps to monitor the message store's use of disk space:

- Set options to monitor message store disk usage
- Lock message store partitions when a disk-usage threshold is reached

Monitoring the Message Store

You can monitor message store disk usage by configuring the following attributes with the **msconfig** utility:

- **alarm.system:<alarmtype>.statinterval** Specifies the length of time, in seconds, between disk availability checks. For example, to set the system to monitor disk space every 600 seconds, enter the following command:

```
msconfig set alarm.system:diskavail.statinterval 600
```

- **alarm.system:<alarmtype>.threshold** Specifies a percentage of disk space that must be available or a warning is generated. For example, it is recommended that disk-space usage should not exceed 75%; correspondingly, the following command generates a warning whenever the amount of disk space available falls below 25%:

```
msconfig set alarm.system:diskavail.threshold 25
```

- **alarm.system:<alarmtype>.warninginterval** Specifies an interval, in hours, between the repetition of disk availability alarms. For example, the following command sets an interval of one hour between one disk availability warning and another.

```
msconfig set alarm.system:diskavail.warninginterval 1
```

- **alarm.system:<alarmtype>.description** A description of the disk availability alarm. For example, the following sets the description of an availability alarm for a message-queue alarm:

```
msconfig set alarm.system:diskavail.description "Percentage message-queue partition disk space available"
```

Monitoring Message Store Partitions

By default, partition monitoring is in effect so that when a message-store partition uses more than a specified percentage of available disk space, the partition is locked and any incoming messages are held in the MTA message queue. Two **msconfig** options control partition monitoring:

- **checkdiskusage**

checkdiskusage enables partition monitoring. It takes a boolean value; the default is 1 (monitoring is enabled).

- **diskusagethreshold**

diskusagethreshold specifies a disk-usage threshold beyond which the partition is locked. It takes an integer value from 1 to 99; the default value is 99.

As a partition approaches the threshold specified in **diskusagethreshold**, the message-store daemon checks the partition with increasing frequency, ranging from once every 100 minutes to once every minute. If disk usage goes higher than the threshold specified in **diskusagethreshold**, the message-store daemon:

- Locks the partition; incoming messages are held in the MTA message queue and are not delivered to mailboxes in the message store partition until it is unlocked.
- Logs a message to the default log file.
- Sends an email notification to the postmaster. (You can change the recipient of the notification by setting the **msconfigalarm.noticercpt** option.)

In setting the **diskusagethreshold** option, specify a usage percentage that is low enough to allow time for repartitioning or assigning more disk space to the local message store. For example, if a partition fills up disk space at a rate of 2 percent per hour and it takes an hour to allocate additional disk space, set the disk-usage threshold to a value lower than 98 percent.

Protecting Mailboxes from Deletion or Renaming

You might want to protect some mailboxes from deletion or modification except by the administrator. The following procedure describes how to do this.

If someone other than an administrator attempts to delete, modify, or rename a protected mailbox, the error message "**mailbox is pinned**" is displayed.

- Set the **store.pin** configuration option by using the following format:

```
msconfig set store.pin <mailbox1>%<mailbox2>%<mailbox3>
```

where *mailbox1*, *mailbox2*, and *mailbox 3* are the mailboxes to be protected (you can use spaces in mailbox names), and @ is the separator between each mailbox.

In this example above, the mailboxes specified in **<mailbox1>@<mailbox2>@<mailbox 3>** are not per user mailboxes, as in **user/user1/Drafts**. Instead, the mailboxes specified are for all users using a certain directory, such as **Drafts**. An administrator can therefore prevent users from renaming or deleting the **Drafts** or **Backup** folder for all users by running:

```
msconfig set store.pin Drafts%Backup
```

Reducing Message Store Size Due to Duplicate Storage

When a message is sent to multiple recipients, that message is placed in each recipient's mailbox. Some messaging systems store separate copies of the same message in each recipient's mailbox. By contrast, the Oracle Communications Messaging Server strives to retain a single copy of a message regardless of the number of mailboxes in which that message resides. It does this by creating hard links to that message in the mailboxes containing that message.

When other messaging systems are migrated to the Messaging Server, these multiple message copies may be copied over with the migration process. With a large message store, this means that a lot of messages are duplicated unnecessarily. In addition, multiple copies of the same message can be accumulated in normal server operation, for example, from IMAP append operations or other sources.

Messaging Server provides a command called **relinker** that removes the excess message copies and replaces them with hard links to a single copy.

Note: The **relinker** feature is intended to repair the situation where the normal single-copy nature of the message store has become broken for some reason. You should only need to use the relinker if you have done something which could have caused duplicate messages to become individual copies instead of the normal single-copy. This feature is not the normal way the store normally accomplishes its single-copy feature. You should not need to keep the real-time relinker feature enabled for long periods of time. You should not need to use the relinker command on an ongoing basis. You should only need this feature if you have done (or will soon be doing) something which would break the single-copy feature of the store. For more information about single-copy, see the discussion in "[How the Message Store Works](#)."

Topics:

- [Relinker Theory of Operations](#)
- [Using Relinker in the Realtime Mode](#)
- [Configuring Relinker](#)

Relinker Theory of Operations

The relinking function can be run in the command or realtime mode. When the **relinker** command is run, it scans through the message store partitions, creates or updates the MD5 message digest repository (as hard links), deletes excess message files, and creates the necessary hard links.

The digest repository consists of hard links to the messages in the message store. It is stored in the directory hierarchy `partition_path/=md5`. This directory is parallel to the user mailbox hierarchy `_partition_path/=user` (see "Message Store Directory Layout"). Messages in the digest repository are uniquely identified by their MD5 digest. For example, if the digest for `fredb/00/1.msg` is `4F92E5673E091B43415FFFA05D2E47EA`, then `partition/=user/hashdir/hashdir/=fredb/00/1.msg` is linked to `partition/=md5/_hashdir/_hashdir/4F92E5673E091B43415FFFA05D2E47EA.msg`. If another mailbox has this same message, for example, `partition_path/=user/hashdir/hashdir/gregk/00/17.msg`, that message will also be hard linked to `partition_path/=md5/4F/92/4F92E5673E091B43415FFFA05D2E47EA.msg`. This is shown in Figure 69-1.

Figure 69-1 Message Store Digest Repository



For this message, the link count will be three. If both messages are deleted from the mailboxes of fredb and gregk, then the link count will be one and the message can be purged.

The **relinker** process can also be run in the realtime mode for similar functionality. See "Using Relinker in the Realtime Mode" for details.

Using relinker in the Command Line Mode

relinker scans through a message store partition, creates or updates the MD5 message repository (as hard links) and deletes excess message files. After **relinker** scans a store partition, it outputs statistics on the number of unique messages and size of the partition before and after relinking. To run more quickly on an already hashed store, **relinker** only computes the digest of the messages not yet present in `=md5`. It also has

an option to erase the entire digest repository (which doesn't affect the user mailboxes).

The syntax for the command is as follows:

```
relinker [-ppartitionname] [-d]
```

where *partitionname* specifies the partition to be processed (default: all partitions) and **-d** specifies that the digest repository be deleted. Sample output is shown below:

```
relinker
Processing partition: primary
Scanning digest repository...
Processing user directories.....
-----
Partition statistics Before After
-----
Total messages 4531898 4531898
Unique messages 4327531 3847029
Message digests in repository 0 3847029
Space used 99210Mb 90481Mb
Space savings from single-copy 3911Mb 12640Mb
-----
relinker -d
Processing partition: primary
Purging digest repository...
-----
Partition statistics Before After
-----
Message digests in repository 3847029 0
-----
```

relinker can take a long time to run, especially for the first time if there are no messages are in the repository. This is because it has to calculate the digest for every message (if the **relinker** criteria is configured to include all messages-see "[Configuring Relinker](#)" for information on configuring **relinker** criteria.) For example, it could take six hours to process a 100 Gigabyte message store. However, if run-time relinking is enabled see "[Using Relinker in the Realtime Mode](#)."

If the **relinker** command line mode is used exclusively, and not the run-time option, it is necessary to purge the digest repository (**=md5**), otherwise messages purged in the store (**=user**) will not become available disk space since they still have a link in the digest repository (they become orphaned). If you are just performing a one-time optimization of the store-for example after a migration-you can run **relinker** once, then delete the entire repository with **relinker -d**. For repeated purging during migration, it is sufficient to just run the **relinker** command repeatedly, since each time it runs it also purges the expired or orphaned messages from the repository.

It is safe to run multiple instances of **relinker** in parallel with each processing a different partition (using the **-p** option). Messages are only relinked inside the same partition.

Using Relinker in the Realtime Mode

The **relinker** function can be enabled in the realtime mode by setting the **msconfig** option **store.relinker.enable** to **1**. Using **relinker** in the realtime mode will compute the digest of every message delivered (or restored, IMAP appended, and so forth) which matches the configured **relinker** criteria ("[Configuring Relinker](#)"), then look in the repository to see if that digest is already present. If the digest is present, it creates a link to it in the destination mailbox instead of creating a new copy of the message. If

there is no digest, it creates the message and adds a link to it in the repository afterwards.

stored scans the digest repositories of each partition and purges the messages having a link count of 1, or which don't match the **relinker** criteria. The scan is done one directory at a time over a configurable time period. This is so that the I/O load is evenly distributed and does not noticeably impact other server operations. By default the purge cycle is 24 hours, which means messages can still be present on the disk for up to 24 hours after they have been deleted from the store or have exceeded the configured maximum age. This task is enabled when the **relinker** realtime mode is enabled.

Configuring Relinker

Table 69–1 shows the options used to set **relinker** criteria.

Table 69–1 *relinker msconfig options*

Option	Description
store.relinker.enable	Enables real-time relinking of messages in the append code and stored purge. The relinker command-line tool may be run even if this option is off. However since stored will not purge the repository, relinker -d must be used for this task. Turning this option on affects message delivery performance in exchange for the disk space savings. Default: 0
store.relinker.maxage	Maximum age in hours for messages to be kept in the repository, or considered by the relinker command-line. -1 means no age limit, that is, only purge orphaned messages from the repository. For relinker it means process existing messages regardless of age. Shorter values keep the repository smaller thus allow relinker or stored purge to run faster and reclaim disk space faster, while longer values allow duplicate message relinking over a longer period of time, for example, when users copy the same message to the store several days apart, or when running a migration over several days or weeks. Default: 24
store.relinker.minsize	Minimum size in kilobytes for messages to be considered by run-time or command-line relinker . Setting a non-zero value gives up the relinker benefits for smaller messages in exchange for a smaller repository. Default: 0
store.relinker.purgecycle	Approximate duration in hours of an entire stored purge cycle. The actual duration depends on the time it takes to scan each directory in the repository. Smaller values will use more I/O and larger values will not reclaim disk space as fast. 0 means run purge continuously without any pause between directories. -1 means don't run purge in stored (then purge must be performed using the relinker -d command). Default: 24

Specifying Administrator Access to the Message Store

This information describes how to grant store privileges to the message store for your Oracle Communications Messaging Server installation. See ["Managing Message Store Partitions and Adding Storage"](#) for conceptual information.

Topics:

- [Overview](#)
- [Adding an Administrator Entry](#)
- [Modifying or Deleting an Administrator Entry](#)

Overview

Message store administrators can view and monitor user mailboxes and specify access control for the message store. Store administrators have proxy authentication privileges to any service (POP, IMAP, HTTP, or SMTP), which means they can authenticate to any service using the privileges of any user. These privileges allow store administrators to run certain utilities for managing the store.

Note: Other users might also have administrator privileges to the store. For example, some administrators may have these privileges.

See also ["Protecting Mailboxes from Deletion or Renaming."](#)

Adding an Administrator Entry

To add an administrator entry at the command line, enter:

```
msconfig set store.admins <adminlist>
```

where *adminlist* is a space-separated list of administrator IDs. If you specify more than one administrator, you must enclose the list in quotes. In addition, the administrator must be a member of the Service Administrator Group, in the LDAP user entry: **memberOf: cn=Service Administrators,ou=Groups,o=usergroup**. You must restart **imapd** for the system to recognize the change in **store.admins**.

Modifying or Deleting an Administrator Entry

To modify or delete an existing entry in the message store Administrator UID list at the command line, use the same command:

```
msconfig set store.admins <adminlist>
```

where *adminlist* is a space-separated list of administrator IDs who should be included in the modified list. If you specify more than one administrator, you must enclose the list in quotes. In addition, the administrator must be a member of the Service Administrator Group, in the LDAP user entry: **memberOf: cn=Service Administrators,ou=Groups,o=usergroup**. You can delete members from the list, but the modified list must contain at least one administrator ID.

You must restart **imapd** for the system to recognize the change in **store.admins**.

Valid Message Store UIDs and Folder Names

This information describes valid constructions for message store UIDs and folder names. Note that folder and mailbox are used synonymously.

Topics:

- [Message Store User ID](#)
- [Message Store Mailbox Name for Commands](#)
- [Valid UIDs](#)

Message Store User ID

The message store user ID is a mail user's unique identifier in the message store. In the default domain, this is the same as the user's LDAP **uid** attribute. In hosted domains, this is *uid@domain* where *uid* is the **uid** LDAP attribute and *domain* is the canonical domain name.

Message Store Mailbox Name for Commands

Some message store commands require that you specify a mailbox name. The required form of the name is **user/userid/mailbox** where *userid* is the message store user ID (see "Message Store User ID") and **mailbox** is a user's mailbox. Specifying INBOX sometimes implies all the user's mailboxes in the message store. For example, the following command removes the INBOX and all the folders of user **joe**.

```
mboxutil -d user/joe/INBOX
```

Note that in the context of message stores, folders and mailboxes are synonymous.

Valid UIDs

Valid and invalid **UID** characters are controlled separately by the MTA and message store mechanisms. That means **UID** character limitations are specified by the union of MTA and message store limitations. The following characters and strings are invalid as **UIDs** in the message store:

- `% ? * & / : \`
- ASCII values less than 20 or greater than 7E hexadecimal (see `man ascii`)
- A leading '-' is prohibited because it is reserved for negative rights
- A leading 'group=' is prohibited because it is reserved for group IDs

- The following UIDs are reserved: 'anonymous' 'anybody' 'anyone' and 'anyone@domain'
- The maximum supported length for a UID is 127 bytes

The following characters are invalid in *UIDs* in the MTA:

<space> \$ ~ = # * + % ! @ , { } () / \ < > ; : " ` [] & ?

The list of characters forbidden by the MTA can be modified by setting the **ldap_uid_invalid_chars** option with a string of the forbidden characters using decimal ASCII values, however, you are strongly advised not to change the default constraint. The default setting is as follows and reflect the characters listed above:

```
ldap_uid_invalid_  
chars=32,33,34,35,36,37,38,40,41,42,43,44,47,58,59,60,61,62,63,64,91,92,93,96,123,  
125,126
```

Valid mail folder names. The following characters are invalid as folder names:

% * ? and ASCII values less than 20 or greater than 7E hexadecimal (see **man ascii**).

In addition, folder names must be valid UTF-7 sequences.

Message Store Automatic Failover with Database Replication

This chapter includes the following topics:

- Overview of Message Store Database Replication
- Configuration Options
- Configuring Message Store Database Replication
- Message Store Automatic Failover

Overview of Message Store Database Replication

Berkeley Database provides support for building high availability (HA) applications based on replication. Message store database replication uses Berkeley Database HA facilities and low cost NFS storage devices to build an HA message store.

The **mboxlist** database is a transactional database. Database changes are written to the transaction logs. You can replicate the database by transporting the transaction log records from one site to another. Berkeley Database HA architecture supports single writer (master) and multiple reader replication. You must perform all database updates on the master. Replicas are available for read only activity. When the master fails, an election takes place, and one of the replicas will take over as master.

Either all replicas are on SPARC CPUs or all replicas are on x86 CPUs. Replication between these two types of CPUs is not supported by Oracle Communications Messaging Server.

The database replication and message store failover feature requires that replicas be on the same platform. Either all replicas are SPARC or all replicas are x86_64. It is also best practice for all replicas to be homogeneous (same hardware/CPU/OS), unless you are in the middle of a rolling hardware upgrade.

A message store replication group consists of one or more message store nodes. The message store nodes typically run on different physical hosts. The **mboxlist** database is replicated on every node. You can store the database locally. You store the mailbox partitions on remote storage devices running NFS servers. The NFS file systems are always mounted on all of the nodes in a replication group. You can configure the message store nodes with one or more remote hosts. If remote hosts are configured, the message store contacts a remote host to retrieve the replication group data on start up. If a master has not been established in a group, an election is called. A priority value is assigned to a node. When an election is held, the node with the most up-to-date log record and the highest priority becomes the new master. A node with priority 0 cannot be elected.

When the master fails, the replicas will automatically hold an election to select a new master. You are responsible for monitoring the master. The automatic failover facility will redirect incoming connections to a new master. The message store replicas run in read only mode. Any attempt to modify a mailbox on a replica returns an error. You can perform read only operations such as **mboxutil -l** and **imsbackup** on the replicas. You can install Multiple message store nodes on the same host in different zones. Each message store node must have a unique **base.listenaddr**.

Berkeley Database maintains an internal database to keep track of the replication group data. Starting the first node in a replication group for the first time initializes the database. To start up the first node the first time, run **start-msg -m**. Database transaction commits blocks until it either receives enough acknowledgments, or the acknowledgment timeout expires. You can configure the acknowledgment policy and timeout.

A two-site replication group is particularly vulnerable to duplicate masters when there is a disruption to communications between the sites. Two-site replication is disabled by default. You can enable it with the **store.dbreplicate.twosites** option. When this option is disabled, the message store cannot take over as master if the original master fails in a replication group with only two sites. In the event this happens, the message store cluster will be unavailable for write access.

When a replication-aware client, such as the MMP (POP or IMAP), LMTP client or ENS subscriber connects to a replication group, it has to be aware of the hosts in the replication group so if one is down it can attempt to connect to the next one. This is presently accomplished by setting the **proxy:repgroup.storehostlist** option to the same value on all replication-aware clients and back-ends. These clients also remember the last host to which they connected in the replication group so it's not necessary to failover on every connection. In addition, replication-aware clients that perform write operations need to connect to the replication group master. This is accomplished by having the back-end servers refer the client to the replication group master. Details about how each replication-aware client works are as follows:

- **IMAP server**

Advertises OK login referrals (RFC 2221) that point to the master. Third-party use of such referrals is supported. Note that login referrals are also provided when a user logs in into an IMAP server that does not contain their **INBOX** (as may be necessary to access shared folders). Login referrals of the latter form will point to a host that is not in the **storehostlist** for the replication group.

- **POP server**

Advertises referrals with a **SYS/REFER/hostname** extended error on login. This is an Oracle extension to the protocol based on RFC 3206. Third-party use of such referrals is supported.

- **LMTP server**

The referral is indicated via a private protocol extension in the LMTP greeting. Oracle does not support third-party use of our LMTP server.

- **ENS server**

No changes have been made.

- **Message Store ENS publisher**

Publishes to all available hosts in the replication group for the current store by default. The recommended deployment for ENS with store failover is to have **enpd** running on every message store master and replica.

- **ENS and JMQ publishers for Message Store**

The **hostname** attribute will use the replication group name instead of the local host name.
- **ENS C Client API**

If **storehostlist** is configured, it will perform failover and cache the last successful host in the host list.
- **MMP IMAP client**

Performs failover, caches the last successful master, and follows referrals.
- **MMP POP client**

Performs failover, caches the last successful master, and follows referrals.
- **MTA LMTP client**

When the **affinitylist** channel option is set, this performs failover, caches the last successful master, and follows referrals.
- **MTA BURL IMAP client**

Performs failover and caches the last successful host. As this is read-only, it does not follow referrals.
- **Shared folder IMAP client in imapd**

Performs failover, caches the last successful master, and follows referrals.
- **mshttpd IMAP client**

Performs failover, caches the last successful master, and follows referrals.
- **Glassfish MQ (aka JMQ)**

No changes other than the **hostname** attribute mentioned above. Support for Glassfish MQ is deprecated.

Configuration Options

Topics in this section include:

- [Configuration Options](#)
- [Command-line Utilities](#)

Configuration Options

Table 72–1 shows the configuration options (only supported in Unified Configuration), their descriptions, data types, and defaults.

Table 72-1 Configuration Options

Option	Description	Data Type	Default
<code>store.dbreplicate.enable</code>	Enable database replication	boolean	0
<code>store.dbreplicate.port</code>	Replication port number	integer	55000 Note that <code>storehostlist</code> does not support non-default port numbers.
<code>store.dbreplicate.dbremotehost</code>	Remote host name list; this option is deprecated in 8.0.1.	host[:port] [host[:port]]...	In 8.0.1 and later this defaults the value of the <code>proxy:repgroup.storehostlist</code> with the current host omitted from the list.
<code>store.dbreplicate.dbpriority</code>	Host priority	integer	100
<code>store.dbreplicate.ackpolicy</code>	Replication acknowledgment policy	0=none, 1=one, 2=one peer, 3=quorum, 4=all peer, 5=all available, 6=all clients	3
<code>store.dbreplicate.acktimeout</code>	Replication acknowledgment timeout	number of seconds	1 second
<code>store.dbreplicate.twosites</code>	Enable two sites replication group	boolean	0
<code>proxy:repgroup.storehostlist</code>	Sets list of hosts in each store replication group for all relevant servers in the deployment. The preferred master should be listed first.	host[host]	Value of LDAP <code>mailHost</code> attribute for a user.
<code>proxy:repgroup.imapport</code>	The port used to connect to IMAP for this replication group.	unsigned 16-bit integer	Value of <code>base.proxyimapport</code> option.
<code>proxy:repgroup.imapadmin</code>	The administrative user name used when connecting to this replication group.	non-empty-string	Value of <code>base.proxyadmin</code> option.
<code>proxy:repgroup.imapadminpass</code>	The administrative password used to connect to this replication group.	password	Value of <code>base.proxyadminpass</code> option.

Command-line Utilities

- The `imcheck` subsystem option prints the database replication statistics. For additional information see "`imcheck`."

To print the database replication statistics, run the following command:

```
imcheck -s rep
```

- The `start-msg -m` option starts the message store as a replication master. For additional information see "`start-msg`."

To start the message store as a replication master, run the following command:

```
start-msg -m
```

- The **stored -d site** option removes a replication site from the replication database. For additional information, see "stored."

To delete an **mboxlist** replication site called **grumpy** from the cluster, run the following command:

```
stored -d grumpy
```

Configuring Message Store Database Replication

The following configuration examples show you how to configure message store database replication.

To Configure a Three Node Cluster for HA

The following example sets up a cluster with three electable nodes (huey, dewey and louie at **example.com**). This example assumes LMTP has been configured on these back-ends. The message store partition is on a shared storage mounted at **/zfssa/primary**.

On **huey.example.com**:

```
msconfig set store.dbrePLICATE.enable 1
msconfig set proxy:cluster1.storehostlist "huey.example.com dewey.example.com
louie.example.com"
msconfig set partition:primary.path /zfssa/primary
msconfig set task:snapshot.enable 0
msconfig set task:snapshotverify.enable 0
start-msg -m
```

On **dewey.example.com**:

```
msconfig set store.dbrePLICATE.enable 1
msconfig set proxy:cluster1.storehostlist "huey.example.com dewey.example.com
louie.example.com"
msconfig set partition:primary.path /zfssa/primary
msconfig set task:snapshot.enable 0
msconfig set task:snapshotverify.enable 0
start-msg
```

On **louie.example.com**:

```
msconfig set store.dbrePLICATE.enable 1
msconfig set proxy:cluster1.storehostlist "huey.example.com dewey.example.com
louie.example.com"
msconfig set partition:primary.path /zfssa/primary
msconfig set task:snapshot.enable 0
msconfig set task:snapshotverify.enable 0
start-msg
```

In addition, the **storehostlist** has to be set on all front-end servers as well, and the LMTP client has to be configured on the front-ends with the **affinitylist** channel option. Due to the complexity of setting up LMTP, it is recommended to copy the example recipe file **LMTPBackendFailover.rcp** and modify it with appropriate settings for the back-end stores.

To Change the DB Replication Local Instance Port

The replication group info is maintained by BDB. **store.dbrePLICATE.port** is for the local site only.

