



Scalix Server Administration Guide

Version 11.0.2

Scalix Server Administration Guide

Published by Scalix Corporation
1400 Fashion Island Blvd., Suite 602
San Mateo, CA 94404-2061
USA

Contents copyright © 2007 Scalix Corporation.
All rights reserved.

Product Version: 11.0.2

E: 2.28.2007



Notices

The information contained in this document is subject to change without notice.

Scalix Corporation makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Scalix Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Unix is used here as a generic term covering all versions of the UNIX operating system. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

Linux is a registered trademark of Linus Torvalds.

Red Hat, and Fedora are registered trademarks of Red Hat Software Inc. rpm is a trademark of Red Hat Software Inc.

SUSE is a registered trademark of Novell Inc.

Java is a registered trademark of Sun Microsystems Inc.

Microsoft, Windows XP, Windows 2000, Windows NT, Exchange, Outlook, Active Directory, and Internet Explorer are either registered trademarks or trademarks of Microsoft Corporation in the United States and other countries.

Restricted Rights Legend

Use, duplication, or disclosure is subject to restrictions as set forth in contract subdivision (c)(1)(ii) of the Rights in Technical Data and Computer Software clause 52.227-FAR14.

Contents

Introduction to this Guide	8
Contents of this Guide	8
How to Use This Guide	9
Using the CLI	9
Identifying the Instance Home Directory.	9
Related Documentation	10
Introduction To Scalix	11
About Scalix Product Editions	11
About Scalix User Types	13
Required Licenses	14
Basic Administration Tasks.	15
About This Section	15
This Section's Contents Include:	15
Introducing the Management Console	16
About the Scalix Management Console	16
Managing the System with the Management Console	17
Launching the Management Console	17
Getting Started with the Management Console	18
Key Features of the Management Console.	19
Navigating the Management Console	20
Filtering in the Management Console	20
Logging out of the Management Console	23
Managing Basic System Settings	24
Settings Overview	24
Reviewing Current Settings.	25
Configuring General Settings.	26
Configuring Mailbox Settings.	30
Configuring Password Settings	31
Reviewing and Updating Licenses.	33
Modifying Server Settings Globally.	35
Modifying Scalix Administration Settings.	36
Changing the general.cfg File	38
Managing Mailnodes	40

About Mailnodes	40
Reviewing Existing Mailnodes	41
Creating New Mailnodes	42
Changing Address and Domain Rules for Mailnodes	42
Granting Administrative Access	44
About Administrative Roles and Permissions	44
Scalix Admins	45
The Four Administrative Groups.	46
Group Managers	48
Managing Users	49
About User Accounts	49
Creating New Users	50
Changing Passwords	53
Modifying User Information.	54
Changing Users' Types	56
Deleting Users.	56
Setting Individual Mailbox Capacity Limits	57
Unlocking Users after Failed Login Attempts	58
Enabling Caching on Individual Mailboxes	58
Enabling Search Indexing	59
Identifying Delegates	60
Managing Groups	61
About Scalix Groups.	61
Creating New Groups	62
Adding New Users to Existing Groups.	64
Modifying Groups.	65
Deleting Groups.	66
Assigning Group Managers	66
Logging in as a Group Manager	67
Managing Resources	69
About Shared Resources.	69
Setting Up Shared Resources	70
Booking Shared Resources	71
Changing Shared Resource Settings.	71
Managing Server Processes	73
Introducing Server Info	73
Stopping and Starting Services	75

Monitoring Scalix Services and Daemons	76
Monitoring the Active Users.	77
Monitoring the Message Store	77
Monitoring Message Queues	78
Review the Installation Summary	79
Using Management Plugins.	80
About Management Plugins	80
Running Management Plugins	81
Writing Management Plugins	81
Execution Enviroment.	83
Output Format.	83
Deploying Management Plugins	84
Language Considerations	84
Deployment Script	85
Sample and Template Scripts	86
Advanced Administration Tasks	87
About This Section	87
This Section's Contents Include:	87
Running Backups and Recovery	88
Concepts and Strategies	88
Full Backup and Restore	89
Disaster Recovery	98
Export/Import Backup	99
Single User Recovery	100
Managing Public Folders	101
Public Folder Overview	101
Creating Public Folders	102
Listing Public Folders	102
Permissions for Public Folders	103
Maintaining Public Folders.	105
Assigning Email Addresses to Public Folders	106
Synchronizing Public Folders	107
Forwarding Public Folder Items	111
Posting to Public Folders by Email	112
Hiding Public Folders	112
Public Folder Commands.	112
Access Control Lists	114

About ACLs	114
ACL Commands	115
Creating ACLs	115
ACL Address Patterns	118
Combining Users and Permissions	119
Working with Scalix Directories	121
Directory Overview	121
Structure and Functions of Shared Directories	122
Directory Entries	123
Introduction to Directory Commands	123
Listing Fields In The SYSTEM Directory	124
Creating New Directories	124
Adding and Modifying Directory Entries	125
Searching Directories	125
Using the Client Directory Access Server	126
Creating a Redirect Account	128
About the Redirect Account	128
Creating the Redirect Account	129
Changing Hostnames and IP Addresses	130
Overview	130
Changing a Hostname	130
Changing an IP Address	131
Recovering Deleted Items	132
Overview	132
Recovering Deleted Items	133
Changing the Default Hold Period on a Per User Basis	133
Changing the Default Hold Period on a System-Wide Basis	134
Disabling the Recovery Folder Feature	134
Emptying the Recovery Folder	134
Setting Message Delivery Rules on the Router	136
About Message Delivery Rules	136
About the Service Router	136
Configuring Message Delivery Rules	138
Examples of Message Delivery Rules	144
Listing Deferred Mail	147
Working with SIS	148

Overview	148
Document Types Handled	149
Creating Users' Indexes	149
Disabling Indexing	149
Re-Creating the Index	149
Localizing the Search and Index Service	150
Identifying an Individual Subdirectory	150
Upgrading the Search and Index Service	150
Configuration Options	151
Configuration Files	151
System-Wide Configuration Options	153
Client-specific Configuration Options	210
User-Specific Configuration Options	218
Scalix Command Line Reference Guide	230
Introduction	231
Audit Log Commands	232
Client Directory Access (CDA) Commands	232
Configuration and Installation Commands	233
Directory Commands	233
Directory Relay Server Commands	233
Directory Synchronization Commands	234
Error Manager Commands	235
Event Log Commands	235
Internet Address Commands	235
Internet Mail Gateway Commands	235
LDAP Commands	236
Mailbox Access Commands	236
Mailnode Commands	236
Message Store Commands	237
Miscellaneous Commands	237
Public Distribution List (Group) Commands	238
Public Folder Commands	238
Routing Table Commands	240
Service, Queue and Daemon Commands	240
System Configuration and Maintenance Commands	240
User Entry and Management Commands	241
Glossary	242

Introduction to this Guide

About this Guide

This guide outlines the ongoing use and maintenance of a Scalix mail and calendaring system. It outlines day-to-day administrative tasks such as creating and managing users, groups, calendars, contact lists and public folders. In addition, it provides instructions for essential maintenance jobs such as performing backups, and working with operations queues.

It does not cover one-time setup and configuration tasks such as installing anti-virus or spam protection, setting up multi-server environments or routing email. For more on those topics, see the *Scalix Server Setup Guide*.

Contents of this Guide

Included in this guide are the following topics:

- “Introducing the Management Console” on page 16
- “Managing Basic System Settings” on page 24
- “Managing Mailnodes” on page 40
- “Granting Administrative Access” on page 44
- “Managing Users” on page 49
- “Managing Groups” on page 61
- “Managing Resources” on page 69
- “Managing Server Processes” on page 73
- “Using Management Plugins” on page 80
- “Running Backups and Recovery” on page 88
- “Managing Public Folders” on page 101
- “Access Control Lists” on page 114
- “Working with Scalix Directories” on page 121
- “Creating a Redirect Account” on page 128
- “Changing Hostnames and IP Addresses” on page 130

- “Recovering Deleted Items” on page 132
- “Setting Message Delivery Rules on the Router” on page 136
- “Working with SIS” on page 148
- “Configuration Options” on page 151
- “Scalix Command Line Reference Guide” on page 230

How to Use This Guide

This guide uses the following typographical conventions:

Table 1: Typographical Conventions Used in this Guide

Typographical Convention	Explanation
Buttons	The boldface verdana type indicates a button, a link, a field or any other UI element to click or press as well as a keyboard stroke. For example: Click Finish . Or In the User-name field.
<i>Italics</i>	Italics indicate a directory path or a file. For example: Go to <i>/var/opt/scalix</i> .
Code	This Arial font Indicates a piece of code to write or run. For example: Launch <code>scalix-installer.sh</code>
<i>Document Names</i>	References to other documents appear in italic font.
<Angle Brackets>	Values that you need to supply on your own are shown within angle brackets.

Using the CLI

As with any procedure done on the command line, there may be more than one way to accomplish many of the tasks outlined in this manual. In many cases, these procedures are intended only as examples of how to complete a setup or configuration. If another method is more comfortable or more in keeping with your unique setup, it may be the best approach.

In addition, Scalix offers complete man pages for all commands. Please consult them whenever needed.

Identifying the Instance Home Directory

Throughout the various administrative procedures, there are repeated references to the instance’s home directory, known as “~”. The location of this directory varies depending on how you ran your initial setup. For example, if you named the instance when you created it, the home directory becomes `/var/opt/scalix/<instance>/s`, where `<instance>` is a two-letter code created from the first and last letter of the instance name. If the instance is unnamed, the home directory becomes `/var/opt/scalix/<nn>/s` where `<nn>` is the first and last letter of the host name for that instance.

To determine the home directory for a particular instance, look in `/etc/opt/scalix/instance.cfg` for the appropriate value of `OMDATADIR`.

Related Documentation

Other Scalix product manuals include:

- Scalix Installation Guide
- Scalix Migration Guide
- Scalix Server Setup Guide
- Scalix Client Deployment Guide
- Scalix API Guide
- Scalix Evaluation Guide

In addition, there are online help systems in:

- Scalix Management Console
- Scalix Web Access
- Outlook (if enabled for the Scalix connector)

Introduction To Scalix

Capitalizing on a proven technology foundation and the openness of Linux, Scalix gives enterprise customers a simple to manage, highly reliable, and feature-rich Linux email and calendaring platform. This offers superior price and performance advantages with greater security, reliability, performance, openness and flexibility, when compared to other operating and messaging systems.

Based on open standards and a proven email server technology foundation, Scalix enables customers to create a robust and scalable environment that is flexible enough to adapt to their changing needs over time. The Scalix platform scales up to support organizations with hundreds of thousands of users and scales down for offices with fewer than one hundred users, making it a viable alternative for a broad range of organizations.

The Scalix architecture supports virtually any email client and device, without loss of functionality or data integrity. This means full-function support for popular clients like Microsoft Outlook, Novell Evolution and Domino Lotus Notes, as well as the broad range of POP or IMAP clients available. Users can count on advanced features like enterprise calendaring and scheduling with real-time free/busy lookup, contact and task management, public folders, rich text formatting, offline folder synchronization, secure delegate access to calendar and email, email rules and more.

About Scalix Product Editions

Scalix offers three editions of its powerful email and calendaring platform based on Linux and open systems: Scalix *Enterprise Edition*, Scalix *Small Business Edition* and Scalix *Community Edition*.

Scalix Enterprise Edition is the company's flagship product and is ideal for organizations that demand the full range of functionality in a commercial email and calendaring system. It includes multi-server support, unlimited number of *Standard* users, any number of *Premium* users, the full complement of Scalix advanced capabilities, and a wide variety of technical support options.

Scalix Small Business Edition targets organizations getting started with a commercial version of Scalix that do not have the higher end requirements of Enterprise Edition. It is functionally equivalent to Enterprise Edition except that it allows only single-server installations instead of multi-server, and does not include the capabilities for high availability and multi-instance support.

Scalix Community Edition is the free, single-server, unlimited-use version of the Scalix product and is great for cost-conscious organizations that desire a modern email and calendaring system but do not require advanced groupware and collaboration functionality for

their entire user population. It includes unlimited Standard users, twenty-five free Premium users, a subset of Scalix functionality, and fee-based, incident-based technical support.