To change the port number on one node with the cluster running:

1. Stop the Message Store.

```
stop-msg store
```

2. Remove the local site from the group.

```
stored -d
```

3. Change the port number on the local site in Unified Configuration.

```
msconfig set store.dbreplicate.port <newport>
```

Or in legacy configuration.

```
configutil -o store.dbreplicate.port -v <newport>
```

4. Restart the Message Store

```
start-msg
```

5. Run **imcheck -s rep** on the other sites. You should see the new port for this site.

```
imcheck -s rep
```

Message Store Automatic Failover

This section describes Messaging Server's Message Store automatic failover feature and its configuration. It contains the following sections:

- [Basic Requirements](#)
- [Overview of Message Store Automatic Failover](#)
- [Configuring Message Store Automatic Failover](#)

Basic Requirements

In order to use Message Store automatic failover, the following is necessary.

- Messaging Server must be running in Unified Configuration mode.
- Messaging Server must be deployed using LMTP.

Overview of Message Store Automatic Failover

The Message Store automatic failover feature is useful for customers who already have 24/7 operators in their machine room but do not want the extra deployment and configuration expense of Oracle Solaris Cluster and do not want the extra cost and vendor interaction of Veritas Cluster.

The basic model is that all Messaging Server hosts in the deployment need to be manually configured with an ordered list of hostnames for each mailStore. Each hostname corresponds to a separate product installation, but in a given mailStore list all the hosts must use a shared disk (e.g., NFS, filer) for the product data, but have a configuration that is largely identical except for the "base.hostname" setting. The first host in the list is the primary host for that mailStore. The primary host is running and the secondary hosts are not running (on standby).

The MMP, LMTP client, and imapd shared folder support will now automatically use the secondary host for requests related to the primary mailHost; there is no need to refresh or restart these services for this to happen.

Note that a standby hostname does not mean it is necessary to have unused hardware. If multiple IP addresses are used on the same server, then it can support multiple installations. However, if the primary host is dedicated and the secondary host is shared for a mailStore, then the service response time will be reduced when automatic failover happens. Sites need to consider how much spare capacity is needed to service users when there is a hardware outage.

To enable safety during rolling version upgrades where a node is deliberately shutdown during upgrade, we recommend having at least 3 hostnames associated with each mailHost.

Configuring Message Store Automatic Failover

For this example, we assume the following hosts in a deployment:

LDAP mailStore "store1.example.com"

store1a.example.com primary

store1b.example.com

store1c.example.com

LDAP mailStore "store2.example.com"

store2a.example.com primary

store2b.example.com

store2c.example.com

mta.example.com - an MTA configured to use LMTP

mmp.example.com - an MMP

We recommend that customers use Unified Configuration recipes to set up automatic failover to avoid typographical errors when configuring multiple machines by running a recipe similar to the following on all machines in the deployment.

```
set_option("proxy:store1\\.example\\.\\.com.storehostlist",
"store1a.example.com store1b.example.com store1c.example.com");
set_option("proxy:store2\\.example\\.\\.com.storehostlist",
"store2a.example.com store2b.example.com store2c.example.com");
```

Save the recipe above to a plain text file, for example, automatic-failover.rcp

Some extra configuration need be added to LMTP server and all the clients communicating to LMTP server.

To Configure the LMTP Server

1. Add the **automatic-failover.rcp** file configuration to the existing LMTP server configuration. Or, if you prefer, there is also a sample recipe **LMTPBackendFailover.rcp** that configures a backend LMTP server for use with failover. If you wish to use this, you must copy the recipe script and manually add in to your LMTP client IP addresses and mailstore proxy information. This is available in the Messaging Server installed location *MessagingServer_home/lib/recipes/LMTPBackendFailover.rcp*.

2. Run the recipe script on all back-end machines in the deployment by executing the following the command.

```
msconfig run manual-failover.rcp OR LMTPBackendFailover.rcp
```

3. If you are running a compiled a configuration, recompile by running:

```
imsimta cnbuild
```

4. Start the Messaging Server by running:

```
start-msg
```

To Configure the Client

1. Add the **automatic-failover.rcp** file configuration to all clients' configurations. For the LMTP client, you must add some extra configuration.

2. Run the recipe script on all client machines (MMP, LMTP client, and do on) in the deployment by executing the following command:

```
msconfig run automatic-failover.rcp
```

3. For the LMTP client only, you must set the affinity list channel option on the LMTP client channel:

```
msconfig set channel:tcp_lmtpcs.affinitylist  
imsimta cnbuild  
stop-msg  
start-msg
```

4. Stop the Messaging Server Client by running:

```
stop-msg mmp or mta
```

5. Start the Messaging Server Client by running:

```
start-msg mmp or mta
```

Administering Message Store Database Snapshots (Backups)

This information describes the tasks for administering database snapshots. For conceptual information on database snapshots, see ["Message Store Automatic Recovery On Startup."](#)

The primary message store database is critical to smooth operation of the message store. You must always ensure that a snapshot (or backup) of the database is available. If the active database becomes damaged, restarting services allows the message store **stored** process to swap in the best snapshot and enable services to come back on demand.

Topics:

- [To Specify Message Store Database Snapshot Interval and Location](#)
- [Message Store Database Snapshot Recovery and Verification](#)
- [Message Store Database Snapshot Rolling Backup](#)
- [Message Store Database Recovery](#)

To Specify Message Store Database Snapshot Interval and Location

Before doing these tasks, see ["Message Store Database Snapshot Interval and Location."](#) Also, descriptions of the configutil options described in this page are in the *Messaging Server Reference*.

A database snapshot is a hot backup (dynamic backup) of the message store database. The system makes this copy without any locking, and requires that both the database files and transaction logs so the snapshot can be "recovered" into a valid copy of the database.

By default, snapshots are scheduled with the **imdbverify -s** command to run at specific times.

```
local.schedule.snapshot.enable = "1"  
local.schedule.snapshot = "0 2 * * * bin/imdbverify -s -m"
```

By default, the **imdbverify -s** command takes a database snapshot at 2 a.m. (This command uses UNIX crontab format: minute hour day-of-month month-of-year day-of-week command arguments.) The **-m** option is used to verify the snapshot. The **-m** option is not required. See ["Message Store Database Snapshot Recovery and Verification"](#) for more information.

Database snapshots are located in the following base directory:

```
local.store.snapshotpath = "dbdata/snapshots"
```

Change this directory to a different disk than the disk used by the primary database, for both performance and recovery reasons.

You configure the number of snapshots retained over time with the following option:

```
local.store.snapshotdirs = "3"
```

Each snapshot requires as much disk space as the entire database and transaction logs at any given time.

Take enough snapshots such that you both have a recent copy, and copies that go back a day or two, to be sure you can find a database not affected by any odd system problems that are not immediately discovered.

Message Store Database Snapshot Recovery and Verification

The message store has been enhanced to continuously recover archived log files into an up-to-date backup copy of the message store database. If the actual database becomes unusable, then the message store automatically uses this backup database. Having an up-to-date database backup provides the next level of recovery and stability for the message store.

The system automatically runs **imdbverify -m** as specified for rolling backups, and **imdbverify -s -m** as specified under snapshots.

If the verification process detects any errors, the errors are written to the **default** log. Errors in the **default** log mean not only that the snapshot failed, but they could also be pointing to a problem with the active database. (However, at this time, not all verification errors indicate a live database problem.)

Message Store Database Snapshot Rolling Backup

The message store rolling backup operates as a specially designated snapshot, where each transaction log is added to the snapshot and recovered every few minutes to provide a more up-to-date backup. For this reason, always enable snapshots, and rolling backup will be enabled by default.

Rolling backup requires the following three configuration settings, which are enabled by default:

```
local.store.rollingdbbackup = "yes"
```

```
local.schedule.snapshotverify.enable = 1
```

```
local.schedule.snapshotverify = "1,3,5,7,9,11,13,15,17,19,
    21,23,25,27,29,31,33,35,37,
    39,41,43,45,47,49,51,53,55,
    57,59 * * * * bin/imdbverify -m"
```

In this configuration:

- **local.store.rollingdbbackup** enables rolling backup. This means the log archive function running under the **stored** daemon copies the database transaction logs to the rolling snapshot instead of removing them every minute.
- **local.schedule.snapshotverify** verifies addition, which is required to continually roll the log files into the snapshot.

Note: Should a rolling backup fail any of its verifies, each side of the process declares the rolling backup invalid and cleans up the logs. Rolling backup then restarts after the next snapshot is put in place. Rolling backup relies on an initial snapshot taken by the normal snapshot process.

Message Store Database Recovery

When the message store services are started, the message store process **stored** decides if the current database is damaged, and if so, replaces it with the best snapshot. The best snapshot is printed in the logs and any recovery actions are also printed in the **default** log.

Messaging Server and Tiered Storage Overview

This chapter describes the operation of the Oracle Communications Messaging Server message store, its performance characteristics, and how to plan for and allocate store partitions. Additionally, this document describes next generation best practices to meet the storage needs of both ISPs and enterprises.

This chapter contains the following sections:

- [Overview of Messaging Server Storage](#)
- [Message Store and ZFS](#)
- [How the Message Store Works](#)
- [Background: Communication Services Logical Architectures Overview](#)
- [Background: "How Email Works" Introduction to Messaging Server](#)

Overview of Messaging Server Storage

For traditional ISPs that provide web-based email services, the rules of engagement have changed. Thanks to companies such as Google, which can now offer consumers multiple gigabytes of email storage space with unlimited retention, the threat is clear: ISPs, with their much smaller storage allotments (50-100 Mbytes) and automatic purging of messages older than 90 days, need to stay competitive by providing storage and retention capabilities similar to Google, or lose out.

The rules have also changed for enterprises and the corporate messaging market, which are being forced to comply with new regulations requiring email retention for ever longer periods of time. In fact, some companies are faced with the requirement of saving every incoming and outgoing email message forever. The storage requirements for such scenarios can be staggering, to say the least.

As if there weren't enough problems already for ISPs and enterprises, email by its nature is a very I/O intensive application where transaction speed is critical to customer satisfaction. In increasing storage capacity, businesses must ensure that email services maintain acceptable performance levels and do so in a cost-effective manner. To stay competitive, businesses understand that they must purchase large numbers of disk drives to satisfy this new appetite for storage space. Purchasing fiber channel drives is typically cost prohibitive, so ISPs and enterprises will be looking to less expensive alternatives, such as Serial Advanced Technology Attachment (SATA) drives. The challenge with SATA is that in order to reduce cost and provide higher capacity, performance is sacrificed. And there's the dilemma: both ISPs and enterprises need to dramatically increase their email storage capacity while at the same time maintaining acceptable performance levels without "breaking the bank." Customers are under immense legal and financial pressure to find a solution to this problem. The

good news is that Oracle can deliver an excellent storage solution for Messaging Server deployments in the form of the Sun StorageTek 6540 Array.

The Messaging Server message store is one of the highest IOPs applications that exists. In the past, customers have kept all portions of the message store on their highest performing, most expensive disks. Fortunately, you can distributed the message store components onto different performing disks to create a cost-effective but high performing application.

Care must be taken to keep the highest IOP portion of the message store (the database and its store partition indexes) on high performance disks. The message store can generate up to 15+ IOPS per message delivered, typically many small, random writes, and is extremely sensitive to response times. If response time diminishes, it can have a cascading effect through the application. Because of the high IOP needs, the message store is ideal for Oracle's StorageTek 6540 controller.

Message Store and ZFS

Oracle's Communications Suite Deployment Engineering group has performed extensive testing with ZFS and measured its impact on the message store. The Messaging Group believes that in the future most Messaging Server customers will be using the ZFS file system. ZFS changes the workload characteristics on the file system, so that there are fewer I/Os, but the I/Os are bigger. ZFS enables read rates to diminish somewhat, whereas it enables write rates to diminish much more. In addition, ZFS can perform snapshotting and compression, which enhances the ability to back up the application. ZFS is also now supported by Oracle Solaris Cluster software.

How the Message Store Works

The Message Store is a dedicated data store for the delivery, retrieval, and manipulation of Internet mail messages. The Message Store works with the IMAP4 and POP3 client access servers to provide flexible and easy access to messaging. The Message Store also works through the HTTP server (mshttpd) to provide messaging capabilities to Convergence clients in a web browser. The Message Store is organized as a set of folders or user mailboxes. The folder or mailbox is a container for messages. Each user has an INBOX where new mail arrives.

Each IMAP or Webmail user can also have one or more folders where mail can be stored. Folders can contain other folders arranged in a hierarchical tree. Mailboxes owned by an individual user are private folders. Private folders can be shared at the owner's discretion with other users on the same Message Store. Messaging Server supports sharing folders across multiple stores by using the IMAP protocol. There are two general areas in the Message Store, one for user files and another for system files. In the user area, the location of each user's INBOX is determined by using a two-level hashing algorithm. Each user mailbox or folder is represented by another directory in its parent folder. Each message is stored as a file. When there are many messages in a folder, the system creates hash directories for that folder. Using hash directories eases the burden on the underlying file system when there are many messages in a folder. In addition to the messages themselves, the Message Store maintains an index and cache of message header information and other frequently used data to enable clients to rapidly retrieve mailbox information and do common searches without the need to access the individual message files.

A Message Store can contain many message store partitions for user files. A Message Store partition is contained by a file system volume. As the file system becomes full,

you can create additional file system volumes and Message Store partitions on those file system volumes to store new users.

Note: If a message store partition fills up, users on the partition are not able to store additional messages. Address this problem by using one or more of the following approaches:

- Reducing the size of user mailboxes
 - If you are using volume management software, adding additional disks
 - Creating additional partitions and moving mailboxes to the new partitions
-
-

The message store maintains only one copy of each message per partition. This is sometimes referred to as a single-copy message store. When the message store receives a message addressed to multiple users or a group or distribution list, it adds a reference to the message in each user's INBOX. Rather than saving a copy of the message in each user's INBOX, the message store avoids the storage of duplicate data. The individual message status flag (seen, read, answered, deleted, and so on) is maintained per folder for each user.

The system area contains information on the entire message store in a database format for faster access and no loss of service. The information in the system area can be reconstructed from the user area. Messaging Server contains a database snapshot function. When needed, you can quickly recover the database to a known state.

Messaging Server also has fast recovery, so that in case of database corruption, you can shut down the message store and bring it back immediately without having to wait for a lengthy database reconstruction.

Messaging Server Disk Throughput


Disk throughput is the amount of data that your system can transfer from memory to disk and from disk to memory. The rate at which this data can be transferred is critical to the performance of Messaging Server. To create efficiencies in your system's disk throughput:

- Consider your maintenance operations, and ensure you have enough bandwidth for backup. Backup can also affect network bandwidth particularly with remote backups. Private backup networks might be a more efficient alternative.
- Carefully partition the store and separate store data items (such as tmp and db) to improve throughput efficiency.
- Ensure the user base is distributed across RAID (Redundant Array of Independent Disks) environments in large deployments
- Stripe data across multiple disk spindles in order to speed up operations that retrieve data from disk.
- Allocate enough CPU resources for RAID support, if RAID does not exist on your hardware.

Note: Measure disk I/O in terms of IOPS (total I/O operations per second) not bandwidth. You need to measure the number of unique disk transactions the system can handle with a very low response time (less than 10 milliseconds).

Typically, most customers deploy their entire message store on the highest performing disk as shown in Figure 74-1.

Figure 74-1 Typical Disk Storage with Message Store



Messaging Server Disk Capacity

When planning server system disk space, be sure to include space for operating environment software, Messaging Server software, and message content and tracking. Be sure to use an external disk array if availability is a requirement. For most systems, external disks are required for performance because the internal system disks supply no more than four spindles. For the Message Store partitions, the storage requirement is the total size of all messages plus 30 percent overhead. In addition, user disk space needs to be allocated. Typically, this space is determined by your site's policy.

Disk Sizing for MTA Message Queues

The purpose of the Messaging Server MTA Queue is to provide a transient store for messages waiting to be delivered. Messages are written to disk in a persistent manner to maintain guaranteed service delivery. If the MTA is unable to deliver the message, it retries delivery. At some point, if delivery is still unsuccessful, the MTA no longer tries to send the message and returns it to the sender.

MTA Message Queue Performance

Sizing the MTA Message Queue disks is an important step for improving MTA performance. The MTA's performance is directly tied to disk I/O above any other system resource. This means that you should plan on a disk volume that consists of multiple disk spindles which are concatenated and striped by using a disk RAID system. End users are quickly affected by the MTA performance. As users press the SEND button on their email client, the MTA does not fully accept receipt of the message until the message has been committed to the MTA Message Queue. Therefore,

improved performance on the MTA Message Queue results in better response times for the end-user experience.

MTA Message Queue Availability

SMTP services are considered a guaranteed message delivery service. This is an assurance to end users that the messaging server will not lose messages that the service is attempting to deliver. When you architect the design of the MTA Message Queue system, all effort should be made to ensure that messages will not be lost. This guarantee is usually made by implementing redundant disk systems through various RAID technologies.

MTA Message Queue Available Disk Sizing

The MTA Message Queue grows excessively if one of the following conditions occurs:

- The site has excessive network connectivity issues
- The MTA configuration is holding on to messages too long
- Valid problems are occurring with those messages (not covered in this document)
- Message stores are operationally down, for example, for maintenance
- Remote target sites are unavailable or overwhelmed

Performance Considerations for a Message Store Architecture

Message store performance is affected by a variety of factors, including:

- Disk I/O
- Inbound message rate (also known as message insertion rate)
- Message sizes
- Login rate (POP/IMAP/HTTP)
- Transaction rate for IMAP and HTTP
- Concurrent number of connections for the various protocols
- Network I/O
- Use of SSL

The preceding factors list the approximate order of impact to the Message Store. Most performance issues with the Message Storage arise from insufficient disk I/O capacity. Additionally, the way in which you lay out the store on the physical disks can also have a performance impact. For smaller standalone systems, it is possible to use a simple stripe of disks to provide sufficient I/O. For most larger systems, segregate the file system and provide I/O to the various parts of store.

Note: In a deployment using LMTP, the MTA queue is almost always unused.

Messaging Server Directories (General Recommendations for Storage)

Messaging Server uses six directories that receive a significant amount of input and output activity. Because these directories are accessed very frequently, you can increase performance by providing each of those directories with its own disk, or even better, providing each directory with a Redundant Array of Independent Disks (RAID).

The six directories are:

- MTA Queue Directory
- Messaging Server Log Directory
- Mailbox Database Files
- Message Store Index Files
- Message Files
- Mailbox List Database Temporary Directory

MTA Queue Directory

Recommendation: Can be located on slower disks or on shared NAS

In this directory, many files are created, one for each message that passes through the MTA channels. After the file is sent to the next destination, the file is then deleted. The directory location can be changed by making the queue directory a symlink.

Default location of the MTA Queue directory: `/var/opt/sun/comms/messaging/queue`

Messaging Server Log Directory

Recommendation: Can be located on slower disks

This directory contains log files which are constantly being appended with new logging information. The number of changes depend on the logging level set. The directory location is controlled by the `msconfig` option `*.logfile.logdir` (Unified configuration) or the `configutil` option `logfile.*.logdir` (legacy configuration), where `*` can be a log-generating component such as `admin`, `default`, `http`, `imap`, or `pop`. The MTA log files can be changed by making that directory a symlink.

Default location of the Messaging Server log directory:
`/var/opt/sun/comms/messaging/log`

Mailbox Database Files

Recommendation: Keep on a fast disk

These files require constant updates as well as cache synchronization. Put this directory on your fastest disk volume. These files are always located in the `/var/opt/sun/comms/messaging/store/mboxlist` directory.

Message Store Index Files

Recommendation: Keep on fast disk

These files contain meta information about mailboxes, messages, and users. By default, these files are stored with the message files. The `msconfig` option `partition.*.path` (Unified configuration) or the `configutil` option `store.partition.*.path` (legacy configuration), where `*` is the name of the partition, controls the directory location. If you have the resources, put these files on your second fastest disk volume.

Default location of the message store index file:
`/var/opt/sun/comms/messaging/store/partition/primary`

Message Files

Recommendation: Can be located on slower disks

These files contain the messages, one file per message. Files are frequently created, never modified, and eventually deleted. By default, they are stored in the same directory as the message store index files. The location can be controlled with the **msconfig** option **partition:partition_name.messagepath** (Unified configuration) or the **configutil** option **store.partition.partition_name.messagepath** (legacy configuration), where *partition_name* is the name of the partition. Some deployments might use a single message store partition called **primary** specified by the **store.partition.primary.path**. Large sites might have additional partitions that can be specified with **store.partition.partition_name.messagepath**, where *partition_name* is the name of the partition.

Default location of the message files:

/var/opt/sun/comms/messaging/store/partition/primary

Note: To separate the message files from the index files (enabling tiered storage), the **store.partition.*** and **.messagepath** options are key. These options must be correctly configured to put the message files on a SATA file system when you create the partition.

Mailbox List Database Temporary Directory

Recommendation: Can be located on fast disk

This is the directory used by the message store for all temporary files. To maximize performance, this directory should be located on the fastest file system. For Solaris OS, use the **msconfig** or the **configutil** options to configure the **store.dbtmpdir** (same in both Unified configuration and legacy configuration) variable to a directory under **tmpfs**, for example, **/tmp/mbolist**.

Default location of the mailbox list database:

/var/opt/sun/comms/messaging/store/mbolist

The following sections provide more detail on Messaging Server high-access directories.

Additional Information: MTA Queue Directories

In non-LMTP environments, the MTA queue directories in the Message Store system are also heavily used. LMTP works such that inbound messages are not put in MTA queues but directly inserted into the store. This message insertion lessens the overall I/O requirements of the message store machines and greatly reduces use of the MTA queue directory on Message Store machines. If the system is standalone or uses the local MTA for Webmail sends, significant I/O can still result on this directory for outbound mail traffic. In a two-tiered environment using LMTP, this directory is lightly used, if at all. In prior releases of Messaging Server, on large systems this directory set needs to be on its own stripe or volume.

MTA queue directories should usually be on their own file systems, separate from the message files in the message store. The message store has a mechanism to stop delivery and appending of messages if the disk space drops below a defined threshold. However, if both the log and queue directories are on the same file system and keep growing, you will run out of disk space and the message store will stop working.

Additional Information: Log Files Directory

The log files directory requires varying amounts of I/O depending on the level of logging that is enabled. The I/O on the logging directory, unlike all of the other high

I/O requirements of the Message Store, is asynchronous. For typical deployment scenarios, do not dedicate an entire LUN for logging. For very large store deployments, or environments where significant logging is required, a dedicated LUN is in order.

Note: In almost all environments, you need to protect the message store from loss of data. The level of loss and continuous availability that is necessary varies from simple disk protection such as RAID5, to mirroring, to routine backup, to real time replication of data, to a remote data center. Data protection also varies from the need for Automatic System Recovery (ASR) capable machines, to local HA capabilities, to automated remote site failover. These decisions impact the amount of hardware and support staff required to provide service.


Additional Information: mboxlist Directory

The **mboxlist** directory is highly I/O intensive but not very large. The **mboxlist** directory contains the databases that are used by the stores and their transaction logs. Because of its high I/O activity, and due to the fact that the multiple files that constitute the database cannot be split between different file systems, you should place the **mboxlist** directory on its own stripe or volume in large deployments. This is also the most likely cause of a loss of vertical scalability, as many procedures of the message store access the databases. For highly active systems, this can be a bottleneck. Bottlenecks in the I/O performance of the mboxlist directory decrease not only the raw performance and response time of the store but also impact the vertical scalability.

For systems with a requirement for fast recovery from backup, place this directory on Solid State Disks (SSD) or a high performance caching array to accept the high write rate that an ongoing restore with a live service will place on the file system.

Figure 74–2 depicts this configuration.

Figure 74–2 Next Generation Storage with Message Store



Multiple Store Partitions

The Message Store supports multiple store partitions. Place each partition on its own stripe or volume. The number of partitions that should be put on a store is determined by a number of factors. The obvious factor is the I/O requirements of the peak load on

the server. By adding additional file systems as additional store partitions, you increase the available IOPS (total IOs per second) to the server for mail delivery and retrieval. In most environments, you get more IOPS out of a larger number of smaller stripes or LUNs than a small number of larger stripes or LUNs. With some disk arrays, it is possible to configure a set of arrays in two different ways. You can configure each array as a LUN and mount it as a file system. Or, you can configure each array as a LUN and stripe them on the server. Both are valid configurations.

However, multiple store partitions (one per small array or a number of partitions on a large array striping sets of LUNs into server volumes) are easier to optimize and administer. Raw performance, however, is usually not the overriding factor in deciding how many store partitions you want or need. In corporate environments, it is likely that you need more space than IOPS. Again, it is possible to software stripe across LUNs and provide a single large store partition. However, multiple smaller partitions are generally easier to manage. The overriding factor of determining the appropriate number of store partitions is usually recovery time.

Recovery times for store partitions fall into a number of categories:

- First, the **fsck** command can operate on multiple file systems in parallel on a crash recovery caused by power, hardware, or operating system failure. If you are using a journaling file system (highly recommended and required for any HA platform), this factor is small.
- Secondly, backup and recovery procedures can be run in parallel across multiple store partitions. This parallelization is limited by the vertical scalability of the mboxlist directory as the Message Store uses a single set of databases for all of the store partitions. Store cleanup procedures (**expire** and **purge**) run in parallel with one thread of execution per store partition.
- Lastly, mirror or RAID re-sync procedures are faster with smaller LUNs. There are no hard and fast rules here, but the general recommendation in most cases is that a store partition should not encompass more than 10 spindles. The size of drive to use in a storage array is a question of the IOPS requirements versus the space requirements. For most residential ISP POP environments, use "smaller drives." Corporate deployments with large quotas should use "larger" drives. Again, every deployment is different and needs to examine its own set of requirements.

Setting Disk Stripe Width

When setting disk striping, the stripe width should be about the same size as the average message passing through your system. A stripe width of 128 blocks is usually too large and has a negative performance impact. Instead, use values of 8, 16, or 32 blocks (4, 8, or 16 kilobyte message respectively).

MTA Performance Considerations

MTA performance is affected by a number of factors including, but not limited to:

- Disk performance
- Use of SSL
- The number of messages/connections inbound and outbound
- The size of messages
- The number of target destinations/messages
- The speed and latency of connections to and from the MTA
- The need to do spam or virus filtering

- The use of Sieve rules and the need to do other message parsing (like use of the conversion channel)

The MTA is both CPU and I/O intensive. The MTA reads from and writes to two different directories: the queue directory and the logging directory. For a small host (four processors or fewer) functioning as an MTA, you do not need to separate these directories on different file systems. The queue directory is written to synchronously with fairly large writes. The logging directory is a series of smaller asynchronous and sequential writes. On systems that experience high traffic, consider separating these two directories onto two different file systems. In most cases, you will want to plan for redundancy in the MTA in the disk subsystem to avoid permanent loss of mail in the event of a spindle failure. (A spindle failure is by far the single most likely hardware failure.) This implies that either an external disk array or a system with many internal spindles is optimal.

MTA and RAID Trade-offs

There are trade-offs between using external hardware RAID controller devices and using JBOD arrays with software mirroring. The JBOD approach is sometimes less expensive in terms of hardware purchase but always requires more rack space and power. The JBOD approach also marginally decreases server performance, because of the cost of doing the mirroring in software, and usually implies a higher maintenance cost. Software RAID5 has such an impact on performance that it is not a viable alternative. For these reasons, use RAID5 caching controller arrays if RAID5 is preferred.

Background: Communication Services Logical Architectures Overview

You can deploy Communications Services in either a single-tiered or two-tiered logical architecture. Deciding on your logical architecture is crucial, as it determines which machine types you will need, as well as how many. In general, enterprise corporate deployments use a single-tiered architecture while internet service providers (ISPs) and telecommunications deployments use a two-tiered architecture. However, as with all generalities, the exceptions prove the rule. Small ISPs might just as well deploy on a single machine, and larger, centralized enterprises might deploy in a two-tiered architecture for many of the same reasons that ISPs do. As more and more corporations look to offer ease of access to employees working remotely, their deployments will begin to look more and more like an ISP.

Two-tiered Logical Architecture

In a two-tiered logical architecture, the data stores communicate through front-end processes. In the case of Messaging Server, this means MMPs and MTAs are residing on separate machines from the data store processes. A two-tiered architecture enables the mail store to offload important and common tasks and focus on receiving and delivering mail.

There might be some level of cohabitation with other services. For example, you could have the Calendar store and the Message Store on the same machine. Similarly, you could have the calendar front end on the MMP machine. In a two-tiered logical architecture, Directory Server is usually a complex deployment in its own right, with multi-master and replication to a set of load-balanced consumer directories.

Benefits of a Two-tiered Architecture

All services within the Communications Services offering rely on network capabilities. A two-tiered architecture provides for a network design with two separate networks: the public (user-facing) network, and the private (data center) network. Separating your network into two tiers provides the following benefits:

- **Hides Internal Networks.** By separating the public (user-facing) network and the private (data center) network, you provide security by hiding the data center information. This information includes network information, such as IP addresses and host names, as well as user data, such as mailboxes and calendar information.
- **Provides Redundancy of Network Services.** By provisioning service access across multiple front-end machines, you create redundancy for the system. By adding redundant messaging front-end servers, you improve service uptime by balancing SMTP requests to the available messaging front-end hosts.
- **Limits Available Data on Access Layer Hosts.** Should the access layer hosts be compromised, the attackers cannot get to critical data from the access hosts.
- **Offloads Tasks to the Access Layer.** By enabling the access layer to take complete ownership of a number of tasks, the number of user mailboxes on a message store increases. This is useful because the costs of both purchase and maintenance are much higher for store servers than for access layer machines (the second tier). Access layer machines are usually smaller, do not require large amounts of disk (see "[MTA Performance Considerations](#)," and are rarely backed up. A partial list of features that are offloaded by the second tier includes:
 - Denial of service protection
 - SSL
 - Reverse DNS
 - UBE (spam) and virus scanning
 - Initial authentication - Authentications to the Message Store should always succeed and the directory servers are more likely to have cached the entry recently.
 - LMTP - With support for LMTP between the MTA relays and the message stores, SMTP processing is offloaded and the need to do an additional write of the message into the MTA queues on the message stores is eliminated.
- **Simplifies End-user Settings in Client Applications.** By using a two-tiered architecture, end users do not have to remember the physical name of hosts that their messaging and calendar applications connect to. The Access-Layer Application hosts provide proxies to connect end users to their assigned messaging or calendar data center host. Services such as IMAP are connected to the back-end service using LDAP information to identify the name of the user's mailbox host. For calendar services, the calendar front-end hosts provide a calendar lookup using the directory server to create a back-end connection to the user's assigned calendar store host.

This capability enables all end users to share the same host names for their client settings. For example, instead of remembering that their message store is host-a, the user simply uses the setting of mail. The MMP provides the proxy service to the user's assigned message store. You need to provide the DNS and load balancing settings to point all incoming connections for mail to one (or more) MMPs.

By placing Messaging Server into two tiers, more than one Messaging Server back-end server can be used.

An additional benefit of this proxy capability provides geographically dispersed users to leverage the same client application settings regardless of their physical location. Should a user from Europe visit California, the user will be able to connect to the immediate access server in California. The user's LDAP information will tell the access server to create a separate connection on the user's behalf to the user's message store located in Europe. Lastly, this enables you to run a large environment without having to configure user browsers differently, simplifying user support. You can move a user's mailbox from one mail store to another without contacting the user or changing the desktop.

- **Reduces Network HTTP Traffic on the Data Center.** The Messaging Server front-end greatly reduce HTTP traffic to the data center network. HTTP provides a connectionless service. For each HTML element, a separate HTTP request must be sent to the mail or calendar service. These requests can be for static data, such as an image, style sheets, JavaScript files, or HTML files. By placing these elements closer to the end user, you reduce network traffic on the back-end data center.

Horizontal Scalability Strategy

Scalability is critical to organizations needing to make the most cost-effective use of their computing resources, handle peak workloads, and grow their infrastructure as rapidly as their business grows. Keep these points in mind:

- How the system responds to increasing workloads: what performance it provides, and as the workload increases, whether it crashes or enables performance to gracefully degrade.
- How easy it is to add processors, CPUs, storage, and I/O resources to a system or network to serve increasing demands from users.
- Whether the same environment can support applications as they grow from Low-end systems to mid-range servers and mainframe-class systems.

When deployed in a two-tiered architecture, Messaging Server is meant to scale very effectively in a horizontal manner. Each functional element can support increased load by adding additional machines to a given tier.

Scaling Front-end and Back-end Services

In practice, the method for scaling the front-end and back-end services differs slightly. For Tier 1 elements, you start the scaling process when traffic to the front end grows beyond current capacity. You add relatively low cost machines to the tier and load balance across these machines. Thus, load balancers can precede each of the Tier 1 service functions as overall system load, service distribution, and scalability requirements dictate.

For Tier 2 elements, you start the scaling process when the back-end services have exceeded user or data capacity. As a general rule, design the Tier 2 services to accommodate just under double the load capacity of the Tier 1 services. For example, for an architecture designed for 5,000 users, the Tier 1 front-end services are designed to support 5,000 users. The back-end services are then doubled, and designed to accommodate 10,000 users. If the system capacity exceeds 5,000 users, the front-end services can be horizontally scaled. If the overall capacity reaches 5,000 users, then the back-end services can be scaled to accommodate. Such design enables flexibility for growth, whether the growth is in terms of users or throughput.

Implementing Local Message Transfer Protocol (LMTP) for Messaging Server

Best practices say you should implement LMTP to replace SMTP for message insertion. An LMTP architecture is more efficient for delivering to the back-end Message Store because:

- Reduces the load on the stores. Relays are horizontally scalable while stores are not, thus it is a good practice to make the relays perform as much of the processing as possible.
- It reduces IOPS by as much as 30 percent by removing the MTA queues from the stores.
- It reduces the load on LDAP servers.
- The LDAP infrastructure is often the limiting factor in large messaging deployments.
- It reduces the number of message queues.
- It requires a small amount of shared/NFS storage across all inbound MTAs.

You need a two-tiered architecture to implement LMTP.

Background: "How Email Works" Introduction to Messaging Server


This write up is a basic introduction. It does avoid most of the more complicated configurations and mechanisms which are part of the Messaging Server. The key objective is to understand "how email works" from a basic "black box" model. As this section progresses, you will start getting the more complicated information on the internals of the Messaging Server.

What Does Messaging Server Enable Users to Do?

The Messaging Server enables users to interact with their email message data in three different manners:

1. Users can write an email, and then send it to someone(s).
2. Users can receive emails from others.
3. Users can access their mailbox.

Figure 74–3 How Messaging Server Allows Users to Interact with Message Data



The key components of Messaging Server are:

- Message Transport Agent (MTA)
- Message Store Database
- Message Access Services
- IMAP/POP3
- Webmail Access for Convergence (separate product)


Figure 74–3 simplifies the Messaging Server into the following three components:

- Messages are delivered into the Messaging Server via the MTA. The MTA is the SMTP Server which accepts the email message and then route the message to it's destination. The destination for the email could be either the Message Store or to another server somewhere out over the network.
- The Message Store is an Oracle private database, which stores user's mailboxes and message data.
- Users access their mailboxes via either IMAP, POP3 or through Webmail.


A User Decides to Send an Email

Users write their emails using any of the popular email clients out on the market. These email clients could include Mozilla/Thunderbird, Apple Mail, Outlook, or any other IMAP-based email application. Once the email is ready, the user would click the "SEND" button.

The "SEND" button allows the email application to connect with the MTA. The client connects to the MTA using the SMTP protocol.

Figure 74-4 Client Connects to the Messaging Server MTA

On the connection to the MTA, the client will send the message over the network to the MTA. The MTA would then store the message in the Message Queue Disk and release the client connection. The message stored on the queue will then be read and the address of the message will be evaluated by using information stored in the Directory Server using LDAP.

Figure 74-5 MTA Stores Message to Message Queue and Message Evaluated Through LDAP

At this point, the MTA will attempt to deliver the message to one of three places:

1. Into a User's Mailbox
2. Into another Messaging Server
3. Into another site on the Internet (note for security: it is recommend that the MTA forward the message to a "smarthost" MTA to handle the message delivery over the Internet)


It is important to note that once the message is on the Message Disk Queue, and the MTA has responsibility of the message. We guarantee this responsibility by committing the message to disk. It is this commitment which provides the first clue into a our first performance limitation. We write all MTA data to a disk queue, and we therefore are bounded by the disk i/o performance of that disk system. We also need to ensure that the disk system is highly available through RAID. (Typically RAID 0+1 or RAID 5.)

Postal Service Example: An example of this would be a letter that you drop into a postal mailbox. If you drop a letter of at the Post Office, you are assuming that the Postal Service will mail the message to the destination. You will not expect the message to get lost.

The MTA will not remove the message from the Message Disk Queue until a successful delivery. If a delivery attempt fails, it will keep trying. If it fails a final time,

the message will be HELD on the disk queue and a notification sent back to the sender. The MTA logs will record the whole history.

Figure 74–6 How the MTA Delivers a Message




The MTA will attempt to deliver the message in one of three possible ways:

1. Message is Deposited into the User's Mailbox on the Message Store.
2. Message is Sent to another MTA Server on the Same Local Network.
3. Message is sent to the Internet by first sending to the Smarthost.

The Smarthost MTA sits between the Messaging Server inside the protected network and the Internet. Typically this would be in an area outside the internal network. The Smarthost would evaluate the message address by domain name. Using DNS, the MTA would attempt to identify the DNS MX Record and A Record for the destination of the email. It would then attempt to deliver the message to the host of that IP Address using SMTP.


User Receives an Email

Figure 74-7 How a User Receives Email from the MTA



An email sent to a user on the Message Store. In the case of the message arriving from the Internet, the email sender would address the email to "jon@example.com". The email sender's MTA would look up "example.com" in DNS and find either an MX record or an A record for this site.


Figure 74-8 Inbound MTA



The Inbound MTA for the site may be configured to detect SPAM or Viruses. In this case, the message will be sent to a anti-spam/anti-virus verdict engine (such as Cloudmark or Symantec Brightmail). If the email passes, it is evaluated for message routing. The email address is evaluated in the Directory Server. If the email address is valid, it is then sent to the user's mailbox host (mailbox).

User Access Mailbox

Figure 74–9 User Access to Mailbox




If users want to access their email, they use their email client. These clients, such as Mozilla Thunderbird or Apple Mail, connect to the mailbox using the IMAP protocol. In some cases, ISPs limit access to POP3 protocol.

The email client would connect to Messaging Server's Mail Multiplexor Proxy (MMP). The MMP would accept the user's email credentials (their username and password) and validate their login using the Directory Server. Should the authentication be successful, the MMP would then identify the user's backend Message Store server. The MMP would then connect to this backend Message Store on the user's behalf. The user can then access their mailbox data.

The value of the MMP is that it provides 1) Security and 2) Scalability. Security is improved by the use of the MMP by isolating the Message Store from the end-users. This helps to prevent unauthorized access (hacking) into the Message Store server data. Scalability (as seen in the figure below) is improved through Horizontal Scalability. This means that the architecture can grow by adding additional MMPs and backend Message Store servers.

Figure 74–10 Growing Access Layer Capacity



Message Store Command Reference

This chapter describes the Oracle Communications Messaging Server Message Store command reference.

configutil

The **configutil** utility enables you to list and change Messaging Server configuration options.

For a list of all configuration options, see the *Messaging Server Reference*.

All Messaging Server configuration options and values are stored locally in the **msg.conf** and **msg.conf.defaults** files. The **msg.conf.defaults** file must never be edited and contains defaults constructed during initial configuration. The **msg.conf** file contains settings that have been explicitly overridden from their default value using **configutil**. Use the **-H** option to **configutil** to view a setting's default value. Use **configutil** to edit **configutil** settings; do not edit these files directly.

Note: If the administrator has defined any language-specific options (such as messages), you must use the **language** option at the end of the command in order to list or change them. Commands entered without a **language** option are only applied to attributes that do not have a specified language option.

Requirements: Must be run locally on the Messaging Server. You may run **configutil** as **root** or **mailsrv**. If you make changes to the servers, you must restart or refresh the servers, depending on the variable, for the changes to take effect.

Location: *MessagingServer_home/bin/configutil*

You can use **configutil** to perform four tasks:

- Display particular configuration options using **-o option**.
Add **;lang-xx** after the option to list options with a specified language option. For example, **;lang-jp** to list options specified for the Japanese language.
- List configuration option values using the **-p pattern** option. (Can be used with the **-m** option.)

Use **-p pattern** to just list those configuration options whose names contain the pattern specified in *pattern*. * is the wildcard character and is assumed to follow *pattern* if there is no wildcard already in *pattern*.

Use **-m** to show whether or not the listed options are refreshable.

- Set configuration options using the **-o** *option* and **-v** *value* options.
Add **,lang-xx** after the option to set options for a specified language option. For example, **,lang-jp** to set options specified for the Japanese language.
- Import configuration option values from **stdin** using the **-i** option.
Include the **-H** option to show settings with help.

Examples:

```
configutil -H
```

```
Configuration option: alarm.diskavail.msgalarmdescription
  Description: Description for the diskavail alarm.
  Syntax: string
  Default: percentage mail partition diskspace available
alarm.diskavail.msgalarmdescription is currently set to: percentage mail
partition diskspace available
```

```
Configuration option: alarm.diskavail.msgalarmstatinterval
  Description: Interval in seconds between disk availability
checks. Set to 0 to disable checks of disk usage.
  Syntax: int
  Default: 3600
alarm.diskavail.msgalarmstatinterval is currently set to: 3600
[.....]
```

```
### Show help on all configuration parameters ending in ".port"
```

```
configutil -p \*.port -H
Configuration option: local.ens.port
  Description: Port (and optionally, a specific IP address) ENS
server will listen on, in the format of [address:]port, for example,
7997 or 192.168.1.1:7997.
  Syntax: string
  Default: 7997
local.ens.port is currently set to: 7997
```

```
Configuration option: local.service.http.proxy.port
  Description: Configures the port number of the back-end
Messenger Express (HTTP) server with the Messaging Multiplexor.
  Syntax: uint
  Default: 80
local.service.http.proxy.port is currently set to: 80
```

```
Configuration option: local.snmp.port
  Description: SNMP subagent port number.
  Syntax: uint
  Default: 0
local.snmp.port is currently set to: 0
```

```
Configuration option: local.watcher.port
  Description: watcher listen port.
  Syntax: uint
  Default: 49994
local.watcher.port is currently set to: 49994
```

```
Configuration option: local.webmail.cert.port
  Description: Specifies a port number on the machine where the
Messaging Server runs to use for CRL communication. This port is used
locally for that machine only. The value must be greater than 1024.
```

```

        Syntax: int
        Default: 55443
local.webmail.cert.port is currently set to: 55443

Configuration option: local.webmail.da.port
        Description: Delegated Administrator port.
        Syntax: int
        Default: 8080
local.webmail.da.port is currently set to: 8080

Configuration option: local.webmail.sieve.port
        Description: The port of the web container where the Mail Filter
has been deployed.
        Syntax: string
        Default: NULL (unset)
local.webmail.sieve.port is currently unset

Configuration option: metermaid.config.port
        Description: Port number on which MeterMaid listens for
connections.
        Syntax: tcpport
        Default: 63837
metermaid.config.port is currently set to: 63837

Configuration option: service.http.port
        Description: Messenger Express HTTP port.
        Flags: MSG_RESTART_HTTP
        Syntax: uint
        Default: 80
service.http.port is currently set to: 81

Configuration option: service.imap.port
        Description: IMAP server port number.
        Flags: MSG_RESTART_IMAP
        Syntax: uint
        Default: 143
service.imap.port is currently set to: 143

Configuration option: service.pop.port
        Description: POP server port number.
        Flags: MSG_RESTART_POP
        Syntax: uint
        Default: 110
service.pop.port is currently set to: 110

### Show help on store.partition.\*.path

configutil -p store.partition.\*.path -H
Configuration option: store.partition.\*.path
        Description: Controls the store index file directory path.
        Flags: MSG_RESTART_ALL
        Syntax: filepath
        Default: NULL (unset)
store.partition.primary.path is currently set
to: /opt/SUNWmsgsr/data/store/partition/primary
store.partition.three.path is currently set
to: /opt/SUNWmsgsr/data/store/partition/three
store.partition.two.path is currently set
to: /opt/SUNWmsgsr/data/store/partition/two

```

```

Configuration option: store.partition.primary.path
Description: Full path name of the primary partition.
Flags: MSG_RESTART_ALL
Syntax: filepath
Default: <msg.RootPath>/data/store/partition/primary
store.partition.primary.path is currently set
to: /opt/SUNWmsgsr/data/store/partition/primary

```

Syntax

```

configutil [-h] [-f <configfile>] [-o <option>[;<language>]] [-v <value>]
configutil [-f <configfile>] [-p <pattern>] [-H] [-m] [-V]
configutil -i <inputfile>

```

Options

The options for this command are:

Table 75–1 configutil Options

Option	Description
-d	Enables you to delete an option. Used with the -o option.
-f <i>configfile</i>	Enables you to specify a local configuration file other than the default. (This option uses information stored in the ConfigRoot environment variable by default.)
-h	Shows usage statement.
-H	Enables you to get help on options. Used with the -o option.
-i <i>inputfile</i>	Imports configurations from a file. Data in the file to be entered in <i>option=value</i> format with no spaces on either side of the pipe. The <i>inputfile</i> should be specified as an absolute path.
-m	Lists meta data.
-o <i>option</i>	Specifies the name of the configuration option that you wish to view or modify. May be used with the -H , -v , -d options. Configuration option names starting with the word local are stored in the local server configuration file.
-p <i>pattern</i>	Lists only options with the given pattern (* is wildcard). Example: configutil -p *enable -H . If no wildcard is present in the <i>pattern</i> , a wildcard at the end of the <i>pattern</i> will be assumed (effectively changing <i>pattern</i> to <i>pattern*</i>).
-v <i>value</i>	Specifies a value for a configuration option. To be used with -o <i>option</i> .
-V	Enables you to validate the configuration against meta data.

If you specify no command-line options, all configuration options are listed.

Examples

To list all configuration options and their values in both the Directory Server LDAP database and local server configuration file:

```
configutil
```

To import configurations from an input file named **config.cfg**:

```
configutil -i config.cfg
```

To list all configuration options with the prefix **service.imap**:

```
configutil -p service.imap
```

To display the value of the **service.smtp.port** configuration option:

```
configutil -o service.smtp.port
```

To set the value of the **service.smtp.port** configuration option to 25:

```
configutil -o service.smtp.port -v 25
```

To clear the value for the **service.imap.banner** configuration option:

```
configutil -o service.imap.banner -v ""
```

To display the refreshable status of the **service.pop** configuration options:

```
configutil -m -p service.pop
```

This example of the **-m** option could produce the following sample output:

```
service.pop.allowanonymouslogin = no [REFRESHABLE]
service.pop.banner = "%h %p service (%P %V)" [REFRESHABLE]
service.pop.createtimestamp = 20030315011827Z [REFRESHABLE]
service.pop.creatorsname = "cn=directory manager" [REFRESHABLE]
service.pop.enable = yes [NOT REFRESHABLE]
service.pop.enablesslport = no [NOT REFRESHABLE]
service.pop.idletimeout = 10 [REFRESHABLE]
service.pop.maxsessions = 600 [NOT REFRESHABLE]
service.pop.maxthreads = 250 [NOT REFRESHABLE]
```

Language Specific Options

To list or set options for a specific language, append **;lang-xx** immediately after the option with no spaces, where **xx** is the two-letter language identifier. For example, to view the text of the Japanese version of the **store.quotaexceededmsg** message:

```
configutil -o "store.quotaexceededmsg;lang-jp"
```

The semicolon is a special character for most UNIX shells and requires special quoting as shown in the example.

counterutil

The **counterutil** utility displays and changes counters in a counter object. It can also be used to monitor a counter object at fixed intervals. For more information and examples see "[Gathering Message Store Counter Statistics by Using counterutil.](#)"

Requirements: Must be run locally on the Messaging Server.