The following table compares the Scalix product editions in greater detail:

Table 1: Product Editions and their Features

Product Feature	Community Edition	Small Business Edition	Enterprise Edition
User Types			
Standard Users	Free, unlimited	Free, unlimited	Free, unlimited
Premium Users	Included: 25 Max: 25	Included: 50 Max: Unlimited	Min Purchase: 25 Max: Unlimited
Core Functionality			
Email & calendaring Server	Single-server	Single-server	Multi-server
Internal user directory	[X]	[X]	[X]
Choice of GUI-based or command line installation and administration	[X]	[X]	[X]
Unlimited POP/IMAP email client access	[X]	[X]	[X]
Native MS Outlook support (via MAPI)	Premium users only (max 25)	Premium users only	Premium users only
Fully functional AJAX web client (Scalix Web Access)	[X] (group scheduling in calendar for 25 premium users only)	[X] (group scheduling in calendar for all premium users)	[X] (group scheduling in calendar for all premium users)
Native Novell Evolution support	[X] (group scheduling in calendar for 25 premium users only)	[X] (group scheduling in calendar for all premium users)	[X] (group scheduling in calendar for all premium users)
Public folders	Premium users only (max 25)	Premium users only	Premium users only
High availability	Not available	Not available	[X]
Multiple instances per server	Not available	Not available	[X]
Migration tools	Not available	[X]	[X]
Upgrade To Enterprise Edition	Via license key. Re-installation not required	Via license key. Re-installation not required	Not applicable
Mobile Access	[X]	[X]	[X]
Ecosystem Support			

Table 1: Product Editions and their Features

Meta-directory support via LDAP	[X]	[X]	[X]
iCal support	[X]	[X]	[X]
Native Exchange Interoperability (via TNEF)	Not available	[X]	[X]
Active Directory integration with MMC plug-in	Not available	[X]	[X]
Anti-virus	Via flexible 3rd party interface	Via flexible 3rd party interface	Via flexible 3rd party interface
Anti-spam	Via flexible 3rd party interface	Via flexible 3rd party interface	Via flexible 3rd party interface
Archiving	Via flexible 3rd party interface	Via flexible 3rd party interface	Via flexible 3rd party interface
Wireless email & PIM	Email-only via POP/IMAP	Email & PIM via Notify	Email & PIM via Notify
Technical Support			
Community Forum	Free	Free	Free
Knowledgebase, Tech notes	Free	Free	Free
Incident-based Support	Fee-based	Fee-based	Fee-based
Software subscription	Not available	[X]	[X]
Premium 7x24 Support	Not available	[X]	[X]
Cost			
Licensing	Free, unlimited use	\$995 for First 50 Premium Users	Per-user License; No Per-server Fees

About Scalix User Types

Scalix users can be defined as *Standard* or *Premium* users, as defined in the following:

Standard Users

Standard users gain access to a subset of Scalix functionality including email, personal calendar and contacts through Scalix Web Access and Novell Evolution as well as email access using POP/IMAP clients. The ability to deploy standard users is ideal for cost-conscious organizations with users who do not have high-end groupware and collaboration requirements. An unlimited number of standard users may be deployed with any Scalix edition for free.

Premium Users

Premium users have access to the full benefits and functionality of the Scalix email and calendaring system. The following Scalix product capabilities are available only to premium users:

- Native MS Outlook support (via MAPI)
- Group scheduling functionality including free/busy lookup in Outlook, Scalix Web Access and Evolution clients
- Access to public folders
- Wireless email and PIM

Any number of licensed premium users may be deployed with Scalix Enterprise Edition. Scalix Community Edition is limited to a maximum of twenty-five (25) free premium users, who enjoy many of the features available to Enterprise Edition premium users.

Flexible, Cost-Effective Email For Everyone

The distinction between standard and premium users provides organizations with the flexibility to cost-effectively provide email for all users. For example, manufacturers and retailers may desire headquarters staff to be designated as premium users as they require advanced groupware capabilities, while less demanding users, such as shop floor or store personnel, would be satisfied as standard users with only email and personal calendaring capabilities. Similarly, educational institutions may decide that faculty and staff are premium users that need advanced collaboration capabilities while students are standard users that just need email and personal calendaring. There is no cost for deploying standard users with either Scalix Community Edition or Scalix Enterprise Edition.

Required Licenses

Enterprise and Small Business editions require a license key whereas Community Edition does not. Additionally, if you are a Scalix Community Edition customer, you can only perform the “typical” installation, in which all the Scalix components are stored on a single host computer.

If you went through the installation without a license key, your system installed as Community Edition and your users as Standard until the correct license key is entered in the Scalix Management Console. For more on how to enter a license key through the Management Console, see the *Scalix Administration Guide*.

Basic Administration Tasks

About This Section

The tasks outlined in the next few chapters are of a more “basic” nature and apply to all setups, single- or multi-server, and all product editions. All are performed in the Scalix Management Console (SAC). These basic tasks include the setting up and managing of users, accounts, groups, services and daemons.

The second half of this manual, Advanced Administration Tasks, covers procedures done only from the command line. Those include scheduling and running backups, managing public folders and more.

This Section’s Contents Include:

Included in this section are the following topics:

- “Introducing the Management Console” on page 16
- “Managing Basic System Settings” on page 24
- “Managing Mailnodes” on page 40
- “Granting Administrative Access” on page 44
- “Managing Users” on page 49
- “Managing Groups” on page 61
- “Managing Resources” on page 69
- “Managing Server Processes” on page 73
- “Using Management Plugins” on page 80

Introducing the Management Console

This chapter introduces the Scalix Management Console (aka SAC), its features, navigation, login and logout procedures.

Contents

This chapter covers the following topics:

- “About the Scalix Management Console” on page 16
- “Managing the System with the Management Console” on page 17
- “Launching the Management Console” on page 17
- “Getting Started with the Management Console” on page 18
- “Key Features of the Management Console” on page 19
- “Navigating the Management Console” on page 20
- “Filtering in the Management Console” on page 20
- “Logging out of the Management Console” on page 23

About the Scalix Management Console

The Scalix Management Console (aka SAC) is a browser-based application that assists in the day-to-day administration and management of a Scalix messaging system. It is a separate component of the system and can be accessed by any of the approved browsers (see the Installation Guide for details) on either Microsoft Windows or Linux workstations.

The console enables many common system administration tasks on all three editions of the Scalix system. It also provides access to all servers on a network through a single GUI interface. This enables both single-server management and global changes with the ease of one centralized interface.

By contrast, the command line interface only manages one server at a time.

For server setup tasks or high-end, advanced maintenance, use the extensive Linux-based command line interface, which is introduced in the second half of this guide. The CLI provides a full set of commands and allows the setup and running of administration scripts.

Managing the System with the Management Console

Once you have installed and configured your Scalix server, and migrated over existing user data, you can begin basic maintenance processes. The Management Console enables a wide range of maintenance tasks that fall into these categories:

- User account management
- Group management (public distribution lists)
- Management of shared resources such as conference rooms, printers and projectors
- Starting and stopping server services and daemons
- Monitoring queues
- Changing server configurations
- Monitoring the system to assess the state of processes, resources and the load on Scalix queues

You also can extend the console to perform many customized tasks through the use of Management Plugins, which are scripts that perform many frequently-repeated actions such as checking CPU or disk usage, scanning logs, etc.

Launching the Management Console

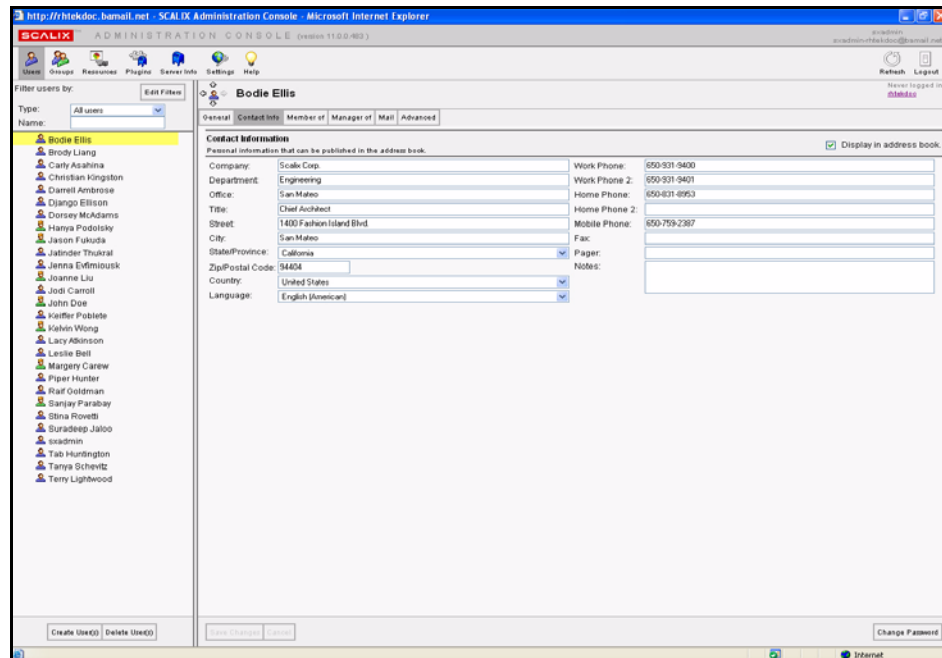
To start the Management Console:

- 1 Start a Web browser.
See the *Scalix Installation Guide* for a complete listing of approved browsers.
- 2 In the browser **Location/Address** field, type the URL of the Scalix server.
The URL must include a “/sac” extension, as shown here:
`http://<your_scalix_mailserver_FQDN>/sac`
- 3 When the Scalix login page appears, enter your login ID (sxadmin@fqdn) in the **Name** field.
(This was created during Scalix installation.)
- 4 Enter the related sxadmin password in the **Password** field.
- 5 Click **Go**.
 - If the **Go** button is inactive, click the checkbox by “Not using a secure HTTPS connection. Click to continue”. The Go button should become active.

Note

This checkbox appears if the Apache server (part of the Scalix system) was not set up to support SSL.

- 6 If your login is successful, the Management Console appears in the browser.



Alert

Because different users have different levels of access to the Management Console, not everybody will see the same features. The screenshot is above is for a fully enabled administrator.

Getting Started with the Management Console

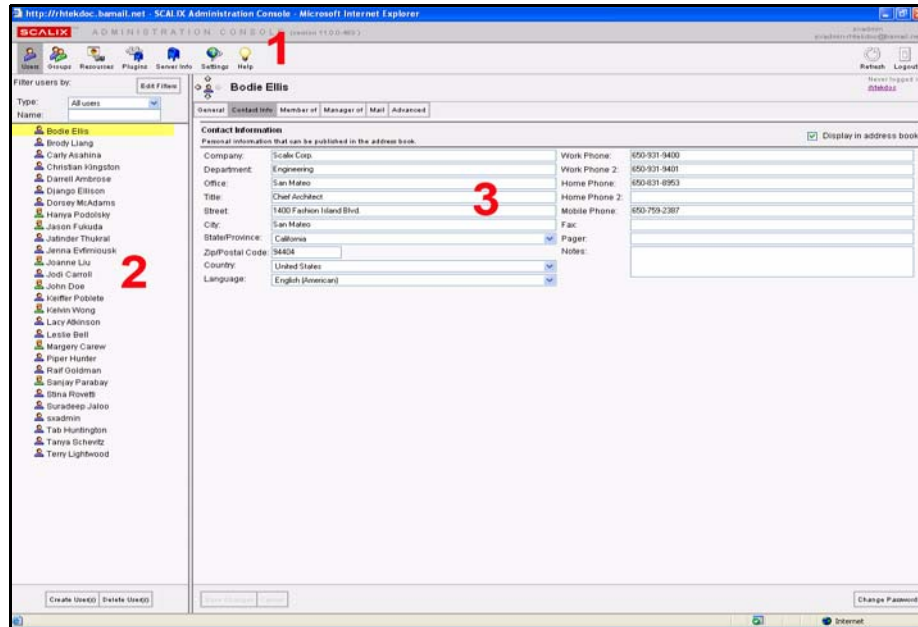
You can get started on either of two task-specific tracks:

- If you migrated a lot of existing users (and their mail data) from a previous mail server, you need to perform the following basic tasks:
 - Add any new users or groups
 - Edit any existing users or groups
 - Check the status of Scalix services and queues, including a quick check of the main Message Store
 - And use the Management Console to review (and fine-tune) a subset of Scalix settings
- Or, if you are setting up Scalix as a brand-new server and have no existing mail users, you must do most of the same tasks, but without editing any existing user records. In this scenario, adding users is your primary workload.
 - Add any new users or groups
 - Check the status of Scalix services and queues
 - Review and adjust a subset of Scalix settings

Key Features of the Management Console

Use of the Scalix Management Console is efficient and uncomplicated. As shown in the following illustration, there are three areas in the Management Console interface. Use each one (as numbered) to complete basic tasks.

- 1 Toolbar
- 2 Contents pane
- 3 Display pane




To use each section of the Management Console:

- 1 Click one of the toolbar buttons—**Users**, **Groups**, **Resources**, **Plugins**, **Server Info** and **Settings**.
- 2 In the contents pane, a list of items related to that toolbar button appears. For example, after clicking **Users**, a list of users appears in the contents pane, whereas after clicking **Server Settings**, a list of servers displays.
- 3 To view details about any item in the contents pane, click it. Its settings and parameters appear in the display pane to the right. You can make needed additions or modifications there.

Each display pane is organized into tabs that offer task-specific options.

Navigating the Management Console

Once you click a toolbar button and select a contents item to review, the navigation button atop the display pane helps you speedily switch tasks and/or lists.



- Click the up or down arrows to scroll through items in the contents pane to the left.
- Click the left or right arrows (like a browser's Back or Forward buttons) to move to previous and next displayed items.

Filtering in the Management Console

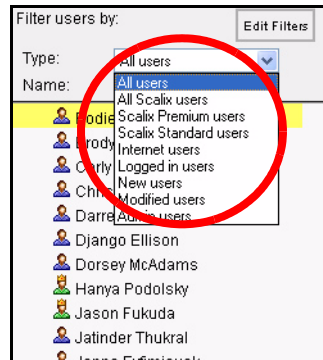
At some point, you may have thousands of Scalix users to manage, and dozens, if not hundreds of users, groups/public distribution lists or servers to monitor. Trying to display or wade through all of them in the contents pane to find the one you need can be cumbersome. To better target the records you need and to speed up response time, you can filter.

Scalix allows you to filter by two main parameters:

- **Type:** From the drop-down list, select whether you want to view all users, Premium users, Standard users, Internet users, Logged in users, etc. For groups, you can choose between direct and effective members, all users and all groups.
- **Name:** Type in part of all of the name of the user or group that you want to find. Only those matching the pattern you type in appear in the contents pane.

To open and use filters:

- 1 On the toolbar, select what you want to filter: **Users**, **Groups**, **Servers**, etc.



- 2 As the contents pane begins to populate, either filter by **Type** of user you want to view, or type in part or all of the **Name** you want to find. The contents pane filters dynamically as you type or select from the drop-down menu.