Location: *MessagingServer_home/bin/*

Syntax

```
counterutil -l
```

```
counterutil -o _counterobject_ [-i _interval_] [-n _numiterations_]
```

```
counterutil -s -o _counterobject_ -c _counter_
```

Options

The options for this command are:

Table 75–2 counterutil Options

Option	Description
-c <i>counter</i>	Specifies a particular counter associated with a counter object.
-i <i>interval</i>	Specifies, in seconds, the interval between reports. The default is 5.
-l	Lists the content of the counter registry.
-n <i>numiterations</i>	Specifies the number of iterations. The default is infinity.
-o <i>counterobject</i>	Displays the contents of a particular counter object at specified intervals (by default, every 5 seconds). The valid counter objects are: imapstat , popstat , httpstat , alarm , diskusage , serverresponse , mmpstat . Each counter object has a list of counters. When you monitor a counter object, counterutil displays the counters associated with it.
-s	Resets the counter to 0.

Examples

To list all counter objects in a given server's counter registry:

```
counterutil -l
```

To monitor the content of a counter object **imapstat** every 5 seconds:

```
counterutil -o imapstat
```

To reset the counter **global.maxconnections**, associated with the counter object **imapstat**, to zero:

```
counterutil -s -o imapstat -c global.maxconnections
```

For usage information on **counterutil**, refer to the "[Gathering Message Store Counter Statistics by Using counterutil.](#)"

deliver

The **deliver** utility delivers mail directly to the message store.

If you are administering an integrated messaging environment, you can use this utility to deliver mail from another MTA, a **sendmail** MTA for example, to the Messaging Server message store.

Note: The **deliver** utility is only for use with files which are already completely and properly formed email messages (RFC 822).

Requirements: Must be run locally on the Messaging Server; the **stored** utility must also be running. Make sure that the environment variable **ConfigRoot** is set to *MessagingServer_home/config*.

Location on UNIX: *MessagingServer_home/bin/*

Syntax

```
deliver [-r <address>] [-m <mailbox>] [-g <flag>] [-q] [-c]
```

```
[-p] [-d] <userid>... | <pattern>
```

You can specify multiple userids. If you specify no options, mail is delivered to the inbox of the user specified in *userid*.

Options

The options for this command are:

Table 75–3 *deliver Options*

Option	Description
-c	Automatically creates the mailbox if it doesn't exist in the message store. <ul style="list-style-type: none"> ▪ The <i>mailbox</i> specified by the -m option will only be created if the sender specified by the <i>authid</i> passed in with the -a option has CREATE ("c") rights. ▪ If no <i>mailbox</i> is specified, the inbox for <i>userid</i> will be created.
-d	This option is recognized but ignored by deliver in order to maintain compatibility with /bin/mail .
-g <i>flag</i>	Sets the system flag or keyword flag on the delivered message. <ul style="list-style-type: none"> ▪ Valid system flags are 'Seen', 'Flagged', 'Deleted', 'Answered' and 'Draft'. ▪ The SEEN ("s"), DELETED ("d") and WRITE ("w") mailbox access control values determine whether the 'Seen', 'Deleted' and all other flags can be set by the sender specified by the <i>authid</i> value that has been passed in with the -a option. ▪ The value of <i>flag</i> will only be set if the sender has the required access rights to set that particular flag on the mailbox. ▪ The -g option can be used multiple times to set more than one flag.
-m <i>mailbox</i>	Delivers mail to a specific <i>mailbox</i> . <ul style="list-style-type: none"> ▪ Attempt to deliver the message to <i>mailbox</i> for each <i>userid</i>. If the access control on <i>mailbox</i> does not grant the POST ("p") right to the sender specified by the <i>authid</i> passed in with the -a option, the message will be delivered to the inbox for the <i>userid</i> instead. ▪ If the <i>mailbox</i> does not exist for the <i>userid</i>, the message will be delivered to the inbox instead.
-p	Deliver to the inbox of the users which match the wildcard <i>pattern</i> . <ul style="list-style-type: none"> ▪ The -c and -m options have no impact when -p is used.
-q	Overrides mailbox quotas. Delivers messages even when the receiving mailbox is over quota.
-r <i>address</i>	Inserts a Return-Path: header containing address.
<i>userid</i> <i>pattern</i>	Deliver to the mailbox of the user(s) specified by <i>userid</i> . Multiple <i>userid</i> values can be specified. <i>pattern</i> will be expanded to the appropriate list of existing users when the -p option is used.

Tip: The `mboxutil -lpx user/userid/mailbox` command can be used to check the access controls on *mailbox* for *userid*.

Debugging and Troubleshooting Options

Table 75–4 *deliver Debugging and Troubleshooting Options*

Option	Description
<code>-t seconds</code>	Deliver as many copies of the message as possible to the <i>userid</i> in <i>seconds</i> time. This option is primarily meant for load-testing a mailbox.

Examples

To deliver the contents of a file named **message** to Fred's **tasks** mailbox:

```
deliver -m tasks fred < message
```

In the above example, if the **tasks** mailbox does not grant "p" rights to the value of the `authid` passed in with the `-a` option or the **tasks** mailbox does not exist, the contents of **message** is delivered to the inbox of the user **fred**.

To deliver the contents of a file named **message** to the inbox of all existing users on the store

```
deliver -p '*' < message
```

To deliver the contents of a file named **message** to the inbox of all existing users in the **example.com** hosted domain.

```
deliver -p '*@example.com' < message
```

Tip: Use the `mboxutil` command to determine which users will receive a copy of the message based on a given user *pattern*, e.g. `mboxutil -lp "user/pattern/INBOX"`

To deliver the contents of a file named **message** to Fred's **Important** mailbox, and set the '**Flagged**' IMAP flag on that new message.

```
deliver -m Important -g '\Flagged' fred < message
```

hashdir

The `hashdir` command calculates the hash where the specified user ID would be found in a store partition. If the user is not in the default domain, you should append `@<domain>` to the user ID.

Syntax

```
hashdir [-a] [-i] <userid>[@<domain>]
```

Options

The options for this command are:

Table 75–5 hashdir Options

Option	Description
-a	Appends the specified user ID to the output.
-i	Allows you to use the command in interactive mode.

Examples

```
hashdir user1
64/b1/
hashdir user1@domain.com
11/05/
```

Note: The **hashdir** command does not validate the input to determine whether it is a valid userid or whether the mailbox exists in the store. Also, the user ID is case sensitive, specifying the wrong case will generate a different hash. To see where the folder exists in the store, the "mboxutil" command may be more useful.

imarchive

The **imarchive** utility supports the AXS-One archiving system for Messaging Server. **imarchive** reads the AXS-One archive report files and performs one of the following actions:

- Marks the messages as archived in the message store. Messages remain in both the message store and archive system. Marking a message prevents it from being re-archived. Marking messages is the default action.
- Creates URL stubs of the messages. The user will see the same header, internal date, save date, and message flags, but instead of the message text, the user will have a clickable URL to retrieve the message text.

Requirements: Must be run locally on the Messaging Server. You can run **imarchive** as **root** or as **mailsrv**.

Location: *MessagingServer_home/bin/*

Syntax

```
imarchive [-s] [-v] [-p <dir>]
```

Options

The options for this command are:

Table 75–6 imarchive Options

Option	Description
-pdir	Processes archive reports in the specified directory
-s	Replaces archived messages with stubs
-v	Provides verbose logging information

If you do not use the **-s** option, **imarchive** takes the default action, marking the messages as archived in the message store.

Example

To stub messages:

```
imarchive -s
```

imcheck

The **imcheck** utility prints mailbox data and metadata. In addition, **imcheck** prints **mbxlist** database data for specified database files and prints database statistics.

Syntax

```
imcheck [-m mailbox [-c|-b messagenum] | -x dir [-c|-b messagenum] | -p partition
|
-f file] [-e | -H] [-D [-S separator]]
```

```
imcheck -q
```

```
imcheck -d <databasename> [-S separator]
```

```
imcheck -s [-n] [subsystem...]
```

Options

The options for this command are:

Table 75–7 imcheck Options

Option	Description
-b <i>messagenum</i>	Prints the contents of the message specified by <i>messagenum</i> .
-c <i>messagenum</i>	Prints cache data for the specified message number.
-d <i>databasename</i>	Prints the contents of the database specified by <i>databasename</i> .
-D	Prints metadata in imsbackup format to verify the success of the backup operation.
-e	Prints the uids of expunged messages.
-f <i>file</i>	Prints metadata for the mailboxes listed in the specified file. You must specify a full path name for <i>file</i> .
-H	Print mailbox header only.
-m <i>mailbox</i>	Prints metadata for the specified <i>mailbox</i> . The <i>mailbox</i> name must be in Modified UTF-7 encoded format. You can use the "-E M-UTF-7" option with the mbxutil command to list mailboxes in Modified UTF-7 format. For information about the output produced by this option, see "imcheck -m Metadata Output."
-n	Specifies that no locking will occur. This option should be used only for debugging database hang problems. Caution Only use this option when the database is hung. It can cause data corruption in the message store.
-p <i>partition</i>	Prints metadata for every mailbox in the specified message store partition.

Table 75-7 (Cont.) imcheck Options

Option	Description
-q	Displays the maintenance queue. See "Displaying the Maintenance Queue" for more information.
-s	Prints database statistics.
-S separator	Specifies an output field separator. The default value is a vertical bar: " "
-x dir	Prints metadata for the mailbox under the directory specified by <i>dir</i> .
subsystem	Prints the replication statistics. subsystem can be spool , lock , log , txn , or rep .

Note: When using the separator, check output consistency by trying the separator string as a single or double quoted string with escape sequences if required.

imcheck -m Metadata Output

The **imcheck -m mailbox** command displays metadata that describes the entire mailbox and metadata that describes each message.

The message-specific data is displayed in a table. Most of the fields in the table are self-explanatory. However, the following fields need further explanation:

HSize - Header size

MT - Message type ID, defined by **configutil** options such as **store.message.type.*** and **store.message.type.enable**. For a list of these options, see the *Messaging Server Reference*.

SFlags - System flags. The system flags are as follows:

```
R : Recent
S : Seen
D : Deleted
A : Answered
F : Flagged
T : Draft
B : Stubbed
C : Archived
E : Encrypted
N : NoLeadingDots
```

In the **imcheck -m mailbox** output, the system-flag metadata is displayed as a character. For example: **S**

indicates that the Seen flag has been set.

UFlags - Flags defined by the user. User-flags are displayed as a hex number. The binary form of the hex number represents the user-flag switches. For example, if the user has defined six flags, the value

3f0000

indicates that all six flags are set.

Examples

To dump metadata to verify the **imsbackup** operation performed on the user **jdoue**'s INBOX, separating the output with a colon (":"):

```
imcheck -D -S ":" -m user/jdoue/INBOX
```

To dump the contents of the **folder.db** database:

```
imcheck -d folder.db
```

To print metadata for the inbox of the user **jdoue**:

```
imcheck -m user/jdoue/INBOX
```

```
-----
Name: user/jdoue
Version: 102
Exists: 10
Flags: 0
Largest Msg: 1094 bytes
Last Append: 20080801062616
Last UID: 1008276527
Oldest Msg: 20000621093214
Oldest Uid: 2
Quota Used: 10061
UID Validity: 1008099930
Cache Offset: 7856
Start Offset: 256
ACL: jdoue      lrswipcda
Userflags:
    f1
    f2
    f3
    f4
    f5
    f6
Subscribed: 0
Partition: primary
Path: /var/opt/SUNWmsgsr/store/partition/primary/=user/64/bl/=jdoue
Msg Path: /var/opt/SUNWmsgsr/store/partition/primary/=user/64/bl/=jdoue

MsgNo Uid Internal-Date Sent-Date Size HSize Cache-Id C-Offset C-Len Last-Updated
Save-Date MT SFlags UFlags Original-Uid Message-id
-----
1 1 20080801061303 20080801061259 752 744 1 16 1324 20080801061303 20080801061303
0 S 0.0.0 1217596383-1 <200808011312.m71DCxJp017783@dumbo.example.com>
2 2 20080801062616 20080801062615 781 772 1 1340 1440 20080801062616
20080801062616 0 S 0.0.0 1217596383-2
<200808011326.m71DQFYb018990@dumbo.example.com>
```

To print metadata for the Sent folder of a Japanese User **jauser**:

```
bash-3.00# ./mboxutil -E "M-UTF-7" -lp user/jauser/*
      msgs  Kbytes last msg          partition  quotaroot mailbox
      ---  ---  ---  ---  ---  ---
      0      0 -          -          primary      5120 user/jauser/INBOX
      0      0 -          -          primary      - user/jauser/&MFQwf3ux-
      0      0 -          -          primary      - user/jauser/&Tgtm+DBN-
      1      0 2009/03/23 11:59 primary      - user/jauser/&kAFP4W4IMH8-
```

```
bash-3.00# ./imcheck -m "user/jauser/&kAFP4W4IMH8-"
-----
```

```

Name: user/jauser/&kAFP4W4IMH8-
Version: 102
Exists: 1
Flags: 0
Largest Msg: 898 bytes
Last Append: 20090323115905
Last UID: 1
Oldest Msg: 20090323115905
Oldest Uid: 1
Quota Used: 898
UID Validity: 1237768457
Cache Offset: 7856
Start Offset: 256
ACL: jauser      lrswipcda
Subscribed: 1
Partition: primary
Path:
/opt/SUNWmsgsr/data/store/partition/primary/=user/20/b0/=jauser/=&k+A+F+P4+W4+I+M+
H8-
Msg Path:
/opt/SUNWmsgsr/data/store/partition/primary/=user/20/b0/=jauser/=&k+A+F+P4+W4+I+M+
H8-

MsgNo  Uid  Internal Date      Sent Date      Size HSize CacheOff  Last Updated
Save Date MT SFlags UFlags   Original-Uid   Message-id
-----
          1      1 20090323115905 20090323115905      898   549      7856 20090323115905
20090323115905 0 S      0.0.0 1237768457-1
<cd8b67384048.49c77989@dumbo.example.com>

```

imdbverify

The **imdbverify** utility takes and verifies message store database snapshots. For more information on running **imdbverify**, see "[Administering Message Store Database Snapshots \(Backups\)](#)."

Syntax

```
imdbverify -s [-m]
```

Options

The options for this command are:

Table 75–8 *imdbverify* Options

Option	Description
-s	Takes a full snapshot of the message store database and saves it in the configured destination. Incremental snapshot is performed when this option is not specified.
-m	Verify all database files.

imexpire

imexpire automatically expires messages in the message store based on administrator-specified criteria. The "[impurge](#)" command purges the messages.

Topics:

- [Expire Actions](#)
- [Expire Criteria](#)
- [Requirements](#)
- [Location](#)
- [Syntax](#)
- [Options](#)
- [Examples](#)
- [Additional Functionality](#)

Expire Actions

The expire phase can perform one of these actions:

- Discard - removes messages from the mailbox immediately.
- Archive - archives messages by using the AxsOne archive store.
- Fileinto - moves messages to a specified mailbox folder.
- Report - prints the specified mailbox name, **uid**, and **uid** validity to **stdout**.

By default, **imexpire** discards messages.

For details about how the **imexpire** actions can be performed, see "[Message Store Maintenance Queue](#)," "[Configuring Message Expiration \(Tasks\)](#)," and "[Message Store Message Expiration](#)."

Expire Criteria

The expire criteria can be set with **configutil** options or in a file called **store.expirerule**. Here are some of the criteria that can be specified:

- Folder pattern
- Number of messages in the mailbox
- Total size of the mailbox
- Age, in days, that messages have been in the mailbox. This criterion dates the age of a message from the time it first arrives in the message store (is first received by the user).
- Age, in days, that messages have been saved in a particular mailbox folder. This criterion dates the age of a message from the time it is moved to a particular folder or re-saved in a folder.
- Size of message and grace period (days that a message exceeding the size limit will remain in the message store before removal)
- Whether a message has been flagged as **seen** or **deleted**
- By message header field such as subject or message ID
- According to a sieve script, as defined in RFC 3028

You can use **imexpire** to install a local expire rule file (**store.expirerule**) without conflicting with existing expire rules. If an expire rule file configured for the same partition or mailbox is executing while you try to install a new expire rule file, a warning message appears and the new expire rule file is not installed. Use the **imexpire -i** option to install a local expire rule file.

You can exclude a particular user or mailbox folder from all expire criteria by setting the *exclusive* expire rule for that user or mailbox without specifying any other rules in the expire rule file.

Note: The functionality of **imexpire** has been expanded and the interface has changed since earlier versions of Messaging Server. However, this version continues to support older **imexpire** configurations.

Requirements

imexpire must run locally on the Messaging Server; the **stored** utility must also be running.

Location

The location of **imexpire** is *MessagingServer_home/bin*.

Syntax

```
imexpire [-n] [-d] [-v <num>] [-p <partition> | -u <user>] [-t <num>][-r <num>]
[-m <num>] [-f <file>]
```

```
imexpire -i {-p <partition> | -x <mailbox>} -f <file>
```

Options

The options for this command are:

Table 75–9 imexpire Command Options

Option	Description
-f file	Use the expire rules specified in <i>file</i> . All other expire rules are ignored. When used with the -i option, -f file refers to the expire-rule file to be installed. Use a full path name to specify <i>file</i> . The expire rules in <i>file</i> must use the same format as the rules in the global expire configuration file.
-i	Install a local expire-rule configuration file. This option must be used with either the -p partition option to specify a message store partition or the -x mailbox option to specify a mailbox. In addition, it must be used with the -ffile option to specify the expire rule file to install.
-n	Trial run only - do not perform expire. A description of what would happen without this flag is output.
-v 0 1 2	Log expire statistics. The number specifies the log level, where 0 = store level statistics (default setting) 2 = mailbox level statistics 3 = message level statistics Messages are logged to the log file by default. When the -d option is used, messages go to stdout .

Table 75–9 (Cont.) imexpire Command Options

-d	Display debug output to stdout/stderr .
-p <i>message_store_partition</i>	Expire the specified message store partition.
-u <i>user</i>	Expire the specified user.
-t <i>num</i>	Maximum number of threads per process. Default is 50 .
-r <i>num</i>	Maximum number of threads per partition. Default is 1 .
-m <i>num</i>	Maximum number of rules in a policy. Default is 128 .
-x <i>mailbox</i>	Name of the mailbox to which the local expire rules apply. For example: user/joe/INBOX . This option is used with the -i option to install a local configuration file that will expire messages in the specified mailbox and its sub-folders.

Examples

Install a local expire rule configuration file for the user `jdoe`. These expire rules will apply to `jdoe`'s memos folder.

```
imexpire -i -x user/jdoe/memos -f /home/jdoe/store.expirerule
```

Additional Functionality

Messaging Server provides the following additional functionality to **imexpire**:

- [Attributes for Spam and Virus Scanning Through an MTA Channel](#)
- [Sieve Body-Test Functionality](#)

Attributes for Spam and Virus Scanning Through an MTA Channel

To facilitate spam and virus scanning through an MTA channel, the following attributes have been added to **imexpire**:

Table 75–10 imexpire Attributes

Attribute	Description
channel	An MTA channel.
rescanhours	Rescan messages that have not been scanned for the specified number of hours.

Sieve Body-Test Functionality

Sieve body-test functionality has been added to the **imexpire** utility. The test can find and perform actions on existing email messages in the Message Store on keywords in the body part.

To enable a Sieve body-test, set **ENABLE_SIEVE_BODY=1** in *option.dat*.

Example usage in **store.expirerule**:

```
folderpattern: *
```

```
sieve: require "body"; body :contains "bug";
```

```
action: discard
```

iminitquota

The **iminitquota** utility reinitializes the quota limit from the LDAP directory and recalculates the total amount of disk space that is being used by the users. It updates the message store **quota.db** database under the **mboxlist** directory in the message store. The **iminitquota** utility should be run after the **reconstruct -q** utility is run.

Location: *MessagingServer_home/bin/*

Syntax

```
iminitquota -a | -q | -s | -u <userid>
```

Options

The options for this command are:

Table 75–11 iminitquota Options

Option	Description
-a	Initializes and updates the quota files for every message store user.
-q	Initializes quota only.
-s	Initializes quota, usage, and overquota status.
-u <i>userid</i>	Reinitializes and updates the quota-related information for the specified user. The <i>userid</i> option specifies the message store id of a user, not the login id of the user.

You must specify either the **-a** or **-u** option with the **iminitquota** command.

immonitor-access

Monitors the status of Messaging Server components-Mail Delivery (SMTP server), Message Access and Store (POP and IMAP servers), Directory Service (LDAP server) and HTTP server. This utility measures the response times of the various services and the total round trip time taken to send and retrieve a message. The Directory Service is monitored by looking up a specified user in the directory and measuring the response time. Mail Delivery is monitored by sending a message (SMTP) and the Message Access and Store is monitored by retrieving it. Monitoring the HTTP server is limited to finding out weather or not it is up and running.

The internal operation of **immonitor-access** is as follows: first it does an ldapsearch of a test user created by the administrator. This checks the Directory Server. It can then connect to the SMTP port and send a message to the mail address to check the dispatcher. Then, it checks Message Access by using the IMAP and POP server to see if the message made it to the Message Store. The command logs a message in the default log file if any of the thresholds are exceeded.

The command creates a report that contains the following information:

- The state of the components
- The response time
- The round-trip time for that service

immonitor-access is typically run by **cron** at scheduled intervals to provide a snapshot of the status of the Message Access and Store components. **immonitor-access** can also

connect to the IMAP/POP service and delete messages with the subject specified by **-k**. If **-k** is not specified, all messages containing the subject header, **immonitor**, are deleted.

The administrator must create a test user for use by this command before it can be executed.

Syntax

```
immonitor-access -C <LMTP_host>[:<port>]= [<threshold>][,STLS|PORT] -u <user_name>
  [-D <threshold>] [-m <test_file>] [-k <subject>] [-dhnv]
  [-A <alert_host>] [-f <alert_from>] [-r <alert_recipients>] [-X] [-T]
```

```
immonitor-access -I <IMAP_host>[:<port>]= [<threshold>][,STLS|PORT]
  -u <user_name> [-w passwd | -j pwdfile]
  [-D <threshold>] [-k <subject>] [-dhnvz]
  [-A <alert_host>] [-f <alert_from>] [-r <alert_recipients>] [-X] [-T]
```

```
immonitor-access -H <HTTP_host>[:<port>]= [<threshold>][,STLS|PORT] [-c
<cookie_name>]
  -u <user_name> [-w passwd | -j pwdfile]
  [-D <threshold>] [-dhnv]
  [-A <alert_host>] [-f <alert_from>] [-r <alert_recipients>] [-X] [-T]
```

```
immonitor-access -L <LDAP_host>[:<port>]= [<threshold>][,STLS|PORT] -u <user_name>
  [-b searchbase -B bindDN] [-w passwd | -j pwdfile]
  [-D <threshold>] [-r <alert_recipients>] [-A <alternate_host>]
  [-f <alert_from>] [-X] [-T] [-dhnv]
```

```
immonitor-access -P <POP_host>[:<port>]= [<threshold>][,STLS|PORT]
  -u <user_name> [-w passwd | -j pwdfile]
  [-D <threshold>] [-k <subject>] [-dhnvz]
  [-A <alert_host>] [-f <alert_from>] [-r <alert_recipients>] [-X] [-T]
```

```
immonitor-access -S <SMTP_host>[:<port>]= [<threshold>][,STLS|PORT] -u <user_name>
  [-D <threshold>] [-m <test_file>] [-k <subject>] [-dhnv]
  [-A <alert_host>] [-f <alert_from>] [-r <alert_recipients>] [-X] [-T]
```

Options

The following list contains valid task options for the command.

Table 75–12 immonitor-access Command Options

Option	Description
-u <i>user_name</i>	The valid test user account to use. This test mail user has to be created by the administrator. If the test mail user is in a hosted domain, <i>user@domain</i> should be specified.
-w <i>passwd</i>	The password corresponding to the user specified with <i>u</i> . To read the input from standard input, "-" can be specified with -w . ANY PASSWORD SPECIFIED ON THE COMMAND LINE WITH -w WILL BE VISIBLE TO OTHER USERS OF THE SYSTEM. Use of the -j option is strongly encouraged. This option, or -j , is mandatory when the -H , -I , or -P options are used, or when -n is used in conjunction with -L .

Table 75-12 (Cont.) immonitor-access Command Options

-j <i>pwdfile</i>	A readable file containing the password corresponding to the user specified with -u . This option, or -w , is mandatory when the -H , -I , or -P options are used, or when -n is used in conjunction with -L . This might be available in a future release.
-C <i>LMTP_host</i> : [<i>port</i>] = [<i>threshold</i>]	Use the LMTP server and the port specified to check if Messaging Server is able to deliver the message to the store. The threshold is specified in seconds.
-H <i>HTTP_host</i> : [<i>port</i>] = [<i>threshold</i>]	Use the HTTP server and the port specified to check if the HTTP server is able to accept requests on the specified port. When -I -H or -P is used, it is necessary to provide the test user password with -w . When -S/-C, -I/-P are specified together, the command does the following: <ul style="list-style-type: none"> + sends mail and retrieves with IMAP and POP + reports the per protocol response time + reports round-trip time o reports delivery time (the time taken to send the mail and be visible to IMAP/POP) Multiple -I, -P , and -S options can be specified, which helps in monitoring Messaging Server on various systems.
-I <i>IMAP_host</i> : [<i>port</i>] = [<i>threshold</i>]	Use the IMAP server and the port specified to check the IMAP component of the Message Access. The threshold is specified in seconds. The threshold involves the time to login, retrieve, and delete the message.
-L <i>LDAP_host</i> : [<i>port</i>] = [<i>threshold</i>]	Use the LDAP server and the port specified to check the Directory Server. The threshold is specified in seconds.
-P <i>POP_host</i> : [<i>port</i>] = [<i>threshold</i>]	Use the POP server and the port specified to check the POP component of the Message Access. The threshold is specified in seconds. The threshold involves the time to login, retrieve, and delete the message.
-S <i>SMTP_host</i> : [<i>port</i>] = [<i>threshold</i>]	Use the SMTP server and the port specified to check if Messaging Server is able to accept mail for delivery. The threshold is specified in seconds.
-A <i>alert_host</i>	The SMTP server to send e-mail alerts to. This option helps in sending alert messages even when the primary mail server is down or heavily loaded. If -A is not specified, the SMTP server on the localhost is used.
-b <i>searchbase</i>	Use search base as the starting point for the searching in the Directory Server. It is the same as -b of ldap-search(1) . If -b is not specified, the utility uses the value of dcRoot of the configuration option local.ugldapbasedn .
-B <i>bindDN</i>	LDAP DN to bind as when performing an LDAP search via -L . If not specified, then the value of configuration option local.ugldapbinddn is used. This option as well as -j or -w is mandatory when the -n option is used.
-d	The debug mode: display the execution steps.
-D <i>threshold</i>	The delivery (also called round-trip time) threshold. The time taken to send the mail and the mail being visible to POP/IMAP. This option can be used only when -I/-P and -S/-C are used.
-f <i>alert_from</i>	When immonitor-access sends out an e-mail, it usually is sent as root@domainname . Specify this option to send out an e-mail as different user: -f user@red.iplanet.com
-h	Prints command usage syntax.

Table 75–12 (Cont.) immonitor-access Command Options

-k <i>subject</i>	Header subject of the messages to be sent/deleted. The utility, by default, uses the string "immonitor:<date>" as the subject in the header sent out with the -S option. If -k is specified, the string "immonitor:subject" is used in the subject header. This option can be used with -z to delete messages, if -k is not specified, all messages with the Subject header containing "immonitor" are deleted.
-m <i>test_file</i>	The file that is mailed to the test user. You can get response and round-trip times for various mail sizes with this option. Specify only text files as non-text files result in unexpected behavior. If -m is not specified, the mailfile.txt file in MessagingServer_home/lib/locale/C/mailfile.txt is used as the mail file.
-n	Operate without a Messaging Server configuration. Useful when running the utility on a remote system. This might be available in a future release.
-r <i>alert_recipients</i>	A comma-separated list of mail recipients who will be notified. If this option is not specified, the command reports the alert messages on the standard output.
-v	Run in verbose mode, with diagnostics written to standard output.
-z	Delete messages containing the string specified by -k in the subject header. If -k is not specified, all messages with the subject header containing "immonitor" are deleted. Use -z only with -I or -P . Do not use -z with -S or -C as this can cause unexpected results.
-x	Enable SASL.
-T	Enable SSL.

The default ports are:

SMTP = 25

IMAP = 143

POP = 110

LDAP = 389

LMTP = 225

HTTP = 80

If either the port or threshold is not specified, default ports with the default threshold of 60 seconds is assumed. The threshold specified can be a decimal number.

Output

The command generates a report containing the various protocol execution times. For example:

```
SmtP Statistics for: thestork:25
Connect Time: 2.122 ms
Greeting Time: 5.729 ms
Helo Time: 2.420 ms
Mail From: Time: 2.779 ms
Rcpt To: Time: 4.128 ms
```

```
Data Time: 1.268 ms
Sending File Time: 94.156 ms
Quit Time: 0.886 ms
Total SMTP Time: 113.488 Milliseconds
```

If the alert recipients are specified and any of the threshold values are exceeded, the command mails the report containing the service name and the response time:

ALERT: <service> exceeds threshold Response time=secs/Threshold=secs

Note that in case of times reported for IMAP, the individual times might not add up to the exact value shown by the "Total IMAP time". This occurs because the message does not get to the store immediately. The utility loops until the message is found. Typically, the search time indicates only the successful search time. However, the total time includes each of the individual sleep and search times.

With POP, the utility needs to login and logout multiple times before the message is actually found in the store. Thus, the total time here is the accumulated time for all the logins and log outs.

Example 1: To monitor the SMTP, IMAP, and POP with the threshold 250 milliseconds more than the default value (60 seconds) on localhost use:

```
immonitor-access -S localhost:=60.25 -I localhost:=60.25 -P localhost:=60.25 -u
test_user -w passwd
```

This example assumes that **test_user** exists with password "**passwd**."

Example 2: To monitor LDAP on localhost with no Messaging Server configuration use:

```
immonitor-access -n -L 127.0.0.1:389 -b o=usergroup -u test_user -B "cn=directory
manager" -w passwd
```

The above example assumes that the Directory Manager password is **passwd** and that there is a user named **test_user**.

Example 3: To check remote connectivity and latency issues on **remotehost** use:

```
immonitor-access -S remotehost:=10 -I remotehost:=10 -P remotehost:=10 -u tester
-w passwd
```

Example 4: To run **immonitor-access** on a normal port use:

```
immonitor-access -u admin -w password -P host:110
```

Example 5: To run **immonitor-access** using STLS:

```
immonitor-access -u admin -w password -P host:110 -X
immonitor-access -u admin -w password -P host:110=STLS -X
```

Example 6: To run **immonitor-access** using SSL via STLS:

```
immonitor-access -u admin -w password -P host:110=STLS -T -X
```

Example 7: To run **immonitor-access** on the SSL port:

```
immonitor-access -u admin -w password -P host:995 -T
immonitor-access -u admin -w password -P host:995=PORT -T
```

Example 8: To run **immonitor-access** using SASL on SSL:

```
immonitor-access -u admin -w password -P host:995 -T -X
immonitor-access -u admin -w password -P host:995=PORT -T -X
```

Exit Status

The exit status is 0 if no errors occur. Errors result in a non-zero exit status and a diagnostic message being written to standard error. A different exit status is returned when various thresholds are exceeded.

Table 75–13 *Exit Status Description*

Status	Description
0	Successful execution with no errors or thresholds exceeded
1	Exceeded threshold of a service
2	Errors
64	Usage errors
75	Insufficient virtual memory

An alert message is written to the console when the response time of any server exceeds the threshold.

An error message is written to the console when any of the servers cannot be reached.

Warnings

The password passed with **-w** can be visible to a user using the **ps(1)** command. It is strongly advised that you create a test user to be specifically used by the monitoring utilities.

It is recommended that you use **-w** and enter the password through standard input. However, if the utility is executed through **cron**, the password can be stored in a file. This file can be redirected as the standard input for the utility.

```
cat passwd_file | immonitor-access -w -
immonitor-access -w - ... < passwd_file
```

Do not use the echo command such as:

```
echo password | immonitor-access .. -w - ..
```

because the **ps** might show the echo's arguments.

To delete the test mail sent by the **-S** option, invoke the **immonitor-access** command with the **-z** option separately. Do not use the two together.

impurge

The **impurge** command can be used to manually purge unused cache records and message files in mailboxes when the purge server daemon process is not running (**local.purge.enable** is disabled). For more information see the *Messaging Server Reference*. The **impurge** command has the following options:

Requirements: Must be run locally on the Messaging Server.

Location: *MessagingServer_home/lib/*

Syntax

```
impurge [-d] [-r hours] [-e]
```

Options

The options for this command are:

Table 75–14 impurge Options

Option	Description
-d	Enables debug mode
-r hours	Exists after the specified number of hours
-e	Exits when the work queue is empty

If both **-r** and **-e** are specified, **impurge** exits when the queue is empty or time has expired.

Only one **impurge** process can be executed at a time. Attempting to run the **impurge** command whilst the purge server daemon is running will result in the following error message in the Messaging Server debug logs:

```
[24/Feb/2009:14:43:59 +1100] hostname impurge[17471]: General Error: Could not get
purge session lock.
Possibly another impurge is running
```

imquotacheck

The **imquotacheck** utility administers user quotas and domain quotas in the message store. This utility can also compare mailbox size with a user's assigned quota. As an option, you can email a notification to users who have exceeded a set percentage of their assigned quota.

The **imquotacheck** utility lists users in the local mboxlist database. To list users in the LDAP directory, use the **imquotacheck -a** option.

Requirements: Must be run locally on the Messaging Server.

Location: *MessagingServer_home/bin/*

Syntax

To report quota and usage information to **stdout**:

```
imquotacheck [-u <user> | -d <domain>] [-p <partition>] | -p <partition> ] [-a]
```

To enforce a domain quota:

```
imquotacheck -f [-d <domain>]
```

To notify users when they have exceeded a set percentage of their assigned quota (to deliver over-quota notification messages):

```
imquotacheck -n [-e] [-d <domain>] [-r <rulefile>] [-t <message template>] [-l]
```

Options

The options for this command are:

Table 75–15 imquotacheck Options

Option	Description
-a	Lists users in the LDAP directory. If the -a option is not used, imquotacheck lists users in the message store (the local mboxlist database).
-d domain	Searches for users only in the specified domain. When used with the -f option, the -d <domain> option enforces quotas for the specified domain. When used with the -n option, the -d <domain> option delivers over-quota notifications for users in the specified domain.
-f	Enforces domain quotas. If the domain is over quota and the mailDomainStatus attribute is currently set to active, the value will be reset to overquota, which will prevent mail from being accepted by the message store. If the domain is not over quota and the mailDomainStatus attribute is set to overquota, then the value will be changed to active, and mail will be accepted.
-h	Displays help for this command.
-n	Sends notification messages based on the rules defined in the <rulefile>. If you do not define any rules with the -r option when you use the -n option, the default <rulefile> is used. If you have not set up a default <rulefile> and you do not define any rules when you use the -n option, you do not receive an error. Use of imquotacheck -n is mutually exclusive with dynamic notification (that is, when store.quotanotification is enabled). See "Methods of Notification."
-p partition	Reports quota for the specified partition.
-r rulefile	Specifies the set of rules to be used when you want to calculate quota usage. This option is used with the -n option. If -r is not specified, a default <rulefile> can be used. To set up a default <rulefile>, copy the "Sample Rulefile" to <MessagingServer_home>/config/imq.rulefile. See "Rulefile Format."
-t message template	Notifies users when their mailbox quota is exceeded. This option is used with the -n option. The message template format is the following: * %U% - user's mailbox id * %Q% - percentage of the used mailbox quota * %R% - quota usage details: assigned quota, total mailbox size, and percentage used. If the -e is specified, mailbox usage of the individual folders are also reported. * %M% - current mailbox size * %C% - quota attribute value. If -t is not specified, a default message file will be mailed. To setup a default message file, copy the "Notification File" to <MessagingServer_home>/config.
-u user	Obtains the quota status of the specified user id. Cannot be used to specify multiple users.
-e	Report disk usage of the individual folders in notification messages.
-l	Log users to action log who are above quota.

Examples

To send a notification to all users in accordance to the default **rulefile**:

```
imquotacheck -n
```

To send a notification to all users in accordance to a specified **rulefile**, **myrulefile**, and to a specified mail template file, **mytemplate.file**:

```
imquotacheck -n -r myrulefile -t mytemplate.file
```

To enforce the domain quota for the **example.com** domain:

```
imquotacheck -f -d example.com
```

To list the usage of all users whose quota exceeds the least threshold in the **rulefile**:

```
imquotacheck
```

To list per folder usages for user **user1**:

```
imquotacheck -u user1
```

To only list the users in domain **example.com**:

```
imquotacheck -d example.com
```

The **imquotacheck** command has stable output formats that we support and allow customers to parse.

Sample **imquotacheck** output:

Name	Quota(K)	Usage(K)	%	Quota#	Usage#	%	OverDate	WarnDate
joe	-	6	-	-	1	-	-	-
mary	100	110	110	-	49	-	10/14/15	10/14/15
big	52428800	420	0	-	1673	-	-	-
small	-	1014B	-	-	3	-	-	-

Rulefile Format

The **rulefile** format is organized into two sections: a general section and a rule name section. The general section contains attributes that are common across all rules.

Attributes that are typically specified in the general section are the **mailQuotaAttribute** and the **reportMethod**. In the rule name section, you can write specific quota rules for notification intervals, trigger percentages, and so on. Attributes that are typically specified in the rule name section are **notificationTriggerPercentage**, **enabled**, **notificationInterval**, and **messageFile**. Note that the attributes and corresponding values are not case-sensitive. The following rulefile format is used:

```
[General]
mailQuotaAttribute = [value]
reportMethod = [value]

[rulename1]
attrname=[value]
attrname=[value]

[rulename2]
attrname=[value]
attrname=[value]

[rulename3]
attrname=[value]
attrname=[value]
```

[Table 75–16](#) shows the attributes, whether they are required, the default value, and the description.

Table 75–16 *rulefile Attributes*

Attribute	Required Attribute ?	Default Value	Description
mailQuotaAttribute	No	Value in quotadb	Specifies the name of the custom mailquota attribute. If not specified, the value in quotadb is used.
reportMethod	No	Not Applicable	You can provide your own code to customize the output of the quota report. The value of this attribute is specified as <code><library-path:function></code> , where <code><library-path></code> is the path of the shared library and <code><function></code> is the name of the report function. See " reportMethod Signature " for more information about what options your function must accept and return.
notificationTrigger Percentage	Yes	Not Applicable	Specifies the consumed quota percentage that will trigger notification. Value should be unique and an integer.
messageFile	No	<code><MessagingServer_home>/config/imq.msgfile</code>	Specifies the absolute path to the message file.
notificationInterval	Yes	Not Applicable	Indicates the number of hours before a new notification is generated.
enabled	No	0 (FALSE)	Indicates if the particular rule is active. Applicable values are 0 for FALSE and 1 for TRUE.
notificationMethod	No	Not Applicable	You can provide your own code to perform the overquota notification. The value of this attribute is specified as <code><library-path:function></code> , where <code><library-path></code> is the path of the shared library and <code><function></code> is the name of the report function. See " notificationMethod Signature " for more information about what options your function must accept and return.

reportMethod Signature

If you override the `imquotacheck reportMethod()` with your own function, it must be defined as:

```
int symbol(QuotaInfo* info, char** message, int* freeflag)

typedef struct QuotaInfo {
    const char* username; /* user name (uid or uid@domain) */
    long quotakb; /* quota in kbytes */
    long quotams; /* quota in number of messages */
    ulong usagekb; /* total usage in kbytes */
    ulong usagemsg; /* total usage in number of messages */
    FolderUsage* folderlist; /* folder list (for -e) */
    long num_folder; /* number of folders in the folderlist */
    long trigger; /* not used */
    const char* rule; /* not used */
}

typedef struct FolderUsage {
    const char* foldername;
    ulong usagekb; /* folder usage in kbytes */
}
```

```
}

```

The address, **message**, points to the output message. The report function is expected to fill the value of ***message** and allocate memory for **message** when necessary. The **freeflag** variable indicates if the caller is responsible for freeing allocated memory for ***message**.

The return values are 0 for success and 1 for failure.

The **imquotacheck** function will invoke the **reportMethod** to generate the report output. If the **reportMethod** returns 0 and ***message** is pointing to a valid memory address, **message** will be printed.

If the ***freeflag** is set to 1, the caller will free the memory address pointed to by **message**. If the **-e** option is specified, the quota usage for every folder will be stored in the **folderlist**, an array in **FolderUsage**; the **num_folder** variable is set to the number of folders in the **folderlist**.

notificationMethod Signature

If you override the **imquotacheck notificationMethod()** with your own function, it must be defined as:

```
int symbol(QuotaInfo* info, char** message, int* freeflag)

typedef struct QuotaInfo {
    const char* username; /* user name (uid or uid@domain) */
    long quotakb; /* quota in kbytes */
    long quotams; /* quota in number of messages */
    ulong usagekb; /* total usage in kbytes */
    ulong usagemsg; /* total usage in number of messages */
    FolderUsage* folderlist; /* folder list (for -e) */
    long num_folder; /* number of folders in the folderlist */
    long trigger; /* the exceeded notificationTriggerPercentage */
    const char* rule; /* rulename that triggered notification */
}

typedef struct FolderUsage {
    const char *foldername;
    ulong usagekb; /* folder usage in kbytes */
}
```

The address, **message**, points to the notification message. The notification function is expected to fill in the value of this variable and allocate the memory for the message when necessary. The **freeflag** variable indicates if the caller is responsible for freeing the memory allocated for **message**.

The return values are 0 for success and 1 for failure.

If the notification function returns a 0, and ***message** is pointing to a valid address, the **imquotacheck** utility will deliver the message to the user. If the ***freeflag** is set to 1, the caller will free the memory address pointed to by **message** after the message is sent.

If the **-e** option is specified, the quota usage for every folder will be stored in the **folderlist** variable, an array of **FolderUsage** structure; the **num_folder** variable is set to the number of folders in the **folderlist**.

Note: If the **messageFile** attribute is also specified, the attributed **messageFile** will be ignored.

Sample Rulefile

In the sample rulefile, the following files **libzz.so**, **libzz.sl**, and **libzz.dll** are library files. The */xx/yy* are directory paths that should be replaced by the relative paths to where these library files are located on the server.

```
#
# Sample rulefile
#
[General]
mailQuotaAttribute=mailquota
reportMethod=/xx/yy/libzz.so:myReportMethod [for Solaris only ]

[rule1]
notificationTriggerPercentage=60
enabled=1
notificationInterval=3
notificationMethod=/xx/yy/libzz.so:myNotifyMethod_60

[rule2]
notificationTriggerPercentage=80
enabled=1
notificationInterval=2
messageFile=/xx/yy/message.txt

[rule3]
notificationTriggerPercentage=90
enabled=1
notificationInterval=1
notificationMethod=/xx/yy/libzz.so:myNotifyMethod_90
#
# End
#
```

Threshold Notification Algorithm

1. Rule precedence is determined by increasing trigger percentages.
2. The highest applicable threshold is used to generate a notification. The time and the rule's threshold are recorded.
3. If users move into a higher threshold since their last quota notification, a new notification will be delivered based on the current set of applicable rules. This notice can be immediately delivered to any user whose space usage is steadily increasing.
4. If usage drops, the notification interval of the current rule (lower threshold) will be used to check the time elapsed since the last notice.
5. The stored time and threshold for the user will be reset to zero if the user's mailbox size falls below all of the defined thresholds.

Notification File

The utility depends on the message file to have at minimum a Subject Header. There should be at least one blank line separating the Subject from the body. The other required headers are generated by the utility. The notification file format is the following:

```
Subject: [Warning] quota reached for %U%
```

```

Hello %U%,
Your quota: %C%
Your current mailbox usage: %M%
Your mailbox is now %Q% full. The folders consuming the most space are:
%R%.

```

Please clean up unwanted diskspace.

```

Thanks,
-Administrator

```

where:

%U% - Species the user ID.

%Q% - Specifies the percentage of the mailbox quota used.

%R% - Specifies quota usage details, including assigned quota, total mailbox size, and percentage used. When -e is specified on the command line, the report shows the mailbox usage of the individual folders.

%M% - Specifies the current mailbox size.

%C% - Specifies quota attribute value.

Note: If an account has less than 1KB of quota usage, **imquotacheck** prints out the usage in bytes (B) rather than kilobytes (KB) shown in the heading.

```

imquotacheck -u testquota
Name      Quota(K) Usage(K) %   Quota#  Usage#  %   OverDate WarnDate
-----
testquota 256000  654B  0   100000   1     0     -     -
-----

```

Note: Localized versions of **imquotacheck** notification incorrectly convert the % and the \$ signs. To correct the encoding, replace every \$ with \24 and replace every % with \25 in the message file.

imsasm

The **imsasm** utility is an external ASM (Application Specific Module) that handles the saving and recovering of user mailboxes. **imsasm** invokes the **imsbackup** and **imsrestore** utilities to create and interpret a data stream.

During a save operation **imsasm** creates a save record for each mailbox or folder in its argument list. The data associated with each file or directory is generated by running the **imsbackup** or **imsrestore** command on the user's mailbox.

Location: *MessagingServer_home/lib/msg*

Syntax

```
imsasm [<standard_asm_arguments>]
```

Options

The options used in the **imsasm** utility are also known as standard-asm-arguments, which are Legato NetWorker backup standards.

Either **-s** (saving), **-r** (recovering), or **-c** (comparing) must be specified and must precede any other options. When saving, at least one *path* argument must be specified. *path* may be either a directory or filename.

The following options are valid for all modes:

Table 75–17 *imsasm* Options

Option	Description
-n	Performs a dry run. When saving, walk the file system but don't attempt to open files and produce the save stream. When recovering or comparing, consume the input save stream and do basic sanity checks, but do not actually create any directories or files when recovering or do the work of comparing the actual file data.
-v	Turns on verbose mode. The current ASM, its arguments, and the file it is processing are displayed. When a filtering ASM operating in filtering mode (that is, processing another ASM's save stream) modifies the stream, its name, arguments, and the current file are displayed within square brackets.

When saving (**-s**), the following options may also be used:

Table 75–18 *imsasm* Saving Options

Option	Description
-b	Produces a byte count. This option is like the -n option, but byte count mode will estimate the amount of data that would be produced instead of actually reading file data so it is faster but less accurate than the -n option. Byte count mode produces three numbers: the number of records, i.e., files and directories; the number of bytes of header information; and the approximate number of bytes of file data. Byte count mode does not produce a save stream so its output cannot be used as input to another asm in recover mode.
-o	Produces an "old style" save stream that can be handled by older NetWorker servers.
-e	Do not generate the final "end of save stream" Boolean. This flag should only be used when an ASM invokes an external ASM and as an optimization chooses not to consume the generated save stream itself.
-i	Ignores all save directives from .nsr directive files found in the directory tree.
-f proto	Specifies the location of a .nsr directive file to interpret before processing any files. Within the directive file specified by <i>proto</i> , <i>path</i> directives must resolve to files within the directory tree being processed, otherwise their subsequent directives will be ignored.
-p ppath	Prepends this string to each file's name as it is output. This argument is used internally when one ASM executes another external ASM. <i>ppath</i> must be a properly formatted path which is either the current working directory or a trailing component of the current working directory.

Table 75–18 (Cont.) imsasm Saving Options

Option	Description
-t <i>date</i>	The date after which files must have been modified before they will be saved.
-x	Crosses file system boundaries. Normally, file system boundaries are not crossed during walking.