The different filters by **Type** are:

Table 1: Filters by Type

User Type	Filter
All Users	Lists all users in the system (the default)
All Scalix Users	Lists all Scalix users, whether Premium or Standard
Scalix Premium Users	Lists users who log in to an Enterprise Edition system, or who have use of one of the 25 "Premium" accounts on a Community Edition system
Scalix Standard Users	Lists only those who log in to the Community Edition with Scalix Web Access or an IMAP mail client, or who log in as "Standard" users on an Enterprise Edition system
Internet Users	Lists all users without local mailboxes, but whose address is included in the Scalix Directory. All such non-Scalix users are labeled with a distinct icon
Logged In Users	Lists all users currently connected to Scalix
New Users	Lists all users of any type added during this session
Modified Users	Lists all users modified during the current session
Admin Users	Lists all "administration access" users (no matter what their level of permissions)

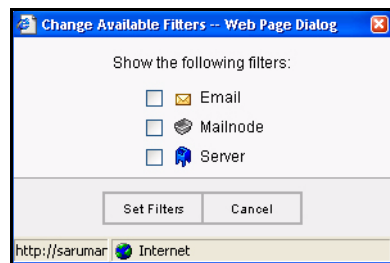
Adding Other Filter Options

If the filters offered in the drop-down menu are not exactly what you want, you can add others.

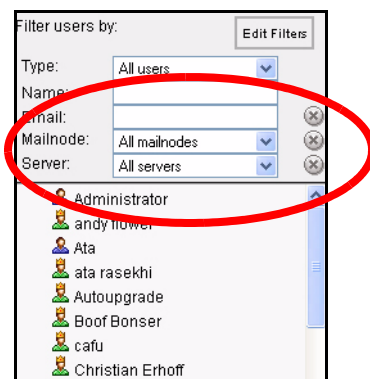
To add other filtering options:

- 1 Above the contents pane, click **Edit Filters**.

- 2 The Edit Filters dialog box appears.



- 3 Click the checkbox by each filtering option you want to activate. The options are:
 - **Email:** Enter either the full address or the domain name to narrow the number of results that display
 - **Mail node:** Enter the name of the node through which the user's account is routed
 - **Server:** Enter the server on which the user's account is located
- 4 Click **Set Filters**.
- 5 The newly activated filtering options appear in the contents pane.



To close the new filter, click the X to the right.

To delete text from the field, click the smaller X that appears when you begin typing.

To use the other filtering options:

- 1 Click in the **Email**, **Mail Node** or **Server** fields and type in the address or select from the drop-down menu the name of the user, node or server you want to find.
When searching by name or email address, partial entries or the wild card symbol narrow the field enough to help you locate the resource you need.

Combining Filters

Often, the best practice is to use more than one filter at a time. For example, you can narrow the field of possible returns by filtering out all but Premium users, then type in the first few letters of the name of the user you want to find.

To combine filters:

- 1 Pick a **Server**, **Node** or **Type**, or enter part of a name.

- 2 As you type, the Management Console dynamically filters the records based on your text entry.
- 3 Typing in more fields or applying more menu commands continues to dynamically focus the current list in the contents pane.

Alert


When you apply filters, remember that the more you specify filtering characteristics, the fewer results you'll get. If you want to locate all groups with the word "money" in their name, you won't need any other filter besides the Name field. If you want to locate all users named "smith" in the engineering node on the ENGRG server, you can set up those filters, but this excludes any other "Smith" engineers on other nodes or servers.

Logging out of the Management Console

Although it is safe to "log out" of Scalix by simply quitting your browser, we recommend following the logout process to prevent later problems.

To log out:

- 1 In the upper right-hand corner, click **Logout**.



- 2 A logout confirmation dialog box appears.
- 3 At the confirmation prompt, click **Yes**.

The Management Console connection is cleanly closed and the Management Console Login page displays.

You can log in to the Management Console (on this or on another server), or exit the browser.

Managing Basic System Settings

This chapter covers basic system settings performed through the Scalix Management Console. That includes such settings as domain names, the format for email addresses and authentication IDs, password configurations, license management, mail node settings, global server settings, caching, and more.

Contents

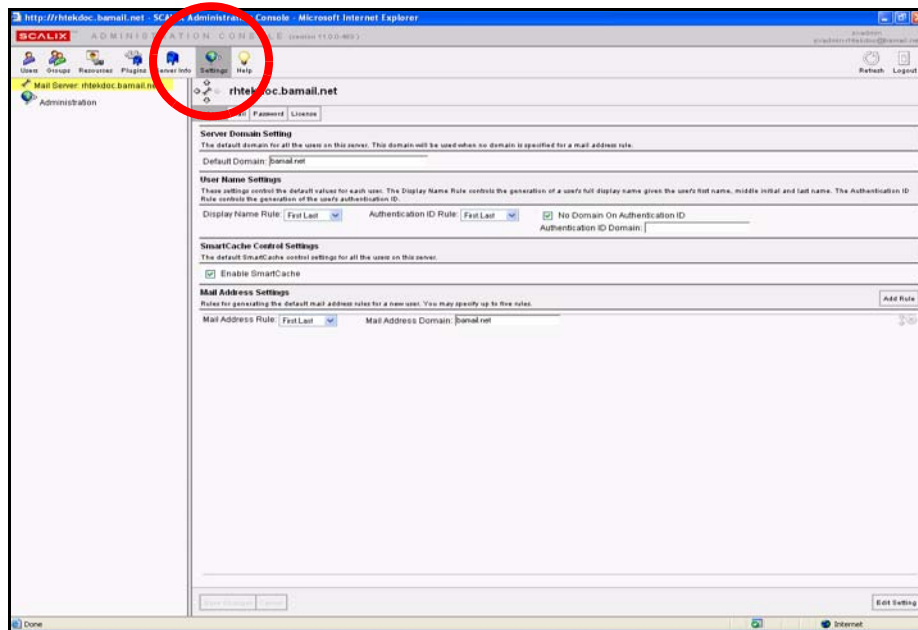
This chapter includes the following information:

- “Settings Overview” on page 24
- “Reviewing Current Settings” on page 25
- “Configuring General Settings” on page 26
- “Configuring Mailbox Settings” on page 30
- “Configuring Password Settings” on page 31
- “Reviewing and Updating Licenses” on page 33
- “Modifying Server Settings Globally” on page 35
- “Modifying Scalix Administration Settings” on page 36
- “Changing the general.cfg File” on page 38

Settings Overview

The Settings feature in the Management Console provides an efficient approach to many ongoing server settings and configurations for email addresses, passwords and authentica-

tion IDs, licenses, mail nodes, global server settings, caching, and more.



Reviewing Current Settings

With the settings feature, you can review settings for all servers in the network.

To review current server configuration/setting options:

- 1 In the Management Console toolbar, click **Settings**.
- 2 The contents pane displays:
 - [Server Name]: The name of the server(s) you are administering
 - **Administration**: A workspace with several general options
 - **Global**: With two or more servers, a **Global** entry also displays. It enables modifications to all servers at the same time.
- 3 In the contents pane, select the name of the server whose settings you want to review.
- 4 That server's settings options appear in the display pane to the right.
- 5 The settings options are organized into four tabs:

Table 1: Server Settings Options Tabs

Tab Name	Tasks Performed
General	<ul style="list-style-type: none"> • Change the default domain name for a Scalix server. • Change user name-specific settings (for display name and authentication ID). • Change the default format for email addresses.

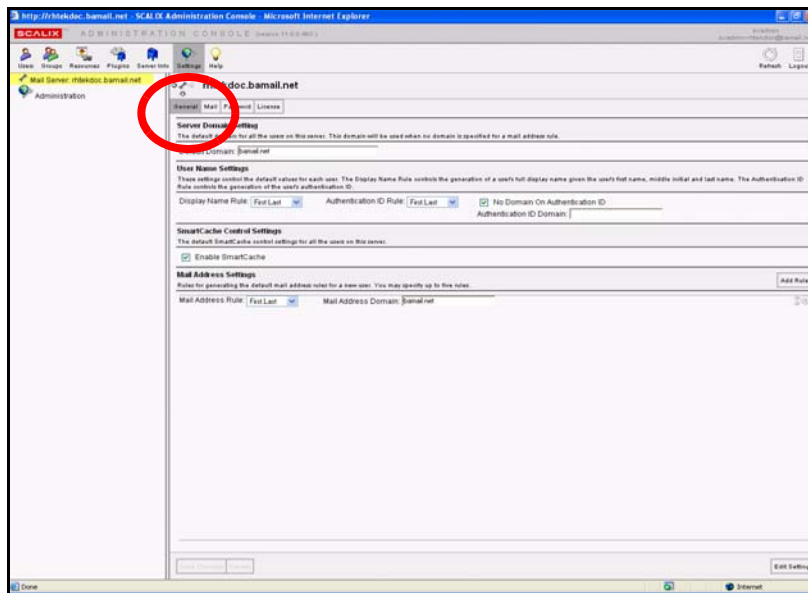
Table 1: Server Settings Options Tabs

Mail	<ul style="list-style-type: none"> Set a general limit for system-wide user mailbox capacity. Set an individual's mailbox capacity, over-riding the system setting. Program the system response to a full mailbox. Determine the frequency of "out of office" mailings from user's clients.
Password	<ul style="list-style-type: none"> Set the formatting and length of user-set passwords. Determine the lifetime of a password. Accept/block re-use of a current password. Set the number of failed login attempts before a user is rejected.
License	<ul style="list-style-type: none"> Add a new Scalix license or review the conditions of the current license.

6 Each of these tasks is detailed separately in the following sections.

Configuring General Settings

The General settings tab of the Management Console enables configuration of domain names, email addresses, authentication IDs and caching.



Changing Server Domain Name Settings

By default, Scalix uses a server's domain name to generate email addresses and authentication IDs for all users on that server. Using the Management Console, that setting can be changed, though.

For example, if the domain name is "founder.net", the default format for new email addresses will be [first.lastname]@founder.net. (Note that previous domain names remain in effect with pre-existing mail accounts.)

For administrators who want to change the format, Scalix allows two types of address generation rules:

- Global (system-wide) address-generation rules
- Mail node-specific address generation rules (For information on mail nodes, see “Managing Mailnodes” on page 40.)

Scalix stores up to five global address-generation rules, but only one mail node address-generation rule—which overrides any global rules currently in use.

To change server domain settings:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to change the default server domain setting.
- 3 In the display pane, select the **General** tab.
- 4 In the field titled, **Default Domain**, delete the current domain name.
- 5 Type the replacement name.
- 6 At the bottom of the display pane, click the now-active **Save Changes** button.

Setting Rules for User Names and Authentication IDs

You can modify the rules that control how a new user’s name is formatted, both as a display name, which determines how a user name displays on email messages, and as an authentication ID, which can be used to sign on to the system. The display name and authentication ID both take from the user’s name (first and last and - optionally - middle initial) in different formats that can be customized.

There are three settings you can change:

- **Display name rule:** The “friendly” name that displays on emails. This is not the email address
- **Authentication ID rule:** The name or identification used to sign on to the system
- **Domain on authentication ID rule:** Determines whether the user must enter the full domain name upon sign-in

Modifying Display Name Settings

To modify display name rule settings:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server whose display name rules you want to modify.
- 3 In the display pane, click the **General** tab.
- 4 Under the **User Name Settings** options, open the **Display Name Rule** drop-down list.
- 5 Select the preferred user name format.
- 6 At the bottom of the display pane, click the now-active **Save Changes** button.

To modify authentication ID rule settings:

- 1 In the Management Console toolbar, click **Settings**.

- 2 In the display pane, click the **General** tab.
- 3 Under the **User Name Settings** options, open the **Authentication ID Rule** drop-down list.
- 4 Select the preferred authentication ID format.
- 5 At the bottom of the display pane, click the now-active **Save Changes**.

To require use of the full domain name when logging in:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the display pane, click the **General** tab.
- 3 Under the **User Name Settings** options, remove the check mark from the box before the **No Domain on Authentication ID** field.
If the "check box is marked, unmark it to enable the use of the full domain name.
- 4 At the bottom of the display pane, click the now-active **Save Changes** button.

Alert

If you want full domain names used in authentication IDs, be sure to communicate to your users that they must update their client account settings accordingly.

Enabling Caching for all Users on a Server

To improve the speed and responsiveness of email service in Outlook, use the Scalix SmartCache feature. This creates copies of all users' mailboxes on their client machines as well as on the server, allowing them to work off of the local machine, speeding up client performance and preserving bandwidth.

With SmartCache, the system checks in to the server only when sending and receiving new messages, meaning fewer trips to and from the sever.

The server-wide setting done here can be overridden on a user-by-user basis. In addition, there are different ways to prepare a user's mailbox for caching. For more on how to set individual users' SmartCache settings, see "Enabling Caching on Individual Mailboxes" on page 58.

To enable SmartCache for all users on a server:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to enable caching.
- 3 In the display pane, click the **General** tab.
- 4 Under the **SmartCache Control Settings** options, put a check mark before **Enable SmartCache**.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Setting Rules for Email Address Formats

By default, Scalix assigns email addresses in the form of `firstname.last@domain.com`. But you can change that.

The mail address settings option allows you to select from 12 different name formats including first.last, first.middleinitial.last, last.first and more. You can pick any rule, then modify the domain name appended to the address. You can create up to five rules at a time.

To set rules for email address formats:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to set email address formats.
- 3 In the display pane, click the **General** tab.
- 4 Under the **Mail Address Settings** options, open the **Mail Address Rule** drop-down menu.
- 5 Select the format you prefer.
- 6 Review the text in the **Mail Address Domain** field and change it if needed. (It is based on the domain name recorded in the **General** tab.)
- 7 At the bottom of the display pane, click the now-active **Save Changes** button.

Note

The previous procedure assumes a server-wide application. If you customize an e-mail address formatting on a mail node, and a user is associated with that selected mail node, then his/her address is generated not by global server-wide rules, but by the mail node-specific rule.

Adding Other Email Address Settings

If needed, you can add additional email address settings.

To add other, alternate email address formats:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to add email address settings.
- 3 In the display pane, click the **General** tab.
- 4 Under the **Mail Address Settings** options, click **Add Rule**.
- 5 A new, blank row appears below the existing row(s).

Mail Address Settings			
Rules for generating the default mail address rules for a new user. You may specify up to five rules.			
Mail Address Rule:	First.Last	Mail Address Domain:	click.bamail.net
Mail Address Rule:	First.Last	Mail Address Domain:	

- 6 Fill in the settings—both address rule and domain name.
- 7 At the bottom of the display pane, click the now-active **Save Changes** button.

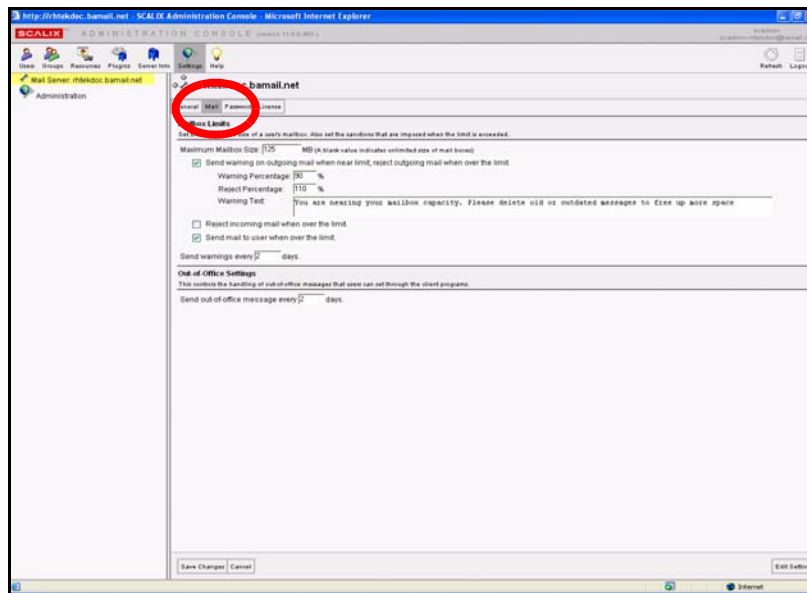
You can use these alternate mail addresses later to create or modify user accounts.

Note

The previous procedure assumes a server-wide application. If you customize an email address formatting on a mail node, and a user is associated with that selected mail node, then his/her address is generated not by global server-wide rules, but by the mail node-specific rule.

Configuring Mailbox Settings

Using the Management Console, you can set limits on mailbox capacity, program warnings to alert users when they are nearing their capacity, limit the number of out-of-office notifications sent, and more.



Setting Size Limits on Mailboxes

The Management Console enables limiting the size of users' mailboxes to prevent anybody from overusing limited server resources.

If needed, you can override these system settings on an individual mailbox level. For more information on how to override system-wide limits, see "Setting Individual Mailbox Capacity Limits" on page 57.

To set capacity limits on all mailboxes on the system:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to set mailbox size limits.
- 3 In the display pane, click the **Mail** tab.
- 4 Under the **Mailbox Limits** options, locate the **Maximum Mailbox Size** field and type in a number representing megabytes of data.
- 5 Click the checkboxes by the alert responses you want to activate.

If you check **Send warning on outgoing mail...** fill in the percentage at which you want automatic warnings sent and mail rejected. Then type in a short warning message that clearly outlines what will happen when their Inbox reaches capacity.

- 6 In the **Send Warnings Every...** field, type the number of days between automatic alerts.
- 7 At the bottom of the display pane, click the now-active **Save Changes** button.

Limiting Out-of-Office Replies

To avoid inundating recipients with out-of-office replies, limit the number of replies sent during a given time period. You can set the system to send only one out-of-office reply per day, or every other day, or any other number of days in between.

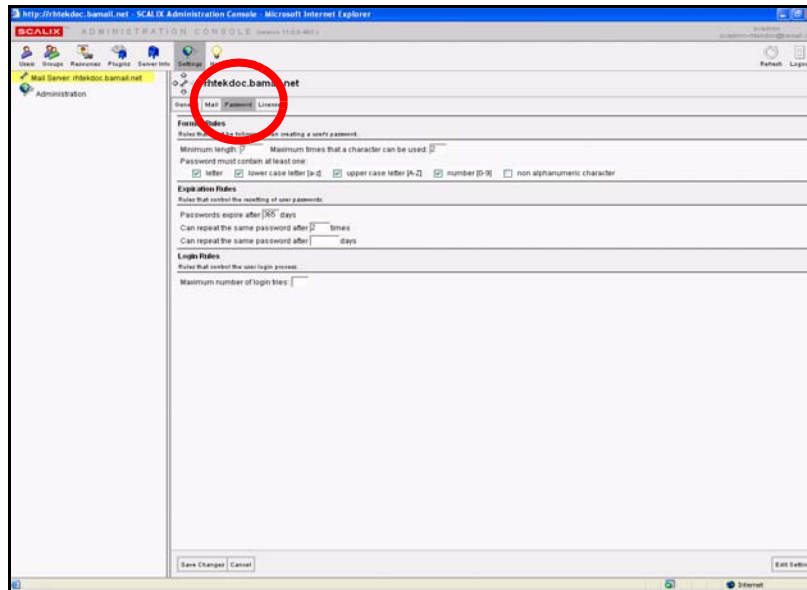
For example, if someone sends five emails on the first day a user is on vacation, they receive only one out-of-office reply.

To limit out-of-office replies:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to limit out-of-office replies.
- 3 In the display pane, click the **Mail** tab.
- 4 Under the **Out-of-Office Settings** options, type the number of days between which an out-of-office reply should be sent.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Configuring Password Settings

Using the Management Console, you can set password parameters, including the minimum number of characters required, character inclusions, repetition patterns and expiration periods. You also can set the number of times an attempted connection fails before the system blocks further attempts by the same user.



Creating Password Format Rules

Password format rules establish the minimum length of passwords, the maximum number of times any one character can be used, and many other character inclusion rules. For example, you can require that passwords have both lower case and upper case letters as well as numbers or other characters.

To set password format rules:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to set password format rules.
- 3 In the display pane, click the **Password** tab.
- 4 Under the **Format Rules** options, fill in the blanks, including:
 - Enter a minimum number of characters, letters or numbers.
 - Enter the number of times a character can be repeated, if at all.
 - Check all the character-inclusion options to be applied.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Password Expiration Rules

Password expiration rules establish the number of days after which passwords expire, and the number of days and times after which users can re-use old passwords.

To set password expiration rules:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to set password expiration rules.
- 3 In the display pane, click the **Password** tab.
- 4 Under the **Expiration Rules** option, fill in the blanks, including:

- The time period (in days) before passwords automatically expire.
- The number of times after which a user can re-use an old password.
- The number of days after which a user can re-use the same password.

5 At the bottom of the display pane, click the now-active **Save Changes** button.

Login Rules

Login rules set the number of failed login tries that trigger a system rejection of that user.

If a user is rejected, he or she is warned that they must wait a set period of time before trying again.

To set login rules:


- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, select the server for which you want to set login rules.
- 3 In the display pane, click the **Password** tab.
- 4 Under the **Login Rules** option, type the number that represents the limit.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Alert

If a user exceeds the maximum number of login tries, they are locked out of the system. This requires that they contact the system administrator and request that their mailbox be unlocked. This unlocking procedure is detailed in "Unlocking Users after Failed Login Attempts" on page 58.

Reviewing and Updating Licenses

You may want to review your current Scalix license, and at the same time, upgrade the license (such as when you need to increase the number of users in your system). You can use the Management Console to read the license conditions and to enter and import the text of a license upgrade.



Reviewing Current Licenses

Using the Management Console, you can review your license and, if needed, reassess your company's needs, extend or upgrade your license.

To review a current license:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, click the **Mail Servers** globe icon. For more on viewing global settings by clicking the mail servers globe icon, see "Modifying Server Settings Globally" on page 35. To see the license information for an individual server, click that server's name in the contents pane.
- 3 The **License** tab displays:
 - License type and lifespan
 - Network identifiers
 - Numbers of Premium and Standard users..

Note


Licenses for all servers should be the same across a Scalix system.

Importing Licenses into Scalix

When first installing Scalix or upgrading a license, license keys are delivered by e-mail. Using the Management Console, you can import the license text into the server.

To copy and import the license text into the server:

- 1 Open the e-mail message containing the license key text.
- 2 In the Management Console toolbar, click **Settings**.
- 3 In the contents pane, click the **Mail Servers** globe icon. For more on creating global settings by clicking the mail servers globe icon, see the next section, "Modifying Server Settings Globally" on page 35.
- 4 In the lower right-hand corner of the **License** tab, click **Install License**.
- 5 An Install License window appears.



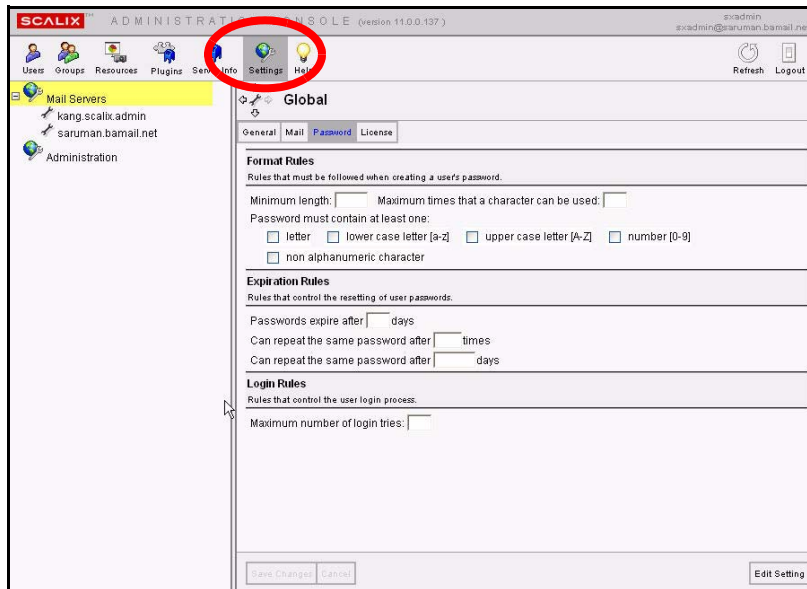
- 6 Copy the entire license text from the e-mail.

- 7 Click in the Install License window and paste the text.
- 8 Click **Install License**.
- 9 After the license is successfully imported, the **License** tab reappears, displaying a summary of the new license. This license remains in effect until upgraded, if needed.

Modifying Server Settings Globally

If you are managing two or more Scalix servers, you can take advantage of the Global settings option, which makes universal changes to all servers. Customizing settings globally helps reduce repetitious tasks and ensure consistency in configurations and settings across all Scalix servers.

If you are working globally, Scalix displays a server-like set of options, starting with the General tab— just as it does for individual servers. You can then modify needed settings at a global cross-server level instead of repeatedly adjusting them for each server.



Alert

If you first make changes to all servers "globally", you can later override any server-specific settings by making individual modifications to a single server.

To do global settings:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, click the **Mail Servers** globe icon.

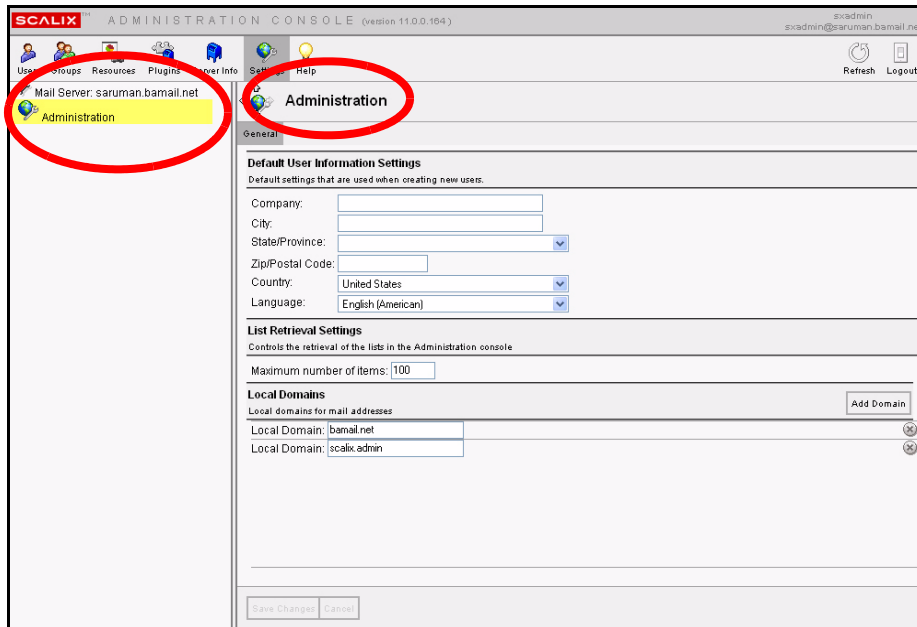


- 3 The Global Settings options appear in the display pane to the right. A full set of tabs lets you modify the same range of settings for all the Scalix servers that you would modify for a single server. For more on each setting, see the corresponding sections above.

Modifying Scalix Administration Settings

Administration settings are default values that the Management Console uses when creating users. They include customizing default user information and changing the list retrieval account, among other settings.

These settings do not affect users created outside the Scalix system.