When recovering (**-r**), the following options may also be used:

Table 75–19 imsasm Recovering Options

Option	Description
-i <i>response</i>	Specifies the initial default overwrite response. Only one letter may be used. When the name of the file being recovered conflicts with an existing file, the user is prompted for overwrite permission. The default response, selected by pressing Return , is displayed within square brackets. Unless otherwise specified with the -i option, n is the initial default overwrite response. Each time a response other than the default is selected, the new response becomes the default. When either N , R , or Y is specified, no prompting is done (except when auto-renaming files that already end with the rename suffix) and each subsequent conflict is resolved as if the corresponding lower case letter had been selected. The valid overwrite responses and their meanings are: <ul style="list-style-type: none"> ▪ n-Do not recover the current file. ▪ N-Do not recover any files with conflicting names. ▪ y-Overwrite the existing file with the recovered file. ▪ Y-Overwrite files with conflicting names. ▪ r-Rename the conflicting file. A dot "." and a suffix are appended to the recovered file's name. If a conflict still exists, the user will be prompted again. ▪ R-Automatically renames conflicting files by appending a dot "." and a suffix. If a conflicting file name already ends in a <i>suffix</i>, the user will be prompted to avoid potential auto rename looping conditions.
-m <i>src=dst</i>	Maps the file names that will be created. Any files that start exactly with <i>src</i> will be mapped to have the path of <i>dst</i> replacing the leading <i>src</i> component of the path name. This option is useful if you wish to perform relocation of the recovered files that were saved using absolute path names into an alternate directory.
-z <i>suffix</i>	Specifies the suffix to append when renaming conflicting files. The default suffix is R .
<i>path</i>	Restricts the files being recovered. Only files with prefixes matching <i>path</i> will be recovered. This checking is performed before any potential name mapping is done with the -m option. When <i>path</i> is not specified, no checking is performed.

Examples

To use **imsasm** to save the mailbox **INBOX** for user **joe**, the system administrator creates a directory file *backup_root/backup/DEFAULT/joe.nsr* with the following contents:

```
imsasm: INBOX
```

This causes the mailbox to be saved using **imsasm**. Executing the **mkbackupdir** utility will automatically create the **.nsr** file. See "[mkbackupdir](#)."

imsbackup

The **imsbackup** utility is used to write selected contents of the message store to any serial device, including magnetic tape, a UNIX pipe, or a plain file. The backup or selected parts of the backup may later be recovered via the **imsrestore** utility. The **imsbackup** utility provides a basic backup facility similar to the UNIX **tar** command.

When **imsbackup**, **imsrestore**, **imsimport** or any processing intensive operation takes significantly more system resources than normal, and continues doing so longer than the **msprobe** interval, there may be a temporary backlog of DB transaction log files to be cleared. If there are more files than specified in **local.store.maxlog**, then **msprobe** may erroneously restart all the processes during a restore. To prevent this from happening, disable **msprobe** during the **imsbackup**, **imsrestore**, and **imsimport**.

Location: *MessagingServer_home/bin*

For more information about **imsbackup** and backing up the message store, see "[Backing Up and Restoring the Message Store](#)" in the *Messaging Server System Administrator's Guide*.

Syntax

```
imsbackup -f file [-b <blockfactor>] [-d <date>] [-m <link_count>]
[-e <encoding>] [-u <file>] [-n <path>] [-ivlgrp] <name>...
```

Options

The options for this command are:

Table 75–20 *imsbackup Options*

Option	Description
-b <i>blockfactor</i>	Everything written to the backup device is performed by blocks of the size 512 x <i>blockfactor</i> . The default is 20 .
-d <i>date</i>	Date from which messages are to be backed up, expressed in YYYYMMDD [:HHMMSS]; for example, -d 19990501:131000 would backup messages stored from May 1, 1999 at 1:10 pm to the present. The default is to back up all the messages regardless of their dates.
-e <i>encoding</i>	Mailbox name encoding (example: IMAP-mailbox-name)
-f <i>file</i>	Specifies the file name or device to which the backup is written. If file is "-", backup data is written to stdout .
-g	Debug mode. The output is written in the default log file not to the stdout . For the stdout , one should use -v .
-i	Ensures that all messages are backed up for a partial restore. Backs up every message multiple times if necessary.
-l	Used to autoloading tape devices when end-of-tape is reached.
-m <i>link_count</i>	Specifies the minimum link count for hashing.
-n <i>path</i>	Specifies the networker backup path.
-r	Specifies read-only mode (does not upgrade or repair the mailbox automatically).

Table 75–20 (Cont.) imbackup Options

-u <i>file</i>	Specifies a backup object file. This file contains the object names (entire message store, user, group, mailbox, and so on) to backup. See <i>name</i> for a list of backup object
-v	Executes the command in verbose mode.
-x	Backs up expunged messages.
<i>name</i>	Can be 1) logical pathname of the backup object, 2) user ID, 3) message store mailbox name. Backup objects and paths: <ul style="list-style-type: none"> ■ Entire message store: / ■ Message store partition: /<i>partition_name</i> (default: /primary) ■ Backup group---a group of users defined with regular expressions in a configuration file. See "To Create Backup Groups" for details. Path: /<i>partition_name/backup_group</i> (/primary/user represents all users under primary). ■ User: /<i>partition_name/backup_group/user_ID</i> ■ Mailbox: /<i>partition_name/backup_group/user_ID/mailbox_name</i> ■ Message: /<i>partition_name/backup_group/user_ID/mailbox_name/msgIDuser_IDs</i>: can be any user ID in the message store. If the user is not in the default domain, the user ID must be fully qualified (example: Wally@example.com). If user is in the default domain, the user ID can stand alone (example: <i>Wally</i>). ■ <i>mailbox</i>: An email folder. It is specified using the following message store internal name: user/user_ID/folder_name. Note that user is a message store keyword.

Examples

The following example backs up the entire message store to /dev/rmt/0:

```
imsbackup -f /dev/rmt/0 /
```

The following backs up the mailboxes of user ID **joe** to /dev/rmt/0:

```
imsbackup -f /dev/rmt/0 /primary/user/joe
```

The following example backs up all the mailboxes of all the users defined in the backup group **groupA** to **backupfile**:

```
imsbackup -f- /primary/groupA > backupfile
```

imsconnutil

Monitors user access of the message store. **imsconnutil** can provide the following information:

- Users currently logged in on IMAP or any HTTP webmail client.
- The last access time (log in or log out) for a specified user.
- For IMAP: lists the authentication method, the IP address from which the users are logged in, the IP address to which users are connected, and the port on which they are logged to and from.

This command requires root access by the system user, and you must set the configuration variables **local.imap.enableuserlist**, **local.http.enableuserlist**, and **local.enablelastaccess** to 1.

See "Monitoring User Access to the Message Store" for more information and examples.

Location: *MessagingServer_home/bin*

Syntax

```
imsconnutil -c|-a [-s <service>] [-u <uid>] [-f <filename>]
```

```
imsconnutil -k {-u <uid> | -f <filename>}
```

Options

The options for this command are:

Table 75–21 imsconnutil Options

Option	Description
-c -a -k	At least one of -c or -a, or -k must be used.
-a	Last IMAP, POP, or http web mail client access (log in or log out) of user(s). -s does not affect the output of -a.
-c	List IMAPusers currently connected.
-k	Disconnect users from IMAP and POP. Users logged on to Communications Express lose the underlying IMAP connection and, thus, are also disconnected. The -k option must be used with either the -u uid option to specify users to disconnect or the -f filename option to specify the users listed in filename. The -k option requires that IMAP IDLE be configured. See "Configuring IMAP IDLE" for more information.
-s service	Can specify either IMAP or HTTP as service to monitor. Only applies to -c option. POP is not available because POP users do not typically stay logged on.
-u uid	Specify a UID to monitor. If -u and -f are not listed, then all users are monitored.
-f filename	File containing UIDs to monitor. Each UID must be on its own line.

Examples

- List every user ID currently logged into IMAP and http.

#imsconnutil -c

- List last IMAP, POP, or Messenger Express access (log in or log out) of every user ID.

#imsconnutil -a

- List access history (last log off or log on) of all user IDs. Lists current user IDs logged into IMAP and http.

#imsconnutil -a -c

- List IMAP users currently logged on the message store.

#imsconnutil -c -s imap

- Reveal whether user ID George is logged onto IMAP or not.

#imsconnutil -c -s imap -u George

- Reveal whether user ID George is currently logged onto IMAP or Messenger Express, and lists the last time George was logged on or off.

#imsconnutil -c -a -u George

- Disconnect the user ID George.

#imsconnutil -k -u George

Notes

The "rehostuser" utility makes use of the **imsconnutil** command. If you are planning to use the **rehostuser** command, for now, you must use the **configutil** command to explicitly configure **service.imap.ensidle=1**. This is because **imsconnutil** decides to use ENS or JMQ based on the configuration of the **service.imap.ensidle** option. While the **imapd** daemon treats that variable as "on" by default, if you have configured **ibiff**, **imsconnutil** treats the option as "off" by default.

imscripter

The **imscripter** utility connects to an IMAP server and executes a command or a sequence of commands.

May be run remotely.

Location: *MessagingServer_home/bin/*

Syntax

```
imscripter [-h]
imscripter [-s <server>] [-p <port>] [-u <userid>] [-x <password>]
imscripter [-s <server>] [-p <port>] [-u <userid>] [-x <password>] -f <scriptfile>
imscripter [-s <server>] [-p <port>] [-u <userid>] [-x <password>] -c <command>
[<command data>]
imscripter [-s <server>] [-p <port>] [-u <userid>] [-x <password>] -c <command> -f
<command data file>
```

Options

The options for this utility are:

Table 75–22 imscripter Command Options

Option	Description
-c <i>command</i>	Executes <i>command</i> - see discussion of commands below.
-f <i>file</i>	The <i>file</i> may contain one or more commands, or a list of mailboxes on which commands are to be executed. See discussion of commands below.
-h	Displays help for this command.
-p <i>port</i>	Connects to the given port. The default is 143 .
-s <i>server</i>	Connects to the given server. The default is localhost . The server can be either a host name or an IP address.
-u <i>userid</i>	Connects as <i>userid</i> .

Table 75–22 (Cont.) imscripter Command Options

Option	Description
-v <i>verbosity</i>	String containing options for printing various information. The options are as follows: E-Show errorsI-Show informational messages P-Show prompts C-Show input commands c-Show protocol commands B-Show BAD or NO untagged responses O-Show other untagged responses b-Show BAD or NO completion results o-Show OK completion results A-Show all of the above The letters designating options can be entered in any order. The default is EPBibo .
-x <i>passwd</i>	Uses this password.

Supported Commands

The imscripter command supports the following commands:

```
create <mailbox>
delete <mailbox>
rename <oldmailbox> <newmailbox> [<partition>]
getacl <mailbox>
setacl <mailbox> <userid> <rights>
deleteacl <mailbox> <userid>
```

If one or more of the above variables are included, the option executes the given command with that input. For example, **create lincoln** creates a mailbox for the user **lincoln**.

If the **-f file** option is used, the option executes the command on each variable listed in the file. For example:

```
cat folders-to-create
xyz
abc
def
imscripter -u <userid> -x <password> -c create -f folders-to-create
#
```

Raw IMAP Commands

In addition to the commands which imscripter interprets (above), it will pass thru any raw IMAP command prefixed with an exclamation mark ("!"). For example:

```
imscripter -u <userid> -x <password> -c \!list "" \*
1) <= OK [CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ NAMESPACE UIDPLUS
CHILDREN BINARY UNSELECT SORT CATENATE URLAUTH LANGUAGE ESEARCH ESORT
THREAD=ORDEREDSUBJECT THREAD=REFERENCES ENABLE CONTEXT=SEARCH CONTEXT=SORT WITHIN
SASL-IR SEARCHRES XSENDER X-NETSCAPE XSERVERINFO X-SUN-SORT ANNOTATE-EXPERIMENT-1
X-UNAUTHENTICATE X-SUN-IMAP X-ANNOTATEMORE XUM1 ID IDLE AUTH=PLAIN]
s4u-280rd-zone1-bur02.east.sun.com IMAP4 service (Oracle Communications Messaging
Exchange Server 7u4-20.01 64bit (built Nov 21 2010))
1) <= OK User logged in
```

```

2) <= CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ NAMESPACE UIDPLUS CHILDREN
BINARY UNSELECT SORT CATENATE URLAUTH LANGUAGE ESEARCH ESORT THREAD=ORDEREDSUBJECT
THREAD=REFERENCES ENABLE CONTEXT=SEARCH CONTEXT=SORT WITHIN SASL-IR SEARCHRES
XSENDER X-NETSCAPE XSERVERINFO X-SUN-SORT ANNOTATE-EXPERIMENT-1 X-UNAUTHENTICATE
X-SUN-IMAP X-ANNOTATEMORE XUM1 ID IDLE
2) <= OK Completed
2) <= LIST (\NoInferiors) "/" INBOX
2) <= LIST (\HasNoChildren) "/" Drafts
2) <= LIST (\HasNoChildren) "/" Sent
2) <= LIST (\HasNoChildren) "/" Trash
2) <= LIST (\HasNoChildren) "/" abc
2) <= LIST (\HasNoChildren) "/" def
2) <= LIST (\HasNoChildren) "/" test2
2) <= LIST (\HasNoChildren) "/" xyz
2) <= OK Completed
3) <= BYE LOGOUT received
3) <= OK Completed
#

```

Interactive Mode

If you issue the `imscripter` command without the `-f` and/or `-c` switches, it will go into interactive mode. If you did not specify `-u`, it will prompt for `userid`. If you did not specify `-x`, it will prompt for password. You can then type `imscripter` commands or raw IMAP commands prefixed by `!` and see the response. See discussion of commands above.

For example:

```

imscripter
s4u-280rd-zone1-bur02 userid: <userid>
s4u-280rd-zone1-bur02 password for <userid>:
1) <= OK [CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ NAMESPACE UIDPLUS
CHILDREN BINARY UNSELECT SORT CATENATE URLAUTH LANGUAGE ESEARCH ESORT
THREAD=ORDEREDSUBJECT THREAD=REFERENCES ENABLE CONTEXT=SEARCH CONTEXT=SORT WITHIN
SASL-IR SEARCHRES XSENDER X-NETSCAPE XSERVERINFO X-SUN-SORT ANNOTATE-EXPERIMENT-1
X-UNAUTHENTICATE X-SUN-IMAP X-ANNOTATEMORE XUM1 ID IDLE AUTH=PLAIN]
s4u-280rd-zone1-bur02.east.sun.com IMAP4 service (Oracle Communications Messaging
Exchange Server 7u4-20.01 64bit (built Nov 21 2010))
2) <= OK User logged in
2) <= CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ NAMESPACE UIDPLUS CHILDREN
BINARY UNSELECT SORT CATENATE URLAUTH LANGUAGE ESEARCH ESORT THREAD=ORDEREDSUBJECT
THREAD=REFERENCES ENABLE CONTEXT=SEARCH CONTEXT=SORT WITHIN SASL-IR SEARCHRES
XSENDER X-NETSCAPE XSERVERINFO X-SUN-SORT ANNOTATE-EXPERIMENT-1 X-UNAUTHENTICATE
X-SUN-IMAP X-ANNOTATEMORE XUM1 ID IDLE
2) <= OK Completed
imscripter> !select inbox
3) <= FLAGS (\Answered \Flagged \Draft \Deleted \Seen $Forwarded)
3) <= OK [PERMANENTFLAGS (\Answered \Flagged \Draft \Deleted \Seen $Forwarded
\*)]
3) <= EXISTS
3) <= RECENT
3) <= OK [UNSEEN 14]
3) <= OK [UIDVALIDITY 1266949413]
3) <= OK [UIDNEXT 354]
3) <= OK [READ-WRITE] Completed
imscripter> quit
5) <= BYE LOGOUT received
5) <= OK Completed
#

```

Individual Command Mode

As with interactive mode (above), `-s` and `-p` default to localhost and port 143 and `imscripter` will prompt for `-u` and `-x` if not specified.

With `-c`, you can execute an individual `imscripter` command and optionally have it operate on several folders. See the example of creating several folders in the discussion of commands above.

Script Mode

As with interactive mode (above), `-s` and `-p` default to localhost and port 143 and `imscripter` will prompt for `-u` and `-x` if not specified.

With `-f`, you can put `imscripter` and raw IMAP commands in a file and have `imscripter` execute that script. For example:

```
cat script
create nmo
create blah
!list "" *
imscripter -u <userid> -x <password> -f script
#
```

Notice the above displayed no output. It executed the command but did not display the output. Use the `-v` switch it specify which output you want to see. For example:

```
imscripter -u <userid> -x <password> -f script -v A
Processing: script, verbosity: A, server: (<ask user>:143)
connecting to s4u-280rd-zone1-bur02:143...
addcallback CAPABILITY, NO, OK BAD
  1) => LOGIN <userid> * *notice imscripter hid the password*
  1) <= OK [CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ NAMESPACE UIDPLUS
CHILDREN BINARY UNSELECT SORT CATENATE URLAUTH LANGUAGE ESEARCH ESORT
THREAD=ORDEREDSUBJECT THREAD=REFERENCES ENABLE CONTEXT=SEARCH CONTEXT=SORT WITHIN
SASL-IR SEARCHRES XSENDER X-NETSCAPE XSERVERINFO X-SUN-SORT ANNOTATE-EXPERIMENT-1
X-UNAUTHENTICATE X-SUN-IMAP X-ANNOTATEMORE XUM1 ID IDLE AUTH=PLAIN]
s4u-280rd-zone1-bur02.east.sun.com IMAP4 service (Oracle Communications Messaging
Exchange Server 7u4-20.01 64bit (built Nov 21 2010))
  2) <= OK User logged in
  2) => CAPABILITY
  2) <= CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ NAMESPACE UIDPLUS CHILDREN
BINARY UNSELECT SORT CATENATE URLAUTH LANGUAGE ESEARCH ESORT THREAD=ORDEREDSUBJECT
THREAD=REFERENCES ENABLE CONTEXT=SEARCH CONTEXT=SORT WITHIN SASL-IR SEARCHRES
XSENDER X-NETSCAPE XSERVERINFO X-SUN-SORT ANNOTATE-EXPERIMENT-1 X-UNAUTHENTICATE
X-SUN-IMAP X-ANNOTATEMORE XUM1 ID IDLE
  2) <= OK Completed
  3) Cmd: create nmo
  3) => CREATE nmo
  3) <= NO Mailbox already exists *this command failed because it
succeeded the first time*
  4) Cmd: create blah
  4) => CREATE blah
  4) <= NO Mailbox already exists *ditto*
  5) Cmd: !list "" *
  5) => list "" *
  5) <= LIST (\NoInferiors) "/" INBOX
  5) <= LIST (\HasNoChildren) "/" Drafts
  5) <= LIST (\HasNoChildren) "/" Sent
  5) <= LIST (\HasNoChildren) "/" Trash
```

```

5) <= LIST (\HasNoChildren) "/" abc
5) <= LIST (\HasNoChildren) "/" blah
5) <= LIST (\HasNoChildren) "/" def
5) <= LIST (\HasNoChildren) "/" nmo
5) <= LIST (\HasNoChildren) "/" test2
5) <= LIST (\HasNoChildren) "/" xyz
5) <= OK Completed
7) => LOGOUT
7) <= BYE LOGOUT received
7) <= OK Completed
#

```

imsexport

The **imsexport** utility exports Messaging Server folders into UNIX `/var/mail` format folders.

The **imsexport** utility extracts the messages in a message store folder or mailbox and writes the messages to a UNIX file under the directory specified by the administrator. The file name is the same name as the IMAP folder name. For message store folders that contain both messages and sub-folders, **imsexport** creates a directory with the folder name and a file with the folder name plus a `.msg` extension. The `folder.msg` file contains the messages in the folder. The `folder` directory contains the sub-folders.

Location: `MessagingServer_home/bin`

Syntax

```
imsexport [-g] [-v 0|1|2] [-f format] [-s mailbox] [-e encoding] -u user -d dir
```

Options

The options for this command are:

Table 75–23 imsexport Options

Option	Description
-d <i>dir</i>	Specifies the destination directory name where the folders will be created and written. This is a required option.
-e <i>encoding</i>	Specify an encoding option.
-f <i>format</i>	Specifies the export <i>format</i> , where <i>format</i> is Solaris, BSD, SYSV or RN. Default is Solaris. The format option impacts the date format used in the "From " line in the mailbox format. The formats are: %a %b %e %H:%M:%S %Y %z (Solaris) %a %b %e %H:%M:%S %Y (BSD) %a %b %e %H:%M %Z %Y (SYSV) %a %b %e %H:%M:%S %Z %Y (RN)
-g	Specifies debugging mode.
-s <i>folder</i>	Specifies the source folder to export.

Table 75–23 (Cont.) imsexport Options

Option	Description
-u <i>user</i>	Specifies the message store id for a user. Note that this is not necessarily the login id of the user. The message store id is either <i>userid</i> (for default domain users) or <i>userid@domain</i> (for other users). This is a required option.
-v <i>mode</i>	Specifies verbose mode. The values for <i>mode</i> are 0 , 1 , and 2 . 0 specifies no output. 1 specifies mailbox level output. 2 (default) specifies message level output.

Example

In the following example, **imsexport** extracts all email for user **smith1**. **smith1** is a valid user account in the Sun Java System Messaging Server message store. User **smith1** has three folders on the store: **INBOX** (the normal default user folder), **private**, and **private/mom**. The destination directory will be **/tmp/joes_mail**.

```
% imsexport -u smith1 -d /tmp/joes_mail/
```

imsexport then transfers each message store folder into a **/var/mail** conforming file. Thus you will get the following files:

- **/tmp/joes_mail/INBOX**
- **/tmp/joes_mail/private**
- **/tmp/joes_mail/private.msg**
- **/tmp/joes_mail/private/mom**

imsimport

The **imsimport** utility migrates UNIX **/var/mail** format folders into a Messaging Server message store.

The **imsimport** utility extracts the messages stored in **/var/mail** mailboxes and appends them to the corresponding users' mailbox in the Messaging Server message store. Files in the directory that are not in the standard UNIX mailbox format are skipped. If the corresponding users do not exist in the message store, **imsimport** creates them. When the user quota is exceeded, **imsimport** bypasses the message store quota enforcement, so the user does not receive an "over quota" message.

The **imsimport** utility can be run while Messaging Server is running. If mail delivery is enabled for the mailbox you are importing, old mail can get mixed with new mail, so you might want to hold the delivery of this user during the migration. Mailbox access should not be a problem.

When **imsbackup**, **imsrestore**, **imsimport** or any processing intensive operation takes significantly more system resources than normal, and continues doing so longer than the **msprobe** interval, there may be a temporary backlog of DB transaction log files to be cleared. If there are more files than specified in **local.store.maxlog**, then **msprobe** may erroneously restart all the processes during a restore. To prevent this from happening, disable **msprobe** during the **imsbackup**, **imsrestore**, and **imsimport**.

Note: **imsimport** does not use the IMAP server. However, the **stored** utility must be running to maintain message store integrity. The LDAP server should be running if **imsimport** is expected to create new users.

Location: *MessagingServer_home/bin/*

Syntax

```
imsimport -u <user> -s _file_ [-c y|n] [-d <mailbox>] [-e <encoding>]
[-b] [-g] [-i] [-n] [-v <mode>]
```

Options

The options for this command are:

Table 75–24 *imsimport* Command Options

Option	Description
-b	Subscribes the destination mailbox folder when a new folder is created by the imsimport utility.
-c y n	Provides an answer to the question: "Do you want to continue?" if an error occurs. Specify y for yes, n for no.
-d mailbox	Specifies the destination mailbox where the messages will be stored.
-e encoding	Specify an encoding option.
-g	Specifies debugging mode.
-i	Ignores the content-length field
-n	Creates a new mailbox with a <i>.date</i> extension if the mailbox exists. The <i>.date</i> extension is in the following form: <i>.mmddy.HHMMSS</i> The month is specified by <i>mm</i> . The day is specified by <i>dd</i> . The year is specified by <i>yy</i> . For example, 052097 specifies May 20 in the year 1997. The time of day is specified by <i>HHMMSS</i> . For example 110000 specifies 11:00am.
-s file	Specifies the UNIX folder's file name from which the messages are to be imported. The <i>file</i> option must be a full path name. This is a required option.
-u user	Specifies the message store identification for a user. Note that this should be identical in content and case to the UID attribute of the user's LDAP entry. If it is not the same, users will not be able to login to their INBOXES.
-v mode	Specifies verbose mode. The values for <i>mode</i> are 0, 1, and 2. 0 specifies no output. 1 specifies mailbox level output. 2 (default) specifies message level output.