Customizing Default User Information Settings

If your user base is restricted to a distinct office or area, you can use these options to pre-load user information, saving time and effort later in data entry. This is helpful if most of your users work for the same company or live in the same city, state, or country. You can set the fields to pre-populate with the information that is common to most of your users, and allow the others to overwrite as needed.

You also can preload the language entry if all your users share a common language.

To set default user information:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, click **Administration**.
- 3 Under the **Default User Information Settings** option, type in the information that most of your users have in common or select it from the appropriate drop-down menu. The fields you can set default information for are:
 - Company
 - City

- State/Province
- Zip Code
- Country
- Language

- 4 At the bottom of the display pane, click the now-active **Save Changes** button.

Changing the List Retrieval Count

If you have more than 100 users or groups, you see “Incomplete List” displayed in the contents pane when you first click the relevant toolbar buttons. You can raise the limit with this option, to force Scalix to list more of the available names at one time.

Note that this slows down the loading speed of Scalix records in the Management Console.

To increase the number of user or groups listed in the contents pane:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, click **Administration**.
- 3 Under the List Retrieval Settings option, locate the **Maximum number of items** field and delete the “100” (or whatever number is in this field).
- 4 Type the replacement value (as a whole number).
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Adding Local Domains

If you want to create additional local domains for use in managing your user base, you can add as many as needed, provided your Scalix license incorporates them.

Alert

If you want to add a domain that is not within the terms of your license, you have to get a new license or upgrade the current one.

To add local domains:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, click **Administration**.
- 3 Under the Local Domains options, click **Add Domain**.
- 4 A new, blank row appears.

Local Domains		
Local domains for mail addresses		
Local Domain:	click.bamail.net	X
Local Domain:		X

- 5 Click in the **Local Domain** text field and type the domain name.
- 6 At the bottom of the display pane, click the now-active **Save Changes** button.

These domains are now available in a pull-down menu in the Create New User wizard, when adding or modifying a user account.

Changing the *general.cfg* File

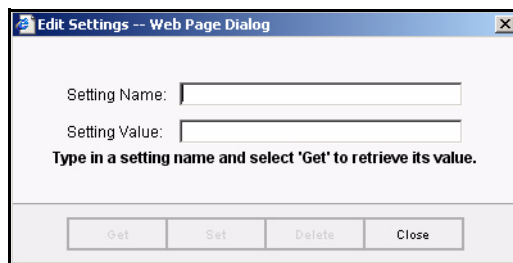
Modifying settings in the *general.cfg* file is best done through the CLI. You can review and modify the general configuration settings stored in the “general.cfg” file with the Management Console. The options are procedurally limited and require a precise knowledge of both the configuration context and the range of extensions you can apply.

This procedure follows two general steps:

- Querying a specific *general.cfg* setting to determine what values have already been set.
- Modifying the setting, if you choose, or entering an all-new setting with its value.

To query a configuration with the Management Console:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, click the name of the server for which you want to make changes.
- 3 In the **General** tab, which displays by default, click the **Edit Settings** button in the lower right corner.
- 4 The Edit Settings dialog box appears.



- 5 In the **Setting Name** field, enter the exact configuration name.
- 6 Click **Get**.
- 7 The query results display in the **Setting Value** field. This field displays the value, such as a whole number or a text entry such as TRUE or FALSE, BEFORE or AFTER, etc.
- 8 If the configuration does not exist or if the text was entered incorrectly, Scalix displays the following message:

No value found for this setting. You may type in a value and click 'Set' to store it.

Entering a New Configuration Setting

If you queried a Scalix general configuration and got the “no value...” message, you can type a value for the named setting and save it to the *general.cfg* file at this time.

To enter a new configuration setting:

- 1 In the Management Console toolbar, click **Settings**.

- 2 In the contents pane, click the name of the server for which you want to make changes.
- 3 In the **General** tab, which displays by default, click the **Edit Settings** button in the lower right corner.
- 4 The Edit Settings dialog box appears.
- 5 With the Edit Settings dialog box open, make sure the exact configuration name is correctly typed in the **Setting Name** field.
- 6 Type the appropriate value in the **Setting Value** field.
- 7 Click **Set** to save the new settings.

After the dialog box closes, the new settings are in effect.

Modifying an Existing Configuration Setting

You can change a configuration value in the Edit Settings dialog box after successfully querying the original setting.

To modify an existing configuration setting:

- 1 In the Management Console toolbar, click **Settings**.
- 2 In the contents pane, click the name of the server for which you want to make changes.
- 3 In the **General** tab, which displays by default, click the **Edit Settings** button in the lower right corner.
- 4 The Edit Settings dialog box appears.
- 5 In the **Edit Settings** dialog box, query the current configuration value.
- 6 Check the results to verify that a change is needed.
- 7 Replace the current value in the **Setting Value** field with your preferred value.
- 8 Click **Set**.

After the dialog box closes, the new settings are in effect.

This guide includes a list of many configuration options and their default values/extensions, plus a range of values where applicable. For more on `general.cfg` settings, see the chapter titled, "Configuration Options" on page 151.

Note

Do not modify settings in the `general.cfg` file without consulting the CLI section of this guide to first understand the ramifications. Call Scalix support if you are unsure of the effects of making changes to the `general.cfg` file.

Managing Mailnodes

This chapter covers mailnodes, a unique Scalix option that organizes an e-mail user base. It explains how to create mailnodes, change their configurations and more.

Contents

This chapter includes the following information:

- “About Mailnodes” on page 40
- “Reviewing Existing Mailnodes” on page 41
- “Creating New Mailnodes” on page 42
- “Changing Address and Domain Rules for Mailnodes” on page 42

About Mailnodes

Mailnodes are a unique Scalix feature that organizes mail user communities into manageable groups. For example, you can organize by work group, employment status or office location. Each Scalix server is associated with a mailnode, which is created during installation, and you can create additional nodes now in the Management Console, including customizing any new mailnodes with a specific Internet address and domain name.

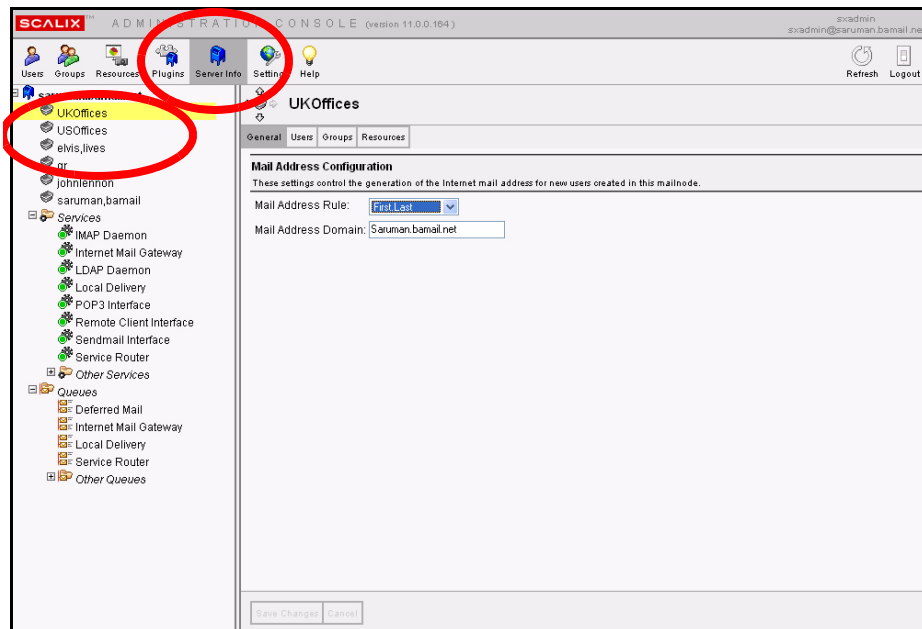
Once you create the needed nodes (a preliminary task), you can do the following, depending on your e-mail user base:

- If starting out with a new system and adding all your users at this time, you can sort your new users into any mailnodes as you work.
- If you already have an established user base, you can't sort the existing user records into new nodes (in the Management Console), but you can sort all newly added users into nodes.

In addition, there are other mailnode settings you can customize. You can:

- Change a mailnode's address configurations
- Review the list of users associated with this node
- Review the groups associated with this node

Note that once a user is associated with a specific node, you cannot move or reassign them in the Management Console. This is done on the Scalix CLI because it requires migrating the user and mail contents/data from one node to another.




Reviewing Existing Mailnodes

Before adding or changing mailnode configurations, it is helpful to review what you have. Using the Management Console, you can review all of your mailnode settings.

To review the status of a mailnode:

- 1 In the toolbar, click the **Server Info** button.
- 2 In the contents pane, select the node whose status you want to review. In a position subsidiary to the blue postal mailbox, you see at least one gray bin. That is a mailnode.



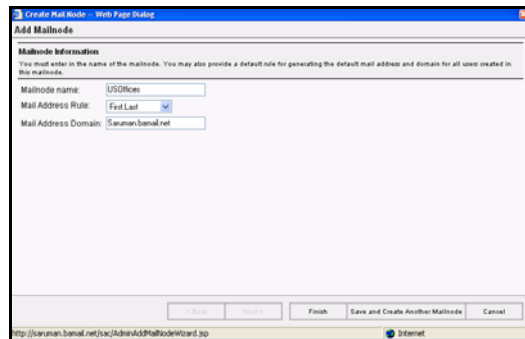
- 3 When you select a node, its options appear in the display pane to the right. The options are organized into three tabs, **General**, **Users**, **Groups** and **Resources**.
- 4 To review general settings, click the **General** tab.
- 5 To see all users assigned to a mailnode, click **Users**.
To see information about any individual user, click its name.
- 6 To see all groups assigned to a mailnode, click **Groups**.
To see information about any individual group, click its name.
- 7 To see all resources assigned to a mailnode, click **Resources**.
To see information about any individual resource, click its name.

Creating New Mailnodes

You can use the Management Console to create additional mailnodes on a server, including customizing new nodes with a specific internet addresses and domain names.

To create additional mailnodes:

- 1 In the toolbar, click the **Server Info** button.
- 2 In the contents pane, select the server on which you want to create additional mailnodes. The server is represented by a blue postal mailbox.
- 3 That server's mailnode options appear in the display pane to the right.
- 4 In the display pane, click the **Mailnodes** tab.
- 5 When the server's mailnode information appears, review the listed nodes. Initially, only one is listed. That is the one created during installation.
- 6 To add another node, click **Add Mailnode** in the lower right corner of this tab.
- 7 The Create Mailnode dialog box appears.



- 8 Fill in the following information:
 - **Mailnode Name:** Enter a single-word name/title for this new node.
 - **Mail Address Rule:** Pick your preference for mail address formatting from this menu.
 - **Mail Address Domain:** Fill in the preferred domain information.
- 9 To close the dialog box and save the node information, click **Finish**.
- 10 Back in the **Mailnodes** tab, click the **Finish** button there, too.
- 11 The new node appears in the tab, but not in the contents pane. To make it display, click the **Refresh List** button at the bottom of the contents pane.

Changing Address and Domain Rules for Mailnodes

Using the Management Console, you can customize nodes to use specific Internet addresses and domain names.

If you create an e-mail address-generation rule on a mailnode, and if a user is associated with that selected mailnode, then his/her address is generated not by global rules, but by the mailnode-specific rule.

Reminder

Nodes can be entered/used only if you've included them in the scope of your Scalix license.

To change node configurations:

- 1 In the toolbar, click the **Server Info** button.
- 2 In the contents pane, select the node on which you want to change configurations.
- 3 In the display pane, select the **General** tab. There are two options you can change:
 - Mail Address Rule
 - Mail Address Domain
- 4 To change the mail address rule, open the **Rule** pull-down menu and choose your preference.
- 5 To change the existing domain (or assign one to this node), type the domain in the **Domain** field.
- 6 At the bottom of the display pane, click the now-active **Save Changes** button.

Granting Administrative Access

This chapter explains the different levels of administrative access to the Management Console: What the different levels of access are, how to grant them, and when to use them. It does not cover how to perform actions once given administrative access. That is addressed in later chapters.

Contents

This chapter includes the following information:

- “About Administrative Roles and Permissions” on page 44
- “Scalix Admins” on page 45
- “The Four Administrative Groups” on page 46
- “Group Managers” on page 48

About Administrative Roles and Permissions

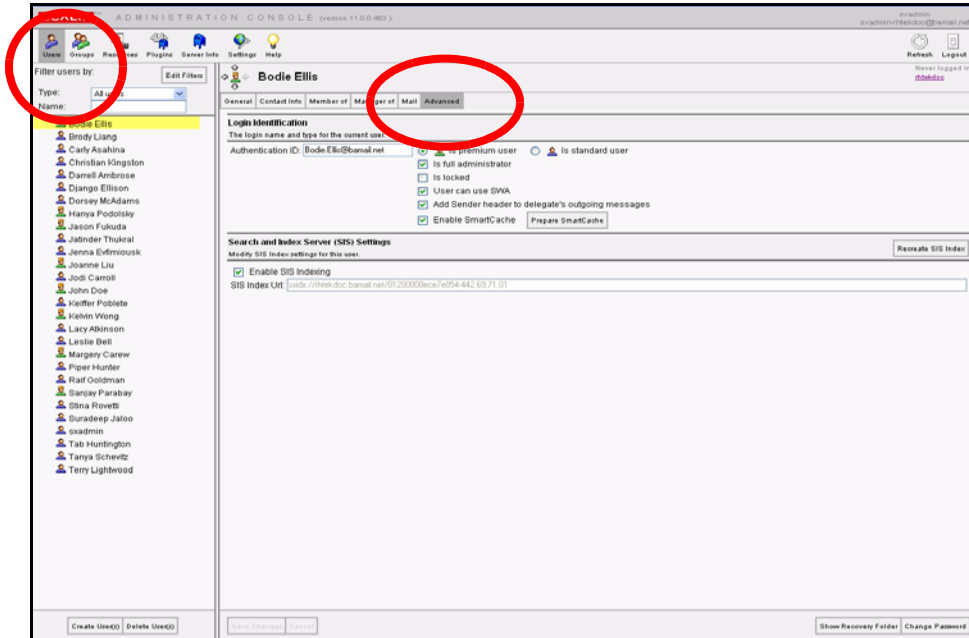
Scalix provides four distinct administrative roles that allow decreasing levels of access and permission to manage some or all aspects of a Scalix system:

- **Root:** Can run any commands on the CLI for all machines in a Scalix network, but cannot sign in to the Management Console because he/she is not a Scalix user. This user does not come into play until the second half of this manual.
- **Full Administrator:** Has complete access to the Management Console and can run most commands on the CLI. One full administrator, `sxadmin`, is created by default during installation. Using the Management Console, you can create as many others as desired.
- **Scalix Admin Groups:** These users have varying levels of access to the Management Console to undertake specific (and limited) tasks. There are four Admin groups, each one overseeing a different aspect of the system: General settings, users, groups and user attributes.
- **Group Manager:** A member of a public distribution list/group who has limited Management Console access. They can add or delete users only to group(s) for which they are specifically designated as managers.

If you sign on to the Management Console as the default administrative user, sxadmin, you can delegate the lesser roles to others.

Note

Users' Management Console passwords are the same as their email client passwords.



Scalix Admins

Scalix Admins, also known as full administrators, have access to all features and aspects of the Management Console. In addition, they can run most commands on the CLI. They also can create and manage other users with lower levels of access.

You can create as many full administrators as you want

Note

Full administrators are not listed in any of the four Scalix Admin groups—and do not need to be.

Creating Scalix Admins

To grant Scalix Admins status, which amounts to full administrative access, the user must have a fully-functioning Scalix account.

To assign full administrator permissions:

- 1 In the Management Console toolbar, click the **User** button.
- 2 The contents pane lists all current users.

Tip

Remember that you can use the Filter menu to make the list more manageable.

- 3 Select the user to which you want to assign full administrator permissions.
- 4 When the user's account information appears in the display pane, click the **Advanced** tab.



- 5 In the Login Identification options, click the checkbox next to **Is full administrator**.
- 6 Click **Save Changes**.

The Four Administrative Groups

Administrative groups are created and managed through the Groups feature in the Scalix Management Console. The four groups have varying levels of access to the Management Console to undertake specific (and limited) tasks. Members of these groups can use the Management Console to perform their duties, but are not permitted CLI access.

All admin group members must have fully-functioning Scalix user accounts and must use one of the approved browsers on either Windows or Linux.

These groups' responsibilities include:

Table 1: The Four Scalix Admin Groups

Roles	Management Console Access
Scalix Admins	Members of this group have permission to see and use all Management Console features.
Scalix User Admins	Members of this group see only the user functions and can: <ul style="list-style-type: none"> • Add new users • Modify existing users • Delete existing user
Scalix Group Admins	Members of this group see only the user and group functions and can: <ul style="list-style-type: none"> • Add new groups • Modify existing groups • Add and delete members of groups
Scalix User Attributes Admins	Members of this group see only user features and can: <ul style="list-style-type: none"> • Edit the personal contact information in a user account

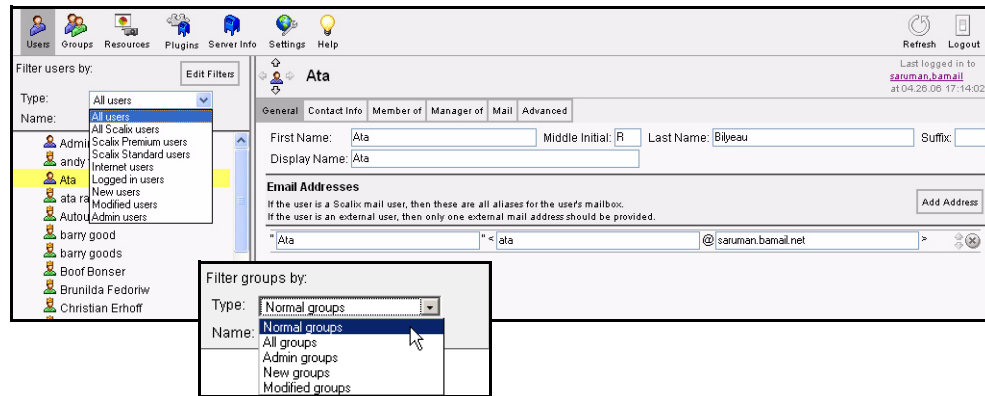
Reviewing Current Admin Groups

By navigating through the filters, menus and tabs, you can see the users assigned to the different administrative groups and their properties.

To view the current Scalix admin groups and their membership:

- 1 In the Management Console toolbar, click the **Groups** button.

- 2 The contents pane displays the current roster of groups.



- 3 Open the **Filter** menu (highlighted above) and choose **Admin Groups**.
- 4 The list of four admin groups appears in the contents pane
- 5 Select each group name to view the members in the display pane.

Assigning Users to Admin Groups

To grant membership in one of the four administrative groups, the user must have a fully-functioning Scalix account.

To assign users to an administrative group:

- 1 In the Management Console toolbar, click the **User** button.
- 2 Select the user to be given administrative access.
- 3 When the account information appears in the display pane, click the **Member Of** tab.
- 4 In the **Member Of Groups** options, open the **Filter users by: Type** menu and choose **All Groups**.
- 5 Locate the relevant group and click the check box by its name.
- 6 Click **Save Changes**.

An alternative method for adding users to administrative groups:

- 1 In the Management Console toolbar, click the **User** button.
- 2 Open the **Filter groups by: Type** menu and choose **Admin Groups**.
- 3 When the four groups appear in the contents pane, select the group you want to add members to.
- 4 That groups options appear in the display pane.
- 5 Click the **Members** tab to view the current roster.
- 6 Open the **Filter Members by: Type** menu and choose **All Users**.
- 7 When this tab fills in with all the current users, use the Filter by **Name** field to isolate your candidate member.

- 8 When the user's account displays, click the empty checkbox to add him or her to the group.
- 9 Click **Save Changes**.

Verifying New Admin Group Members

Before moving on, verify that the user is properly added to the group.

To verify a new member of a group:

- 1 In the Management Console toolbar, click the **User** button.
- 2 Open the **Filter users by: Type** menu and choose **Admin Groups**.
- 3 When the four groups appear in the contents pane, select the group you just added members to.
- 4 That group's options appear in the display pane.
- 5 Click the **Membership** tab to view the current roster.
- 6 Your newly added members should be listed.

Group Managers

The final - and lowest - level of administrative access is the group manager. This user can sign in to the Management Console to see and work with only those groups that he or she manages. This user can add or remove users from a single group/public distribution list.

Creating Group Managers

To grant group manager status, the user must have a fully-functioning Scalix account.

To designate a user as a group manager:

- 1 In the Management Console toolbar, click the **User** button.
- 2 Select the user to be given administrative access.
- 3 When the account information appears in the display pane, click the **Manager Of** tab.
- 4 In the **Member Of Groups** option, put a check mark by the group this user will manager.

Tip

Remember that you can use the group filtering menu to make the list of groups more manageable.

Managing Users

This chapter covers users: The creation, management of, and deletion of user accounts. It also addresses user types, passwords and mailbox capacity limits.

Contents

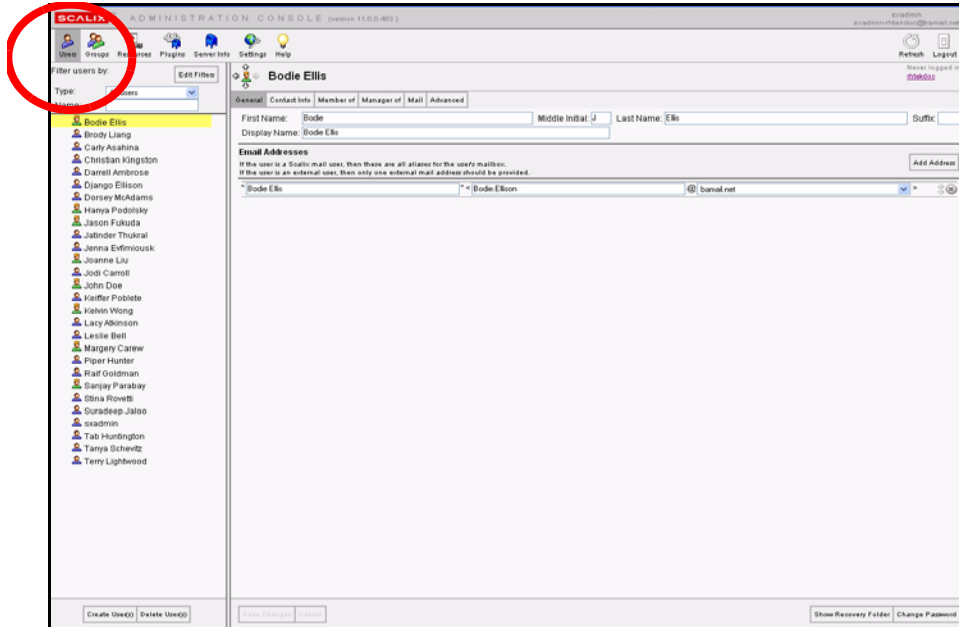
This chapter includes the following information:

- “Creating New Users” on page 50
- “Changing Passwords” on page 53
- “Modifying User Information” on page 54
- “Changing Users’ Types” on page 56
- “Deleting Users” on page 56
- “Setting Individual Mailbox Capacity Limits” on page 57
- “Unlocking Users after Failed Login Attempts” on page 58
- “Enabling Caching on Individual Mailboxes” on page 58
- “Enabling Search Indexing” on page 59

About User Accounts

You can create and add three types of users to the Scalix server:

- **Scalix Standard User:** Has limited access to Scalix features. See About Scalix Product Editions for more information.
- **Scalix Premium User:** Has access to all Scalix features including full Outlook support through MAPI, public folders and group scheduling features. See About Scalix Product Editions for more information.
- **Internet Mail User:** Serves as a contact database alias, only, to redirect mail for a Scalix “user” to an external Internet mailbox.




Creating New Users

Using the Management Console, you can create new Scalix users, set their level of service, give them authentication IDs, enter their personal information, assign them to specific mail-nodes and more.

To create new users:

- 1 In the Management Console toolbar, click the **Users** button.



- 2 All current Scalix users display in the contents pane.
- 3 At the bottom of the contents pane, click the **Create User(s)** button.

- 4 The Create New User wizard appears.

- 5 Select the type of user for this account:
- **Scalix Standard User:** Has limited access to Scalix features.
 - **Scalix Premium User:** Permits use of both Outlook or SWA, and has access to all Scalix features.
 - **Internet Mail User:** Serves as an alias to redirect mail to an external Internet mailbox.

Alert

If you are administering a Scalix Community Edition server, remember that you have a maximum limit of 25 "Premium" user accounts for use. (Upgrading to Enterprise Edition overrides that limit.)

- 6 Enter the **First Name**, **Middle Initial** and **Last Name** for the user you are creating. Scalix uses this information to automatically fill in the following fields:
- Display name
 - Mail address
 - Authentication ID

Note

All three fields adapt the user's First Name, Last Name and Middle Initial according to rules (formulas) set up in the General tab of the Settings options. You can customize these fill-in rules according to your preferences, as detailed in "Reviewing Current Settings" on page 25

If the Display name, Mail address, and/or the Authentication ID fields cannot be edited, this is due to an existing rule used by the server to auto-generate the values in those fields. As noted before, you can modify those rules as detailed in "Reviewing Current Settings" on page 25.

Alert

If illegal characters are entered in this field, the text turns red.

- 7 If you are creating an "Internet mail user" account, you must enter the user's fully qualified email address.
- 8 Open the **Mailnode** pull-down menu and select the appropriate Scalix mailnode, if more than one exists.
- 9 Enter a **Password** (and its confirmation) to be initially linked with the user's Authentication ID.

Note

You cannot proceed through the Create New User Wizard until you successfully enter the password correctly in both password validation text fields.

- 10 Click the checkbox by **User Must Change Password on first login** if you want to require the user to create a new password when they log in to the Scalix messaging system for the first time.

This completes the essential user account information. You now have the following options:

- Click **Next** to add Contact Information for this user.
 - Click **Finish** to add the user to the system. You can add this user to relevant groups at a later time.
 - Click **Save and Create Another User** to add this user to the system (with this level of data) and immediately reopen the Create New User wizard.
- 11 Click the checkbox by **Is Locked**, if you want this user account locked (which disables logins).
 - 12 If you completed the basic user information entry, and clicked **Next**, the second Add User screen appears. This is where you enter contact information.

- 13 Fill in all blank fields, or use any available pull-down menus to fill them in.
- 14 After you complete the new user contact information, you have the following options:
 - Click **Next** to add this user to any groups.

- Click **Finish** to add the user to the system. You can enter contact and group information at any later time.
- Click **Save and Create Another User** to add this user to the system (with this level of data) and immediately reopen the Create New User wizard.

If you completed the contact information entry and clicked **Next**, the third (and final) Group Membership wizard appears.

All of the existing Scalix public distribution lists/groups are listed in this wizard, accompanied by empty checkboxes.

- 15 If you want to list this user in any existing public distribution lists/groups, click the checkbox by each group name.
- 16 After making any group additions, you have two options:
 - Click **Finish** to add the user to the system.
 - Click **Save and Create Another User** to add this user to the system and immediately reopen the Create New User wizard.