Examples

imsimport migrates the specified **/var/mail/folder** for the specified user to the Messaging Server message store. If the destination folder is not specified, **imsimport** calls the destination folder by the same name as the source folder. In the following example, the command migrates the default **/var/mail/INBOX** for the user **smith**, to the **INBOX**.

```
imsimport -u smith -s /var/mail/smith -d INBOX
```

Similarly, if you are trying to move a folder called **test** from **/home/smith/folders/** to the Messaging Server message store, use the following command:

```
imsimport -u smith -s /home/smith/folders/test -d test
```

If a destination folder called **test** already exists in the Messaging Server message store, **imsimport** appends the messages to the existing folder in the mailbox.

imsrestore

The **imsrestore** utility restores messages from the backup device into the message store. See "Backing Up and Restoring the Message Store" for more information.

When **imsbackup**, **imsrestore**, **imsimport** or any processing intensive operation takes significantly more system resources than normal, and continues doing so longer than the **msprobe** interval, there may be a temporary backlog of DB transaction log files to be cleared. If there are more files than specified in **local.store.maxlog**, then **msprobe** may erroneously restart all the processes during a restore. To prevent this from happening, disable **msprobe** during the **imsbackup**, **imsrestore**, and **imsimport**.

Location: *MessagingServer_home/bin*

Syntax

```
imsrestore -f <device>|- [-b <blocking_factor>] [-c y | n]
[-d <path>] [-D] [-e <encoding>] [-h] [-i|-E] [-m <file>]
[-n] [-p] [-P partition] [-r <file>] [-s] [-S] [-t] [-u <file>] [-g] [-v
0|1|2|3|4|5]
[-x] [<name> . . .]
```

Options

The options for this command are:

Table 75–25 *imsrestore Options*

Option	Description
-b <i>blocking_factor</i>	Indicates the blocking factor. Everything read on the device is performed by blocks of the size 512 x <i>blocking_factor</i> . The default is 20. Note: this number needs to be the same blocking factor that was used for the backup.
-c y n	Provides an answer to the question: "Do you want to continue?" if an error occurs. Specify y for yes, n for no.
-d <i>path</i>	Restore all mailboxes under <path> (for multiple incremental restore).
-e <i>encoding</i>	Mailbox name encoding (example: IMAP-mailbox-name).
-E	Updates and expunges existing messages.
-f [<i>device filename -</i>]	Specify the device or file from which to read the backup data. Use -f- to read from stdin (ie, to pipe the data from some other command).
-h	Dumps the header.
-g	Debug mode.
-i	Ignores existing messages. Does not check for existing messages before restore. Note that if you specify the -i option, you may have duplicate messages after the restore, since the -i option supersedes your ability to check for duplicates.

Table 75–25 (Cont.) *imsrestore* Options

Option	Description
-m <i>file</i>	This mapping file is used when renaming user IDs. The format in the mapping file is <i>oldname=newname</i> with one set of names per line. For example: a=x b=y c=z where a , b , and c are old names and x , y , and z are new names. For users in a hosted domain, add @domain_name to the old and new user IDs. This option is used to rename user IDs when restoring from a Messaging Server 5.x, 6.x, or 7.x imsbackup file. Use the -u option when restoring from a SIMS 4.0 imbackup file.
-n	Creates a new mailbox with a .date extension (if the mailbox exists). By default, messages are appended to the existing mailbox.
-p	Restore mailboxes to original partition.
-P	Restore all mailboxes to this partition (ignore -p and mailMessageStore).
-r <i>file</i>	Reference file name (will restore all links in <i>file</i>).
-s	Used to restore a large file without using large file seeking.
-t	Prints a table of contents, but restore is not performed.
-u <i>file</i>	This file contains the object names (entire message store, user, group, mailbox, and so on) to restore. See <i>name</i> for a list of backup objects. For restoring SIMS 4.0 data into a Message Server 5.x, 6.x, or 7.x system, you can specify or rename users with -u file . Users should have one name per line. If you rename users, the format of <i>file</i> is oldname=newname with one set of names per line. For example: joe bonnie jackie=jackie1 where joe and bonnie are restored, and jackie is restored and renamed to jackie1 . Note that full object pathnames are not needed for user IDs.
-v [0 1 2 3 4 5]	Executes the command in verbose mode. 0 = no output 1 = output at mailbox level 2 = output at message level (default) 3 = print meta data (for use with -t only) 4 =print object level meta data (for use with -t only) 5 =print the backup data of mailboxes and messages (for use with -t only)
-x	Restores stubs into message files.
-S	Not Applicable.
<i>name</i>	Can be 1) logical pathname of the backup object, 2) user ID, 3) <i>mailbox</i> . See " imsbackup " for description.

Examples

The following example restores the messages from the file **backupfile**:

```
imsrestore -f backupfile
```

The following example restores the messages for **joe** from the file **backupfile**:

```
imsrestore -f backupfile /primary/user/joe
```

The following example lists the content of the file **backupfile**:

```
imsrestore -f backupfile -t
```

The following example renames users in the file **mapfile**:

```
imsrestore -m mapfile -f backupfile
```

where the **mapfile** format is **oldname=newname**:

```
userA=user1  
userB=user2  
userC=user3
```

mboxutil

The **mboxutil** command lists, creates, deletes, renames, or moves mailboxes (folders). You can use the **mboxutil** command to report quota information, restore expunged messages, and list mailbox subscriptions.

Note: Do not "kill" the **mboxutil** process (**SIGKILL (kill -9)** command) in the middle of execution. Otherwise, you will need to restart the Messaging Server process.

You can use POSIX regular expressions in the **mboxutil** command.

Requirements: Must be run locally on the Messaging Server. The **stored** utility must also be running.

Location: *MessagingServer_home/bin*

Mailbox Naming Conventions

You must specify mailbox names in the following format:

```
user/userid/mailbox
```

where *userid* is the user that owns the mailbox and *mailbox* is the name of the mailbox.

For hosted domains, *userid* is **uid@domain**.

For example, the following command creates the mailbox named **INBOX** for the user whose user ID is **crowe**. **INBOX** is the default mailbox for mail delivered to the user **crowe**.

```
mboxutil -c user/crowe/INBOX
```

Important: The name **INBOX** is reserved for each user's default mailbox. **INBOX** is the only folder name that is case-insensitive. All other folder names are case-sensitive.

Syntax

List mailboxes:

```
mboxutil -l [-B <bound>] [-E <encoding>] [-p <MUTF7 IMAP pattern> | -P <regular expression>] [-x | -s] [-D] [-z]
```

Create mailboxes:

```
mboxutil -c [-E <encoding>] {<mailbox> | -f <file>}
```

Delete mailboxes:

```
mboxutil -d [-E <encoding>] { -p <MUTF7 IMAP pattern> | -P <regular expression> | -f <file> | <mailbox> }
```

Rename mailboxes:

```
mboxutil -r [-E <encoding>] { <oldname> <newname> | -f <file> } [<partition>]
```

Expunge mailboxes:

```
mboxutil -e [-E <encoding>] [-p <MUTF7 IMAP pattern> | -P <regular expression> ]
```

Schedule cleanup:

```
mboxutil -C [-g num] [-E encoding] [-p <MUTF7 IMAP pattern> | -P <regular expression>]
```

List orphan/inactive mailboxes:

```
mboxutil -o [-w <file> ] [-t <number of days>]
```

Restore expunged messages that have not been purged:

```
mboxutil -R [-E <encoding>] <mailbox name>
```

List personal mailbox subscriptions:

```
mboxutil -S [-n [-f <file>] | -u -f <file>]
```

Options

The options for this command are:

Table 75–26 mboxutil Command Options

Option	Description
-B bound	Uses the user specified delimiter while listing the mailboxes. The option is always used with the -l option like <code>mboxutil -l [-B bound]</code> . The given bound can be any string.
-c mailbox	Creates the specified mailbox. A mailbox must exist before creating a secondary mailbox.

Table 75–26 (Cont.) mboxutil Command Options

Option	Description
-C	<p>Schedules cleanup immediately.</p> <p>Cleanup is automatically scheduled when mailboxes are expunged. This command can be used to schedule cleanup manually when:</p> <ul style="list-style-type: none"> ▪ The storage is very full. ▪ store.cleanupsize has been reduced and you don't want to wait for users to do the next expunge. ▪ Or you want to force cleanup on specific users/folders. For more information see "Mechanism to Schedule Cleanup Immediately".
-d mailbox	<p>Deletes the specified mailbox. To delete a user from the message store, use the following value for -d mailbox:</p> <p>user/userid/INBOX</p> <p>For example, to delete the user john from the message store, use -d user/john/INBOX. To delete the mm folder in the user john's mailbox, use -d user/john/mm.</p> <p>The recommended method to delete a user is to mark the user status as deleted in the LDAP directory (by using the Delegated Administrator utility commadmin user delete command or the Delegated Administrator console.). Next, remove resources from the user that have been marked as deleted for a period longer than a specified number of days by using the msuserpurge command for all mail services and csclean for all calendar services. Finally, use the commadmin domain purge command to permanently remove the users.</p> <p>If you use the Delegated Administrator utility as described in the preceding paragraph, you do not have to use the mboxutil -d command to delete a mailbox.</p> <p>Note that -p and -P can be used in conjunction with one another.</p>
-D	<p>Lists deleted mailboxes. When used with the -I option, lists deleted folders (as long as if the you enabled the store.mailboxpurgedelay option before the delete).</p>
-e	<p>Expunges all deleted messages in the message store. This option also can be used with the -p pattern or -P regexp options to expunge all deleted mailboxes with names that match <i>pattern</i> or <i>regexp</i>.</p>
-E encoding	<p>Specifies character set encoding used for the folder name display or the input of localized mailbox names. The default is that the folder name is displayed (or input accepted) in the localized character set (that is, in the form the user would expect to see). Using "-E MUTF-7" causes the folder name to be displayed (or input accepted) in the form used internally by the message store. For more information, see "Localized Mailbox Names in imexpire".</p>
-f file	<p>Specifies a file that stores mailbox names. The -f option can be used with the -c, -r, -S, or -d options.</p> <p>The file contains a list of mailboxes on which the mboxutil command is executed. The following is an example of entries in a data file:</p> <p>user/daphne/INBOXuser/daphne/projxuser/daphne/mm</p>

Table 75–26 (Cont.) mboxutil Command Options

Option	Description
-g <i>num</i>	Minimum cleanup size, <i>num</i> must be > 0. Default is <code>store.cleanupsize</code>
-l	Lists all of the mailboxes on a server. mboxutil -l correctly displays characters associated with the system locale under which mboxutil is being executed. The -P <i>regex</i> option accepts international characters.
-n	Lists personal, non-existing mailbox subscriptions. The -n option is used with the -S option.
-o	Checks for orphaned accounts. This option searches for inboxes in the current messaging server host that do not have corresponding entries in LDAP. For example, the -o option finds inboxes of owners who have been deleted from LDAP or moved to a different server host. For each orphaned account it finds, mboxutil writes the following command to the standard output: mboxutil -d user/userid/ INBOX unless -w is specified
-p <i>pattern</i>	When used with the -l option, lists only those mailboxes with names that match <i>pattern</i> . Can also be used with the -d or -e option to delete or expunge mailboxes with names that match <i>pattern</i> . You can use IMAP wildcards. This option expects a pattern in IMAP M-UTF-7 format. This is not the recommended way to search for non-ASCII mailboxes. To search for non-ASCII mailboxes, use the -P option.
-P <i>regex</i>	Lists, deletes, or expunges only those mailboxes with names that match the specified POSIX regular expression. This option expects the <i>regex</i> in the local character encoding.
-q <i>domain</i>	Obsolete. Use imquotacheck -d domain
-r <i>oldnamenewname</i> [<i>partition</i>]	Renames the mailbox from <i>oldname</i> to <i>newname</i> . Optionally, to move a folder from one partition to another, specify the new partition with the partition option. This option can be used to rename a user. For example, mboxutil -r user/user1/INBOX user/user2/INBOX moves all mail and mailboxes from user1 to user2 , and new messages will appear in the new INBOX. (If user2 already exists, this operation fails.) Folders are cleaned and purged automatically during the rename process. When used with imsbackup and imsrestore , you can use mboxutil -r to transfer ownership of a set of shared folders to another user while preserving the readership settings.
-R <i>mailbox</i>	Restores deleted messages that have not yet been purged. When a mailbox is expunged or expired, the uids of the deleted messages are stored in a store.exp file. The messages are physically removed by impurge depending on the values of store.cleanupsize and store.cleanupage . If expunge or expire is done by mistake, this option can be used to restore the deleted messages that have not yet been purged by impurge into the original mailbox.
-s	When used with the -l option, displays only the mailbox name. No other data is displayed.

Table 75–26 (Cont.) mboxutil Command Options

Option	Description
-S	Lists personal mailbox subscriptions. When used with the -n option, the -S option lists or unsubscribes non-existing mailbox subscriptions.
-t days	When used with the -o option, lists the mailboxes that have not been accessed in a specified number of days (<i>days</i>). Thus, the -t option identifies inactive mailboxes (based on last-accessed date) together with orphaned mailboxes (mailboxes that do not have corresponding user entries in the LDAP directory). To identify (list) the orphaned and inactive mailboxes, use mboxutil -o -w file-t num . To delete these orphaned and inactive mailboxes, use mboxutil -d -f file , where <i>file</i> is the same file as the one passed to -w in the preceding command. To use this feature, the configuration variable local.enablelastaccess must be enabled for at least the number of days specified with the -t option.
-u	Unsubscribes personal non-existing mailbox subscriptions listed in the file specified with the -f file option.
-w file	Used with the -o option. Writes to a file the mailbox names generated by the -o option (which identifies orphaned accounts).
-x	When used with the -l option, displays the path and access control for a mailbox.
-z	When used with the -l option, displays a count of the mailboxes displayed.

Note: When using the separator, check output consistency by trying the separator string as a single or double quoted string with escape sequences if required.

Examples

To list all mailboxes for all users:

```
mboxutil -l
```

To list all mailboxes and also include path and ACL information:

```
mboxutil -l -x
```

To list all mailboxes displaying only the mailbox names:

```
mboxutil -l -s
```

To create the default mailbox named **INBOX** for the user **daphne**:

```
mboxutil -c user/daphne/INBOX
```

To delete a mail folder named **projx** for the user **delilah**:

```
mboxutil -d user/delilah/projx
```

To delete the default mailbox named **INBOX** and all mail folders for the user **druscilla**:

```
mboxutil -d user/druscilla/INBOX
```

To rename Desdemona's mail folder from **memos** to **memos-april**:

```
mboxutil -r user/desdemona/memos user/desdemona/memos-april
```

To schedule a cleanup of Dorothea's mailbox if there are at least 50 expunged messages:

```
$ mboxutil -C -g 50 -p user/dorothea/*
```

To restore messages that were expunged from Desdemona's **memos** mail folder in the last 24 hours:

```
mboxutil -R -t 24 user/desdemona/memos
```

To move the mail account for the user **dimitria** to a new partition:

```
mboxutil -r user/dimitria/INBOX user/dimitria/INBOX partition
```

where **partition** specifies the name of the new partition.

To move the mail folder named **personal** for the user **dimitria** to a new partition:

```
mboxutil -r user/dimitria/personal user/dimitria/personal partition
```

To list orphaned mailboxes and mailboxes that have not been accessed in 60 days:

```
mboxutil -o -w orphanfile -t 60
```

The preceding example writes the list of orphaned and inactive mailboxes to a file named **orphanfile**.

To delete orphaned and inactive mailboxes:

```
mboxutil -d -f orphanfile
```

where **orphanfile** is a file that has stored a list of orphaned and inactive mailboxes identified with the **-o** option.

To list personal, non-existing mailbox subscriptions for user mailboxes listed in a file named **orphanfile**:

```
mboxutil -S -n -f orphanfile
```

To unsubscribe the non-existing mailbox subscription list generated by the previous example:

```
mboxutil -S -u -f orphanfile
```

The **mboxutil** command has stable output formats that we support and allow customers to parse. **mboxutil -l -p user/%/INBOX -s** lists all accounts (INBOXES) in the store. **mboxutil -o -w -t** lists inactive accounts. The formats are the same: one line per mailbox name. To generate an active list, remove the inactive mailboxes from the all mailbox list. Sample **mboxutil** output:

```
user/joe/INBOX
user/mary/INBOX
```

Mechanism to Schedule Cleanup Immediately

Normally, cleanup is scheduled automatically when mailboxes are expunged. This command can be used to schedule cleanup manually when:

- The storage is very full.
- `store.cleanupsize` is reduced.

For example, the following command schedules cleanup of mailboxes with at least 50 expunged messages. Messages are removed when `store.cleanupage` has expired.


```
$ mboxutil -C -g 50
```

mkbackupdir

The `mkbackupdir` utility creates and synchronizes the backup directory with the information in the message store. It is used in conjunction with Solstice Backup (Legato Networker). The backup directory is an image of the message store. It does not contain the actual data. `mkbackupdir` scans the message store's user directory, compares it with the backup directory, and updates the backup directory with the new user names and mailbox names under the message store's user directory.

The backup directory is created to contain the information necessary for Networker to backup the message store at different levels (server, group, user, and mailbox). The following figure displays the structure.

Figure 75–1 Backup Directory Hierarchy



Location: `MessagingServer_home/bin`

The variables in the backup directory contents are:

Table 75–27 mkbackupdir Variables

Variable	Description
<code>BACKUP_ROOT</code>	Message store administrator root directory.
<code>partition</code>	Store partition.

Table 75–27 (Cont.) mkbackupdir Variables

Variable	Description
<i>group</i>	System administrator-defined directories containing user directories. Breaking your message store into groups of user directories allows you to do concurrent backups of groups of user mailboxes. To create groups automatically, specify your groups in the <i>MessagingServer_</i> <i>home/config/backup-groups.conf</i> file. The format for specifying groups is: <i>groupname= patterngroupname</i> is the name of the directory under which the user and mailbox directories will be stored, and <i>pattern</i> is a folder name with IMAP wildcard characters specifying user directory names that will go under the <i>groupname</i> directory.
<i>user</i>	Name of the message store user.
<i>folder</i>	Name of the user mailbox.
<i>mailbox</i>	Name of the user mailbox.

The **mkbackupdir** utility creates:

- A default *group* directory (**ALL**) or the group directories defined in the **backup-groups.conf** configuration file. The following is a sample **backup-groups.conf** file:

```
groupA=a* (regexp)
groupB=b*
groupC=c*
.
.
.
```

- A *user* directory under the backup directory for each new user in the message store.
- A 0 length mailbox file for each mailbox.
- A **.nsr** file for each subdirectory that contains user mailboxes.

The **.nsr** file is the NSR configuration file that informs the Networker to invoke **imsasm**. **imsasm** then creates and interprets the data stream.

Each user mailbox contains files of zero length. This includes the **INBOX**, which is located under the *user* directory.

Note: Make sure the backup directory is writable by the message store owner (mailsrv).

Syntax

```
mkbackupdir [-a <name_of_asm>] [-i | -f | -u] [-g] [-t <number_of_threads>] [-e
<encoding>] [-v] -p <directory>
```

Options

The options for this command are:

Table 75–28 *mkbackupdir Options*

Option	Description
-a <i>name_of_asm</i>	Creates .nsr files using the specified asm name. This can be used for when you have multiple instances of Messaging Server as in symmetric HA environments.
-e <i>encoding</i>	Specify an encoding option.
-f	Backs up the folders only. By default, all mailboxes are backed up.
-g	Executes the command in debug mode.
-i	Backs up the inbox only. By default, all mailboxes are backed up.
-p <i>directory</i>	Specifies the directory for the backup image. This is a required option when local.store.backupdir is not configured. The max_thread option must be set between 1 and 1024. Note: The Networker has a limitation of 64 characters for <i>saveset</i> name. If your default backup directory pathname is too long, you should use this option to specify another pathname.
-t <i>number_of_threads</i>	Specifies the number of threads that can be used to create the backup directory. The default is one thread for each partition, which is usually adequate. If you have many partitions, and you do not want mkbackupdir to consume all your resources, you can lower this number.
-u	User level backup. Instead of backing up each folder as a file, create a backup file per user.
-v	Executes the command in verbose mode.

Examples

To create the **mybackupdir** directory, enter the following:

```
mkbackupdir -p /mybackupdir
```

msprobe

msprobe is a daemon that probes servers to see if they respond to service requests. When **msprobe** detects a possible problem, it can, depending upon other configuration, potentially let the Watcher know (at which point the Watcher can attempt to restart a troubled component) and/or generate an alarm message; see the Watcher options and Alarm options, respectively in the *Messaging Server Reference*.

For additional information on **msprobe**, see the *Messaging Server Reference*.

Location: *MessagingServer_home/lib*

Syntax

```
msprobe [-d] [-n] [-r] [server]
```

Options

The options for this command are:

Table 75–29 msprobe Command Options

Option	Description
-d	Debug mode.
-n	Disable restart.
-r	Report the server, port, <i>response_time</i> , status, message to stdout . Status is: OK , SLOW , or BAD . Response time is in millisecond.
<i>server</i>	The accepted values for <i>server</i> are: imap pop http ens cert job_controller smtp lmt submit metermaid deploymap If server is not specified, all enabled servers are probed.

Example

```
$ msprobe -r
imap,143,0,OK,"No error"
pop,110,-1,BAD,"Connection refused"
ens,7997,0,OK,"No error"
http,8990,0,OK,"No error"
job_controller,27442,0,OK,"No error"
smtp,25,35,SLOW,"No error"
submit,587,1,OK,"No error"

$ msprobe -r imap
imap,143,0,OK,"No error"
```

msuserpurge

When user and domain mailboxes are marked for deletion, the **msuserpurge** command purges those user and domain mailboxes from the message store. Specifically, this command scans the following domain and user status attributes in LDAP for a value of deleted: **inetDomainStatus**, **mailDomainStatus**, **inetUserStatus**, **mailUserStatus**. This command can be run at the command line, or can be scheduled for execution with the **configutil** option **local.sched.userpurge**.

Requirements: If run manually, it must be manually run locally on the messaging server. Make sure that the environment variable **ConfigRoot** is set to *MessagingServer_home/config*.

Location: *MessagingServer_home/lib*

Syntax

```
msuserpurge [-v|-r] [-d <domain>] [-g <days>]
```

Options

The options for this command are:

Table 75–30 *msuserpurge Options*

Option	Description
-v	Verbose (list deleted users).
-r	Report (list all users, do not purge).
-g <i>days</i>	Specify a grace period.
-d <i>domain</i>	Specify a domain.

Examples

```
msuserpurge -d example.com -g 0
```

readership

The owner of an IMAP folder can grant permission for other users to access the folder. A folder that other users are allowed to access is called a shared folder. See "[Managing Shared Folders](#)" for more information. Users can modify the access rights on folders if their mail provides an interface to the **SETACL IMAP** commands. Administrators can use the **readership** utility to set or remove access rights on the folder, and to see how many users other than the owner are accessing a shared folder.

To list the rights on all shared folders, the mail administrator can use the **imcheck -d lright.db** command. To list rights on individual folders, use the **mboxutil -lxp folder** command.

Requirements: Must be run locally on the Messaging Server. The **stored** process must also be running.

Location: *MessagingServer_home/bin/*

Syntax

```
readership [ -d <days> | -p <months> ]
readership -r <folder pattern>
readership -s <foldername> <identifier> <acl_right>
readership -s -m <folder-pattern> <identifier> <acl_right>
readership [-e <encoding>] -s [ -m pattern | mailbox ] <identifier> <acl_right>
```

To remove all rights for a user who has previously been granted access, set that user's rights to a null list, for example:

```
readership -s user/<userid>/<folder> <user-to-remove> ''
```

Options

The options for this command are:

Table 75–31 *readership Command Options*

Option	Description
-d <i>days</i>	Counts as a reader any identity that has selected the shared IMAP folder within the indicated number of days. The default is 30.

Table 75–31 (Cont.) readership Command Options

Option	Description
-p <i>months</i>	Does not count users who have not selected the shared IMAP folder within the indicated number of months. The default is infinity and removes the seen flag data for those users. This option also removes the seen flag data for those users from the store.
-r <i>folder_pattern</i>	Lists the last access time of the folders matching the pattern.
-s <i>folderidentifieracl_right</i>	Sets ACL rights character for folder, where <i>folder</i> is the name of the folder for which you are setting rights, <i>identifier</i> is the person or group to whom you are assigning the rights, and <i>acl_rights</i> are the rights you are assigning. See Table 75–32, "ACL Rights Characters". To remove a rights setting, specify a null set of rights.
-s -m <i>folder_patternidentifieracl_right</i>	Sets ACL rights as -s but on all folders matching the specified pattern.
-e <i>encoding</i>	Specifies character set encoding used for the folder name display or the input of localized mailbox names. The default is that the folder name is displayed (or input accepted) in the localized character set (that is, in the form the user would expect to see). Using " -e MUTF-7 " causes the folder name to be displayed (or input accepted) in the form used internally by the message store.

Table 75–32 ACL Rights Characters

Character	Description
l	lookup - User can see and subscribe to the shared folder. (IMAP commands allowed: LIST and LSUB)
r	read - Users can read the shared folder. (IMAP commands allowed: SELECT , CHECK , FETCH , PARTIAL , SEARCH , COPY from the folder)
s	seen - Directs the system to keep seen information across sessions. (Set IMAP STORESEEN flag)
w	write - Users can mark as read, and delete messages. (Set IMAP STORE flags, other than SEEN and DELETED)
i	insert - Users can copy and move email from one folder to another. (IMAP commands allowed: APPEND , COPY into folder)
p	post - Users can send mail to the shared folder email address. (No IMAP command needed)
k	create - Users can create new sub-folders. (IMAP command allowed: CREATE)
x	delete - Users can delete entries from the shared folder. (IMAP commands allowed: EXPUNGE , set STOREDELETED flag)
a	administer - Users have administrative privileges. (IMAP command allowed: SETACL)
t	delete - For messages, sets or clears \DELETED flag via STORE , or sets \DELETED flag during APPEND/COPY .
e	expunge - Performs EXPUNGE and expunge as a part of CLOSE .

Table 75–32 (Cont.) ACL Rights Characters

Character	Description
n	access - Retrieves annotation information about the folder (see RFC 5257 at http://tools.ietf.org/html/rfc5257#section-4.10).

Example

```
bin/readership -s -m user/kellyc/a\* barbs lr
user/kellyc/abc: OK
user/kellyc/abc/123: OK
user/kellyc/abc/456: OK
```

reconstruct

The **reconstruct** utility rebuilds one or more mailboxes, or the master mailbox file (the mailboxes database), and repairs any inconsistencies. Use this utility to recover from almost any form of data corruption in the message store.

A mailbox consists of files under the user partition directory. The mailboxes database is the mboxlist database.

Requirements: Must be run locally on the Messaging Server host; the **stored** utility must also be running.

Location: *MessagingServer_home/bin*

Syntax

```
reconstruct [-n | -f] [-i] [-x] [-p partition] {-r [<mailbox...>] | <mailbox...>}
```

```
reconstruct [-p <partition> [ -u <user> ] ] -m
```

```
reconstruct -a [-n | -f] [-t <nthreads>] [<mailbox...>]
```

```
reconstruct -A [-u userid]
```

```
reconstruct -l
```

```
reconstruct -q
```

```
reconstruct -s
```

Options

The options for this command are:

Table 75–33 reconstruct Options

Option	Description
-a <i>mailbox...</i>	Addresses a user's mailbox comprehensively. If reconstruct -a -f is run, a full repair and reparse is forced. If reconstruct -a -n is run, only a consistency check is run, printing out errors found in the folder. This can be useful to catch issues before fixing them. Since there are many recovery functions embedded in the store libraries, running -n might still fix and modify the folder to some degree.
-A [-u <i>userid</i>]	Repairs the annotation database.
-i	Sets the store.idx file length to zero before reconstructing.
-f	Forces reconstruct to perform a fix on the mailbox or mailboxes.

Table 75-33 (Cont.) reconstruct Options

Option	Description
-x	Recovers partially delivered messages.
-l	Reconstructs lright.db .
-m	Performs a consistency check and, if needed, repairs the mailboxes database. This option examines every mailbox it finds in the spool area, adding or removing entries from the mailbox's database as appropriate. The utility prints a message to the standard output file whenever it adds or removes an entry from the database. Specifically it fixes folder.db , quota.db , and lright.db
-n	Checks the message store only, without performing a fix on the mailbox or mailboxes. The -n option cannot be used by itself unless a mailbox name is provided. When a mailbox name is not provided, the -n option must be used with the -r option. The -r option may be combined with the -p option. For example, any of the following commands are valid: reconstruct -n user/dulcinea/INBOX reconstruct -n -rreconstruct -n -r -p primary reconstruct -n -r user/dulcinea
-e	Obsolete, see mboxutil -R .
-o	Obsolete, see mboxutil -o .
-o-d filename	Obsolete, see mboxutil -o .
-p partition	The -p option is used with the -m option and limits the scope of the reconstruction to the specified partition. If the -p option is not specified, reconstruct defaults to all partitions. Specifically it fixes folder.db and quota.db , but not lright.db . This is because fixing the lright.db requires scanning the acls for every user in the message store. Performing this for every partition is not very efficient. To fix lright.db run reconstruct -l . Specify a partition name. Do not use a full path name.
-q	Fixes any inconsistencies in the quota subsystem, such as mailboxes with the wrong quota root or quota roots with the wrong quota usage reported. The -q option can be run while other server processes are running.
-r [mailbox]	Repairs and performs a consistency check of the partition area of the specified mailbox or mailboxes. The -r option also repairs all sub-mailboxes within the specified mailbox. If you specify -r with no mailbox argument, the utility repairs the spool areas of all mailboxes within the user partition directory.
-s	Repairs subscriptions.
-t nthreads	Performs a multithreaded reconstruct -a with <i>nthreads</i> threads (if -t is not specified, <i>nthreads</i> is chosen automatically, single-threaded if specific mailbox is specified).
-u user	The -u option is used with the -m option and limits the scope of the reconstruction to the specified user. The -u option must be used with the -p option. If the -u option is not specified, reconstruct defaults to all partitions or to the partition specified with the -p option. Specify a user name. Do not use a full path name.

The *mailbox* argument indicates the mailbox to be repaired. You can specify one or more mailboxes. Mailboxes are specified with names in the format **user/userid/sub-mailbox**, where *userid* is the user that owns the mailbox. For example, the inbox of the user **dulcinea** is entered as: **user/dulcinea/INBOX**.

Examples

The following command performs a reconstruct on a specific mailbox:

```
reconstruct user/dulcinea/INBOX
```

The following checks the specified mailbox, without performing a reconstruct:

```
reconstruct -n user/dulcinea/INBOX
```

The following command checks all mailboxes in the message store:

```
reconstruct -n -r
```

To Rebuild Mailboxes

To rebuild mailboxes, use the **-r** option. You should use this option when:

- Accessing a mailbox returns one of the following errors: "System I/O error" or "Mailbox has an invalid format".
- Accessing a mailbox causes the server to crash.
- Files have been added to or removed from the spool directory.

reconstruct -r first runs a consistency check. It reports any inconsistencies and rebuilds only if it detects any problems. Consequently, performance of the **reconstruct** utility is improved with this release. You can use **reconstruct** as described in the following examples: To rebuild the spool area for the mailboxes belonging to the user **daphne**, use the following command:

```
reconstruct -r user/daphne
```

To rebuild the spool area for all mailboxes listed in the mailbox database:

```
reconstruct -r
```

You must use this option with caution, however, because rebuilding the spool area for all mailboxes listed in the mailbox database can take a very long time for large message stores. (See the discussion about **reconstruct** Performance below.) A better method for failure recovery might be to use multiple disks for the store. If one disk goes down, the entire store does not. If a disk becomes corrupt, you need only rebuild a portion of the store by using the **-p** option as follows:

```
reconstruct -r -p subpartition
```

To rebuild mailboxes listed in the command-line argument only if they are in the **primary** partition:

```
reconstruct -p primary mbox1 mbox2 mbox3
```

If you do need to rebuild all mailboxes in the **primary** partition:

```
reconstruct -r -p primary
```

If you want to force reconstruct to rebuild a folder without performing a consistency check, use the **-f** option. For example, the following command forces a reconstruct of the user folder **daphne**:

```
reconstruct -f -r user/daphne
```

To check all mailboxes without fixing them, use the **-n** option as follows:

```
reconstruct -r -n
```

reconstruct -r -f focuses on fixing the folder index files. It assumes the folder record is good, and the peruser record is good. There are other commands to address these other data areas, such as reconstruct **-m**.

reconstruct -a attempts to address these all at once.

So if you think you need to repair an index file, but you're not sure if it is in the **folder.db**, you should not need to worry about running a **reconstruct -m** first, and whether the index corruption will be handled correctly there, or if both these will result in something that will conflict with the peruser entry later.

If you know your problem is with the index file, and there are no other complications, then you can go ahead and use **reconstruct -r -f** to save time.

Checking and Repairing Mailboxes

To perform a high-level consistency check and repair of the mailboxes database:

```
reconstruct -m
```

To perform a consistency check and repair of the primary partition:

```
reconstruct -p primary -m
```

Note: Running **reconstruct** with the **-P** and **-m** flags together will not fix **lright.db**. This is because fixing the **lright.db** requires scanning the ACLs for every user in the message store. Performing this for every partition is not very efficient. To fix the **lright.db** run **reconstruct -l**

To perform a consistency check and repair of an individual user's mailbox named **john**:

```
reconstruct -p primary -u john -m
```

You should use the **-m** option when:

- One or more directories were removed from the store spool area, so the mailbox database entries also need to be removed.
- One or more directories were restored to the store spool area, so the mailbox database entries also need to be added.
- The **stored -d** option is unable to make the database consistent. If the **stored -d** option is unable to make the database consistent, you should perform the following steps in the order indicated:
 - Shut down all servers.
 - Remove all files in *store_root/mboxlist*.
 - Restart the server processes.
 - Run **reconstruct -m** to build a new mailboxes database from the contents of the spool area.

reconstruct Performance

The time it takes **reconstruct** to perform an operation depends on the following factors:

- The kind of operation being performed and the options chosen
- Disk performance
- The number of folders when running **reconstruct -m**
- The number of messages when running **reconstruct -r**

- The overall size of the message store
- What other processes the system is running and how busy the system is
- Whether or not there is ongoing POP, IMAP, HTTP, or SMTP activity. Note that **reconstruct** is designed to run with the store services up. It is not necessary to keep the store offline to run **reconstruct**.

The **reconstruct -r** option performs an initial consistency check; this check improves **reconstruct** performance depending on how many folders must be rebuilt. The following performance was found with a system with approximately 2400 users, a message store of 85GB, and concurrent POP, IMAP, or SMTP activity on the server:

- **reconstruct -m** took about 1 hour
- **reconstruct -r -f** took about 18 hours

Note: A **reconstruct** operation may take significantly less time if the server is not performing ongoing POP, IMAP, HTTP, or SMTP activity.

refresh

The **refresh** utility refreshes the configuration of the specified messaging server processes (SMTP, IMAP, POP, STORE, HTTP, ENS, SCHED). It is used when an option for one of the services has been modified and you wish this option to take effect.

Location: *MessagingServer_home/bin*

Syntax

```
refresh [ -h | <service>]
```

Options

The options for this command are:

Table 75–34 *refresh* Command Options

Option	Description
-h	Show usage statement and list available services that can be refreshed.
<i>service</i>	Refresh just the specified service. Accepted values are: <ul style="list-style-type: none"> watcher metermaid store imap pop cert http sched dispatcher job_controller mmp mta <p>If no components are specified, all enabled services will be refreshed.</p>

Examples

The following command refreshes the scheduler utility:

```
refresh sched
```

If refresh does not cause the change to take effect, then stop and restart the service.

rehostuser

The **rehostuser** utility enables you to move a Messaging Server user's mail store from one mailhost to another. It also disconnects any active session, locks the store to ensure atomicity of the move from the user's perspective (no loss of data, flag change, and so on), changes the user's LDAP entry, flushes LDAP caches as necessary, and causes any queued mail to be rerouted to the new store.

Requirements: The following setup and configuration is required before you can use **rehostuser**:

1. Ensure that both the source and destination mailhosts are running Messaging Server 8.0.1 or later.
2. Ensure that both the source and destination mailhosts have IMAP IDLE configured. See "[Configuring IMAP IDLE](#)" for more information.
3. Configure **ssh** on the source and destination mail hosts to allow the mail server user to perform a remote login. Even if you start **rehostuser** as **root**, it still needs to run and execute **ssh** as the mail server user.
4. Make sure that **ugldapbinddn** has read and write access to the **mailHost**, **mailUserStatus**, and **mailMessageStore** attributes. New installations of Messaging Server should have this access made as part of the installation or upgrade process. The **rehostuser** utility checks for these permissions at startup and exits with an error if the permissions are insufficient.
5. The source mailhost must be capable of sending mail to the destination mailhost. In particular, the **tcp_intranet** channel on the source mailhost must be open for relaying, the mail must be routable to the new mailhost (directly or following **mx** records), and the destination mailhost must accept mail coming from the source mailhost.
6. **rehostuser** needs to consider LDAP replication delay. If either the source or destination Messaging Server system (or both) are configured to use an LDAP replica server instead of a master, there may be problems where attribute changes do not show up on the replica as quickly as required. This issue can be addressed as follows:
 - a. If a Messaging Server points to a single Directory Server, configure the Directory Server in multi-master replication (MMR) mode so that it writes updates to its local database and replicates the change. (Note that MMR is the recommended way to setup Directory Server as opposed to the old slave/master model where slaves refer writes to the master, then have to wait for replication to get the updated data).
 - b. If Directory Server failover capability is needed, use Directory Server proxy and enable one of the client affinity modes that guarantees that any single client will always see the updated data it just wrote.
 - c. OpenDS 2.0, (possibly released 03/2009), may offer an *assured replication* feature where once a write is committed, it is guaranteed that the whole farm of replicated

Syntax

```
rehostuser -u userid -d mailhost [-p partition] [-c imskonutil] [-b imbackup]
[-s ssh] [-r imsrestore] [-e] [-o iss_src -t iss_dest -y scp_user [-z src_path]
[-w dest_path]]
```

Options

The options for this command are:

Table 75–35 rehostuser Options

Option	Description	Default	Notes
-u	Specifies the user to move.	NA	NA
-d	Specifies the fully qualified domain name of the destination mailhost.	NA	NA
-p	Specifies the destination partition.	NA	NA
-c	Specifies the imskonutil path	<i>MessagingServer_home/bin.</i>	NA
-b	Specifies the imbackup path	<i>MessagingServer_home/bin.</i>	NA
-s	Specifies the ssh path	/usr/bin/ssh	Do not substitute rsh for ssh . The rsh command cannot pass failure status from the imsrestore command.
-r	Specifies the imsrestore path on the remote host.	<i>MessagingServer_home/bin</i>	NA
-e	Specifies extended availability of user's mailbox during move.	NA	In this mode, Messaging Server performs a full backup and a remote restore without locking the user out of the system. Once the full backup and restore are complete, the user is locked out and Messaging Server performs an incremental backup and restore to take care of any changes that occurred since the start time of the full backup.
-x	Moves the expunged message files.	NA	NA
-n	Does not remove the source mailbox.	NA	NA
-o	Moves the ISS index from the specified source host.	NA	NA
-t	Moves the ISS index to the specified destination host.	NA	NA
-z	Specifies the ISS source installation path.	<i>IndexingServer_home</i>	NA
-w	Specifies the ISS destination installation path.	Source installation path	NA
-y	ISS secure copy (scp) user name.	NA	NA

rehostuser Example

This example shows how to move **user2** to mail server **bigdipper** where the Messaging Server software is installed in the **/opt/sun/comms/messaging** directory and the mail server user is **mailuser**.

```
rehostuser -u user2 -d bigdipper.example.com -r
/opt/sun/comms/messaging/bin/imsrestore -s "/usr/bin/ssh -x -l mailuser"
```

```
disconnecting user2
```

```
-----
Tape Version : 2
Backup Date : 2008/01/08 17:45:16
Message Store : host1.example.com
Block factor : 20
-----
/second/user/user2/INBOX restoring...
/second/user/user2/INBOX/1 restoring...
/second/user/user2/INBOX/2 restoring...
/second/user/user2/INBOX/3 restoring...
/second/user/user2/folder20 restoring...
/second/user/user2/folder22 restoring...
/second/user/user2/folder33 restoring...
/second/user/user2/folder38 restoring...
/second/user/user2/folder49 restoring...
Mailbox user2 copied successfully.
Updated LDAP entry for uid=user2, ou=People, o=example.com, o=usergroup
Source mailbox deleted successfully.
```

relinker

Note: The **relinker** feature is intended to repair the situation where the normal single-copy nature of the message store has become broken for some reason. You should only need to use the **relinker** if you have done something which could have caused duplicate messages to become individual copies instead of the normal single-copy. This feature is not the normal way the store normally accomplishes its single-copy feature. You should not need to keep the real-time **relinker** feature enabled for long periods of time. You should not need to use the **relinker** command on an ongoing basis. You should only need this feature if you have done (or will soon be doing) something which would break the single-copy feature of the store. For more information about single-copy, see the discussion in "[How the Message Store Works.](#)"

relinker finds and relinks duplicate messages. Refer to "[Reducing Message Store Size Due to Duplicate Storage](#)" for more information.

Requirements: You may run **relinker** as **root** or **mailsrv**.

Location: *MessagingServer_home/bin*

Syntax

```
relinker [-p <partitionname>] [-d]
```

Options

The options for this command are:

Table 75–36 *relinker Options*

Option	Description
-d	Specifies that the digest repository be deleted.
-p <i>partitionname</i>	Specifies the partition to be relinked. (default: all partitions).

Examples

To relink a message store:

```
relinker
```

```
Processing partition: primary
Scanning digest repository...
Processing user directories...
```

```
-----
Partition statistics          Before      After
-----
Total messages                73         73
Unique messages              41         40
Message digests in repository 1          1
Space used                   55Kb       51Kb
Space savings from single-copy 40Kb       43Kb
-----
```

Note: run-time relinker (`local.store.relinker.enabled`) is not enabled, so the repository will not be automatically purged. When you're done with relinker, remember to purge the repository by running `"relinker -d"`

After the "Scanning digest repository..." text shown above, **relinker** displays another '.' for every 100,000 messages message digests it finds in the repository and another '.' for every 100,000 messages as it is scanning messages. This indicates the progress of the **relinker** command.

To delete the digest repository:

```
relinker -d
```

```
Processing partition: primary
Purging digest repository...
```

```
-----
Partition statistics          Before      After
-----
Message digests in repository 1          0
-----
```

Note that command-line **relinker** is also controlled by the **store.relinker.maxage** (Unified Configuration) or **local.store.relinker.maxage** (legacy configuration) option, which defaults to 24 (hours). To have command-line **relinker** consider all messages in the store, rather than just those delivered in the past day:

For Unified Configuration, run:

```
msconfig set store.relinker.maxage -v 24
msconfig set store.relinker.maxage -1
```

For legacy configuration, run:

```
configutil -o local.store.relinker.maxage 24
configutil -o local.store.relinker.maxage -v -1
OK SET
relinker
```

```
Processing partition: primary
Scanning digest repository...
Processing user directories...
```

```
-----
Partition statistics          Before      After
-----
Total messages                75          75
Unique messages              40          22
Message digests in repository 1           22
Space used                    51Kb       29Kb
Space savings from single-copy 50Kb       73Kb
-----
```

```
Note: run-time relinker (local.store.relinker.enabled) is not
enabled, so the repository will not be automatically purged.
When you're done with relinker, remember to purge the
repository by running "relinker -d"
# msconfig set store.relinker.maxage 24 (Unified Configuration)
# configutil -o local.store.relinker.maxage -v 24 (legacy configuration)
OK SET
#
```

start-msg

The **start-msg** utility starts all of the messaging server processes (**smtp**, **imap**, **pop**, **store**, **http**, **ens**, **sched**), or optionally, one specified service. Some services started by **start-msg** can be controlled by enabling or disabling the following options: **imap.enable**, **pop.enable**, **http.enable**, **sms_gateway.enable**, **snmp.enable**, **imta.enable**, **mmp.enable**, **ens.enable**, **schedule.enable** (Unified Configuration) or **service.imap.enable**, **service.pop.enable**, **service.http.enable**, **local.smsgateway.enable**, **local.snmp.enable**, **local.imta.enable**, **local.mmp.enable**, **local.ens.enable**, and **local.sched.enable** (legacy configuration). The **ha** option starts the server in HA mode. The watcher monitors process restarts processes on failure. The HA agent monitors the watcher process. In HA mode, the watcher will terminate when it gives up on restarting processes to trigger a failover.

Location: *MessagingServer_home/bin*

Syntax

```
start-msg [ -h | -m | ha | <service>]
```

Options

The options for this command are:

Table 75–37 *start-msg Options*

Option	Description
-h	Shows the usage statement: Usage: start-msg [COMPONENTS...] Start all or some Messaging Server components. Usage: start-msg ha Start all Messaging Server components in HA mode.
ha	Starts all Messaging Server components in HA mode.
<i>service</i>	The accepted values for <i>service</i> are: watcher mfagent ens metermaid store imap pop cert http sched dispatcher job_controller snmp sms mmp mta If no components are specified, all enabled services will be started.
-m	Starts the message store as a replication master.

Examples

The following command starts all the Messaging Server processes:

```
start-msg
```

The following command starts the **imap** process:

```
start-msg imap
```

stop-msg

The **stop-msg** utility stops all Messaging Server processes (**smtp**, **imap**, **pop**, **store**, **http**, **ens**, **sched**), or optionally, one specified service. To use **stop-msg component**, the component must be enabled. The **stop-msg** command without arguments shuts down everything started by **start-msg**, including disabled components.

Location: *MessagingServer_home/bin*

Syntax

```
stop-msg [ -h | ha | [-f] <service>]
```

Options

The options for this command are:

Table 75–38 *stop-msg Command Options*

Option	Description
-f	Force stop using SIGKILL .
-h	Shows the usage statement: Usage: stop-msg [OPTIONS] [COMPONENTS...] Stop all or some Messaging Server components. Usage: stop-msg ha Stop all Messaging Server components in HA mode.
ha	Stops all Messaging Server components in HA mode.
<i>service</i>	The accepted values for <i>service</i> are: watcher mfagent ens metermaid store imap pop cert http sched dispatcher job_controller snmp sms mmp mta If no components are specified, all enabled services will be stopped.

Examples

The following command stops all Messaging Server processes:

```
stop-msg
```

The following command stops the **http** service:

```
stop-msg http
```

stored

The **stored** utility starts a daemon that performs the following functions:

- Performs checkpoint database transactions.
- Deadlock detection and rollback of deadlocked database transactions.
- Cleanup of temporary files and lock files on startup.
- Creates a database snapshot archive.

- Database recovery as necessary (see "Message Store Automatic Recovery On Startup.")

If any server daemon crashes, you must stop all daemons and restart all daemons including **stored**.

Requirements: Must be run locally on the Messaging Server.

Location: *MessagingServer_home/lib/*

Syntax

To run **stored** as a daemon process:

```
stored -r | -R | -t [-v] | [-m] -d [site] | [-m]
```

Options

The options for this command are:

Table 75–39 *stored Options*

Option	Description
-r	Removes the database temporary files and synchronizes the database. This cleans up the database environment to prepare for an upgrade, downgrade, or migration.
-R	Removes the database temporary files without synchronizing the database. If stored -r fails because the software has been upgraded with an incompatible libdb , customers can run stored -R to remove the database temporary files. stored -R does not perform recovery.
-t	Checks the status of stored . The return code of this command indicates the status. To print the status, enter: stored -t -v
-v	Verbose output.
-m	Berkeley Database replication master.
-d site	Delete a mboxlist replication site. Site format is host[:port] . Default site is the local site.