For more information about groups, their creation and management, including adding new members, see “Managing Groups” on page 61.

Changing Passwords

To change the password for an existing Scalix user account:

- 1 In the Management Console toolbar, click the **Users** button.



- 2 In the contents pane, select the user whose password you want to change.

Tip

Remember to use the filtering feature if needed. If filtering by name or type doesn't give the results you want, you can filter by mailnode: In the Management Console toolbar, click Server Info. The contents pane displays all servers and their respective nodes. Select the node whose users you want to review. The node is depicted by a gray postal letter bin. When the node options appear, click the Users tab. For more on filtering, see "Filtering in the Management Console" on page 20.

- 3 In the display pane, select the **General** tab.
- 4 In the lower right-hand corner, click **Change Password**.
- 5 The Change User Password dialog box appears.



- 6 Enter the new password and then confirm it by typing it in the field below.
- 7 If desired, click the checkbox by **User must change password on first login**. This prompts the user to create a new password the next time they log in to the Scalix server with this newly changed password.
- 8 Click the now-active **Change Password** button.

Modifying User Information

No matter how you initially created a Scalix user account—with the Management Console, by migrating existing records or a bulk provisioning—you can use the Management Console to open, review and modify the personal information stored in a current user account.


Changing Contact Information

If needed, you can change users' company, address, department, phone numbers and more.

To change the contact information for a user:

- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, select the user whose information you want to change.

- 3 In the display pane, modify the information in the **Contact Info** tab.



- 4 At the bottom of the display pane, click the now-active **Save Changes** button.

Changing Email Addresses

You can change users' email addresses after they have already been assigned.

To change users' email addresses:


- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, select the user whose email address you want to change.
- 3 In the display pane, select the **General** Tab.
- 4 Under the **Email Addresses** options, modify the alias and domain values.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Assigning More than One Email Address to Users

You also can assign more than one Internet address to a user.

To give users' more than one email address:

- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, select the user to whom you want to give the additional email address.
- 3 In the display pane, select the **General** Tab.
- 4 On the right side of the display pane, click the **Add Address** button.



- 5 A blank entry space appears below the current addresses.
- 6 The user's current display name appears by default. If needed, modify the text in the **Display Name** field.
- 7 Enter an alias for this user.
- 8 Open the drop-down menu and choose a valid domain (covered by your Scalix license).
- 9 You can shuffle the entries in the address list. (Whatever address is at the top of the list becomes the default address that the system uses for all messages addressed to or sent by this user.)
- 10 At the bottom of the display pane, click the now-active **Save Changes** button.

Deleting Email Addresses

Over time, you may need to delete users' email addresses.

To delete an address from a user's address list:


- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, select the user whose address you want to delete.
- 3 In the display pane, select the **General** Tab.
- 4 Highlight the email address and click the "X."
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Changing Users' Types

Using the Management Console, you can either upgrade Standard users to Premium or vice versa. For more on the differing levels of service, see "About Scalix User Types" on page 13.

To change a user's account type:

- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, select the user whose type you want to change.
- 3 In the display pane, select the **Advanced** tab.
- 4 Click the radio button by the preferred option: **Premium** or **Standard**.



- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Deleting Users

IMPORTANT: When you delete a user, you remove the record and all related mail and scheduling archives. This task requires a confirmation before it is completed.

To delete a user account from Scalix:

- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, select the user(s) whose account you want to delete. If you want to delete more than one, use the control or shift keys to select as many as desired.
- 3 Click **Delete User(s)**.
- 4 When a Confirmation dialog box appears, click **Yes** to proceed.

Deleting All Users from a Mailnode

The following procedure is “global”—if you follow this procedure, all users assigned to a node are deleted.

To globally delete all users from a mailnode:

- 1 In the Management Console toolbar, click the **Server Info** button.
- 2 In the contents pane, click the mailnode whose users you want to delete.
- 3 In the display pane, click the **Users** tab.
- 4 Locate the user record to be removed and click **Delete Listed Users**.
- 5 At the bottom of the contents pane, locate and click the **Refresh List** button.

Setting Individual Mailbox Capacity Limits

Using the Management Console, you can limit the capacity of individual mailboxes. In addition, you can control the type and frequency of warnings the user gets when theyhe or she exceeds that capacity.

Any individual limits set in the Users screen override server or global limits created in the Settings screens.

Scalix recommends setting all preferences at a global level in a multiple-server environment, at which time you can override this limit at the individual user account level. For example: If all your users have a 100 MB capacity limit, you can easily grant the CEO a Gigabyte of capacity.

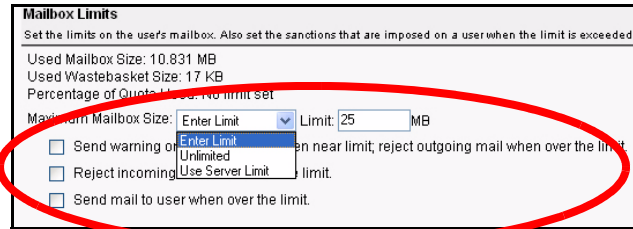
For more on creating global or server-wide mailbox capacity limits, see “Setting Size Limits on Mailboxes” on page 30.

Alert

Once you override global or server-wide mailbox capacity settings on an individual mailbox level, you cannot use the Management Console to revert that user back to the default setting. You must make that change on the command line.

To set limits on individual users' mailboxes:

- 1 In the toolbar, click the **Users** button.
- 2 In the contents pane, select the account for whom you want to set individual limits.
- 3 When the user's options appear in the display pane to the right, click the **Mail** tab.
- 4 In the area of the display pane labelled Mailbox Limits options, click in the **Maximum Mailbox Size** and select one of the three options:
 - From the drop-down manu, select **Enter Limit** then type in a size in MB.
 - From the drop-down menu, select **Unlimited**.
 - From the drop-down menu, select **Server Limit**.



Mailbox Limits
Set the limits on the user's mailbox. Also set the sanctions that are imposed on a user when the limit is exceeded.

Used Mailbox Size: 10,831 MB
Used Wastebasket Size: 17 KB
Percentage of Quota Used: No limit set

Maximum Mailbox Size: Limit: MB

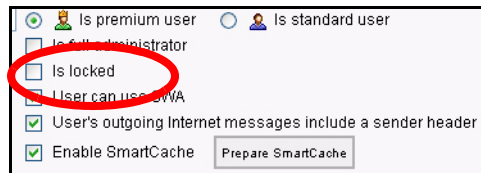
☐ Send warning on near limit, reject outgoing mail when over the limit.
☐ Reject incoming mail when over the limit.
☐ Send mail to user when over the limit.

- 5 Click the check boxes by the alert responses you want to activate.
- 6 Click the now-active **Save Changes** button at the bottom of the display pane.

Unlocking Users after Failed Login Attempts

After a user has attempted to log in to the Scalix system (to their mailbox) and have reached the limit for failed attempts, they are automatically locked out of the system. At that time, a server-generated alert will recommend that they contact their IT department for readmission to the system (and probably a password reset.) You can unlock this user and allow them access to the system by following these steps:

- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, locate the user who is locked out.
- 3 In the display pane, Click the **Advanced** tab.
- 4 Among the Login Identification options, look for the **Is Locked** check box—which should have a check mark.



☒ Is premium user ☐ Is standard user
☐ Is full administrator
☒ Is locked
☐ User can use SVA
☒ User's outgoing Internet messages include a sender header
☒ Enable SmartCache

- 5 Clear the check mark.
- 6 Click **Save Settings**.

Tip

You may want to reset the user's password at this time, which you can do in the **General** tab (as detailed previously in "Changing Passwords" on page 53).

Enabling Caching on Individual Mailboxes

To improve the speed and responsiveness of a user's email service in Outlook, use the Scalix SmartCache feature. This creates a copy of the user's mailbox on his or her client machine as well as on the server, allowing them to work off of the local machine, speeding up client performance and preserving bandwidth.

With SmartCache, the system checks in to the server only when sending and receiving new messages, meaning fewer trips to and from the sever.

You can set SmartCache for all users on a server by enabling it through the server setting explained in “Enabling Caching for all Users on a Server” on page 28. Or you can set it on a user-by-user basis here, which overrides the system-wide setting.

If preferred, you can download the user’s entire mailbox in one single file, which speeds up the initial cache creation process. This method, initiated by clicking the Prepare Smart-Cache button, is optional and best used for large mailboxes. If you choose not to use this method, the cache creates one message at a time the next time that user signs on to Outlook.


The Prepare SmartCache button is only used when the cache is first created, or when the user changes computers and the cache has to be rebuilt. From then on, the user’s messages cache each time they sync with the server.

Alert

Once you override server-wide SmartCache settings on an individual mailbox level, you cannot use the Management Console to revert that user back to the default setting. You must make that change on the command line.

To enable SmartCache for individual users:

- 1 In the Management Console toolbar, click **Users**.
- 2 In the contents pane, select the user for whom you want to enable caching.
- 3 In the display pane, click the **Advanced** tab.
- 4 Under the **Login Identification** options, put a check mark before **Enable SmartCache**.



- 5 At the bottom of the display pane, click the now-active **Save Changes** button.
- 6 At a convenient time, create the mailbox on your local machine by clicking the **Prepare SmartCache** button. This is a resource-intensive procedure, so if your mailbox is large, wait until you have the time and bandwidth to allow the process to complete. The amount of time it takes depends on the size of your mailbox.

Enabling Search Indexing

To enable searching for messages, contacts and appointments in the SWA client, you must enable search indexing for each, individual user. This is done in the Management Console. Also, if the user’s search index corrupts, you can repair it there.

To enable search indexing:

- 1 In the Management Console, click **Users**.
- 2 In the contents pane, select the user for whom you want to enable indexing.

- Under the Search and Index Service (SIS) settings option, put a check mark by **Enable SIS Indexing**.

- At the bottom of the display pane, click the now-active **Save Changes** button.
- After the index creates, its location appears in the **SIS Index URL** field. This location cannot be changed.
- If the index corrupts, click the **Recreate SIS Index** button.

Identifying Delegates

Using the Management Console, you can set the system to identify delegates in the header of an email message. With this setting checked, outgoing messages sent by a delegate have a header identifying the actual sender.

Note

This setting applies only to the SWA client. For more information on how to set this in Outlook, see Outlook's online help system.

To identify messages as from delegates:

- In the Management Console, click **Users**.
- In the contents pane, select the user for whom you want to enable identification as a delegate.
- Under the Login Identification options, put a check mark by **User's outgoing internet messages include a sender header**.

- At the bottom of the display pane, click the now-active **Save Changes** button.

Managing Groups

This chapter covers the creation, administration and use of groups (aka public distribution lists) in the Management Console.

Contents

This chapter includes the following information:

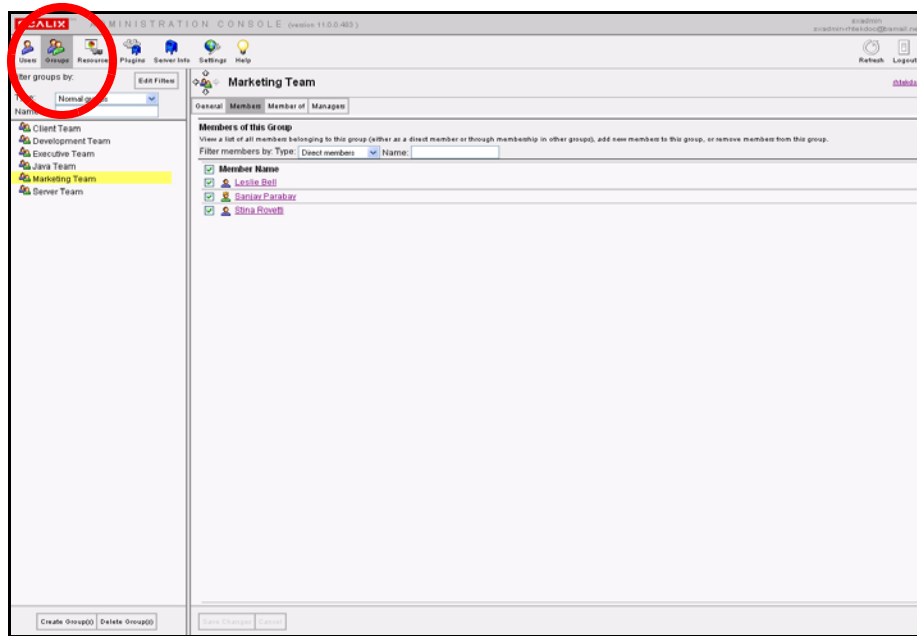
- “About Scalix Groups” on page 61
- “Creating New Groups” on page 62
- “Adding New Users to Existing Groups” on page 64
- “Modifying Groups” on page 65
- “Deleting Groups” on page 66
- “Assigning Group Managers” on page 66
- “Logging in as a Group Manager” on page 67

About Scalix Groups

In the Scalix system, the terms “Group” and “Public Distribution List” are used interchangeably. Groups can contain individual users, other groups, or a combination of both.

There are two levels of group membership:

- **Direct:** User is assigned directly to the group.
- **Effective:** User became a member because another group to which the user belongs was added to this group.



Creating New Groups

To create new groups:

- 1 In the Management Console toolbar, click the **Groups** button.
- 2 At the bottom of the contents pane, click **Create Group(s)**.
- 3 The first of three Create New Group wizard screens appears.

- 4 In the **Group Name** field, enter the new group name.
- 5 From the pull-down menu in the **Group Server Location**, select the appropriate server.

- 6 Click **Next**.
- 7 The second wizard screen is for adding members.

http://click.bamail.net:8080 - Create New Group - Mozilla Firefox

Create New Group

Members
Select the users and groups that should be a member of this group.

Filter groups by: Type: Name:

Fetching Group Members ...

< Back Next > Finish Save and Create Another Group Cancel

http://click.bamail.net:8080/sac/AdminAddGroupWizard.jsp#

- 8 Review the listed names and click the check boxes by the users you want to add to this group. Or click **Select All**.

Note Remember to use the Filter features if needed. For more on filtering, see “Filtering in the Management Console” on page 20.

- 9 Click **Next**.
- 10 The third and final wizard screen appears.

http://click.bamail.net:8080 - Create New Group - Mozilla Firefox

Create New Group

Group Membership
Select the groups that the group should be a member of.

Filter groups by: Type: Name:

☐ Musicians

< Back Next > Finish Save and Create Another Group Cancel

http://click.bamail.net:8080/sac/AdminAddGroupWizard.jsp#

- 11 If this new group is to be a “member” of another existing group, review the listed groups and click the check box by the group’s name.

- 12 Click **Finish**.

Adding New Users to Existing Groups

There are two ways to add new members to existing groups:

- Open each user's account records and add them (through the **Member of** tab) to a group.
- Open a specific group's records and add one or more members through the **Members** tab.

To add users through the Group feature:


- 1 In the Management Console toolbar, click the **Groups** button.
- 2 In the contents pane, select the group to which you want to add members.

Tip	Remember to use the filtering feature if necessary. If filtering by group name or type doesn't give the results you want, you can filter by mailnode: In the Management Console toolbar, click the Server Info button. In the contents pane, a tree displays all servers and their respective mailnodes. Select the node whose groups you want to review. The node is depicted by a gray postal letter bin. When the node options appear, go to the display pane and click the Groups tab. For more on filtering, see "Filtering in the Management Console" on page 20.
------------	---

- 3 In the display pane to the right, click the **Members** tab.
- 4 Open the **Filter by: Type** menu and choose **All users**.
- 5 When the complete users list appears in this tab, scroll down and click the empty checkbox by each new member to add them to the group.
- 6 At the bottom of the display pane, click the now-active **Save Changes** button.

To add a user to a group through their individual account:

- 1 In the Management Console toolbar, click the **Users** button.
- 2 In the contents pane, select the user you want to add to a group.
- 3 In the display pane, click the **Member of** tab.



- 4 In the Filter Groups By Type field, open the pull-down menu and select **All groups**.
- 5 Click the empty check box by the group.
- 6 At the bottom of the display pane, click the now-active **Save Changes** button.

Modifying Groups

There are many modifications you can make to existing groups. They include changing a group's name and its address as well as giving it a second address.

Changing a Group Name

When changing a group's name, the system automatically updates all uses of it.

To change the name of a group:

- 1 In the Management Console toolbar, click the **Groups** button.
- 2 In the contents pane, select the group whose name you want to change.
- 3 In the display pane, select the **General** tab.
- 4 In the Email Addresses option, delete the text in the **Group Name** field and type a replacement.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Modifying a Group Address


To modify the email address used by the group:

- 1 In the Management Console toolbar, click the **Groups** button.
- 2 In the contents pane, select the group whose address you want to change.
- 3 In the display pane, click the **General** tab.
- 4 In the Email Addresses options, modify the alias and domain values.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Adding a Group Address

To add a second email address to be used by the group:

- 1 In the Management Console toolbar, click the **Groups** button.
- 2 In the contents pane, select the group whose address you want to change.
- 3 In the display pane, click the **General** tab.
- 4 On the far right, click **Add Address**.



- 5 A new data row appears below the current email address.
- 6 Enter an alias for the group. (This typically matches the Group Name.)
- 7 Enter a valid domain, if other than the default.

- 8 Use the up and down arrows to shuffle the new address within the group's address list. The address at the top of the list is the default address that the system uses.
- 9 At the bottom of the display pane, click the now-active **Save Changes** button.

Deleting Groups

You can easily delete groups from Scalix if you have appropriate admin permissions.

To delete an existing group:

- 1 In the Management Console toolbar, click the **Groups** button.
- 2 In the contents pane, select the group(s) you want to delete.
- 3 At the bottom on the contents pane, click the now-active **Delete Group(s)** button.
- 4 When the Confirmation dialog box appears, click **Yes**.

Assigning Group Managers

A Scalix administrator can delegate permissions to another user to manage one or more groups. Such users undertake a Group Manager role, which has limited scope.

Group managers can:

- Add or delete members from their group
- Modify group-specific information about members

Group managers must:

- Have a fully active Scalix account
- Be a member of the group they are managing.

Assigning a Group Manager

Before you assign a current user as the new Group Manager to one or more groups, make sure that groups exist in Scalix.

To make a user into a group manager:

- 1 In the Management Console toolbar, click the **User** button.
- 2 In the contents pane, select the user to whom you want to delegate group management permissions.
- 3 In the display pane, click the **Manager Of** tab. Use the filtering options if needed.
- 4 Click the check box by the group name.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.

Determining who is Group Manager

To determine who is group manager:

- 1 In the Management Console toolbar, click the **Groups** button.
- 2 In the contents pane, select the group for which you want to determine who is manager.
- 3 In the display pane, click the **Managers** tab.



- 4 A list of the users who are Group Manager(s) appears in the tab.

Logging in as a Group Manager

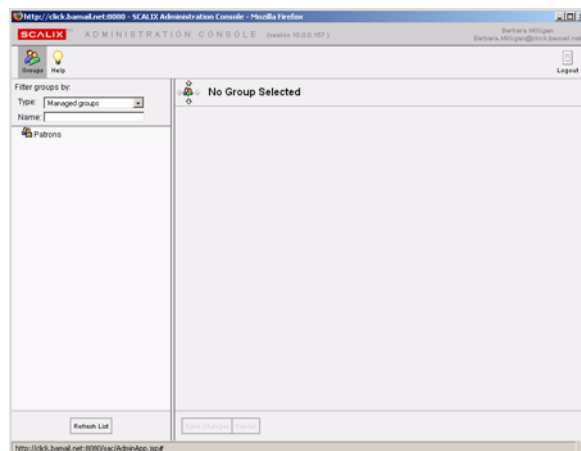
When a group manager logs in to the Management Console, he or she only has access to the Groups module. In this module, the user can view all users in the contents pane (read-only), but can perform limited administration tasks only the group(s) for which he or she is the manager.

If you are a group manager, here is how you can open and manage your group:

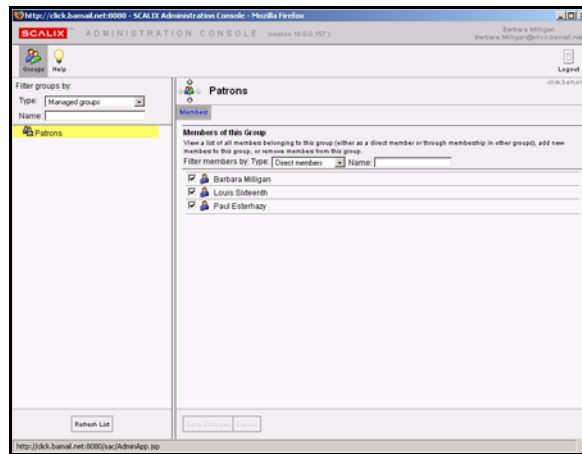
- 1 Log in to the Management Console with your usual Scalix authentication ID and password.

Note A reminder—the texts of your authentication ID (and password) are case-sensitive.

- 2 When the Management Console appears, all that is available is the **Groups** button (pre-clicked) and a list of the groups of which you are a manager.



- 3 Select the listed group, and the Management Console opens the group options in the display pane.



- 4 You can do the following:
 - Sort the members by direct or effective status.
 - Look for new member candidates and add them to the group.
 - Delete current members from the group.
- 5 When you have finished, click **Save Changes** at the bottom of the display pane.
- 6 Click **Log out** to disconnect your browser from Scalix.

Managing Resources

This chapter covers shared resources: How to create them, book them and manage them.

Contents

This chapter includes the following information:

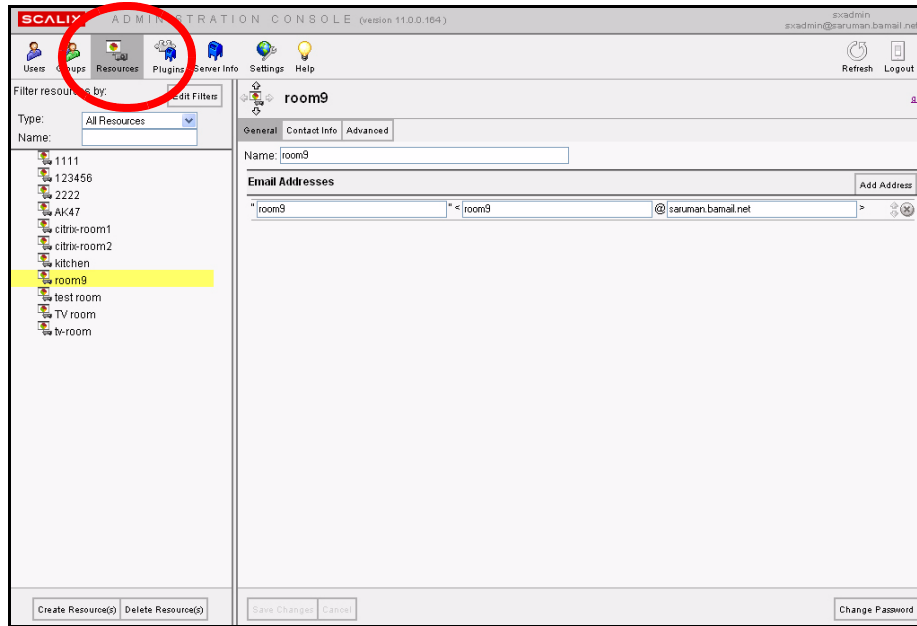
- “About Shared Resources” on page 69
- “Setting Up Shared Resources” on page 70
- “Booking Shared Resources” on page 71
- “Changing Shared Resource Settings” on page 71

About Shared Resources

Using the Scalix Management Console, you can manage shared resources such as conference rooms, company projectors, printers, other equipment, and more. SAC allows you to set up a shared resource so that it can be booked and tracked by users.

The administrator sets up the resource in SAC, then others can reserve or view those reservations in their client, whether that be Outlook or SWA.

Only Premium users can book resources. Standard users cannot.



Setting Up Shared Resources

Any resources to be booked by users must first be set up in the Scalix Management Console.

To set up a shared resource:

- 1 In the SAC toolbar, select **Resources**.
- 2 In the lower left corner, click **Create Resources**.
- 3 The new Create Resource Wizard launches. Fill in the fields. Most are self explanatory. A few that are not include:

- **Mailnode:** From the drop-down menu, select a mailnode through which the resource's reservations should route. For more on mailnodes, see "Managing Mailnodes" on page 40.

Create as premium resource: If this box is NOT checked, every premium user has full access to the resource.

When this box IS checked, select users can log in to the resource account from SWA and Outlook. This enables you to set fine-grained access control on the resource calendar, determining who may book a resource and who cannot. Access is set in the client as you would any folder.

This also enables delegates to open the resource calendar if given that level of access (also set in the client as a folder property), and it enables blocking direct resource booking if you prefer to designate one central person (such as an office admin) to manually approve every booking.

- **Can have recurring events:** When this box is checked, the resource can be reserved in a regular, recurring time slot such as every Friday at 3 p.m.

- **Can book concurrent events:** When this box is checked, the resource can be booked by multiple people for the same time.

- 4 **Click Next or Finish or Save and Create Another Resource.**
- 5 If you clicked **Next**, you see the Contact Information Screen. Here, you can record information about the shared resource such as the telephone number for a conference room or the location of a shared projector. Fill in the fields as desired, then click **Finish** or **Save and Create Another Resource**.

Booking Shared Resources

Shared resources are booked by end users in their client applications, which include Outlook and SWA. Other clients such as Thunderbird and Evolution do not have this feature. Remember that only Premium users can book resources.

To book shared resources in Outlook or SWA:

- 1 Go to your client application.
- 2 In the Folders List, click **Calendar**.
- 3 Reserve the resource in the same way you would add an attendee to a meeting. For more on how to add attendees, see Using the Calendar in the online help system.
- 4 When finish, click **Send**.

Changing Shared Resource Settings

Using the Scalix Management Console, you can modify or delete shared resources.

Modifying Shared Resources

If needed, you can change the properties for any shared resource. That includes its name, email address, location, availability to Standard versus Premium users and more.

To modify a shared resource:

- 1 In the Management Console toolbar, click **Resources**.
- 2 In the contents pane, select the resource you want to modify.
- 3 Its properties appear in the display pane.
- 4 Click through the three tabs - **General**, **Contact Info** and **Advanced** - to display the fields you want to change.
- 5 Type in the new information as needed.
- 6 When finished, click

Changing Passwords

Using the Scalix Management Console, you can change a shared resource's password.

To change a shared resource's password:

- 1 In the Management Console toolbar, click **Resources**.
- 2 In the contents pane, select the resource for which you want to change the password.
- 3 In the lower right-hand corner, click **Change Password**.
- 4 The **Change Password** dialog box appears.
- 5 Type in the new password and then confirm it in the next field down. If desired, click **Resource must change password on first login**.
- 6 When finished, click **Change Password**.

Deleting Shared Resources

If needed, you can delete a shared resource from the system. Once deleted from the Scalix Management Console, it no longer appears in users' clients.

To delete a shared resource:

- 1 In the Management Console toolbar, click **Resources**.
- 2 In the contents pane, select the resource you want to delete.
- 3 In the lower left-hand corner, click **Delete Resource**.

Managing Server Processes

This chapter covers the management of basic several server functions, including mailnodes, services, daemons and queues.

Contents

This chapter includes the following information:

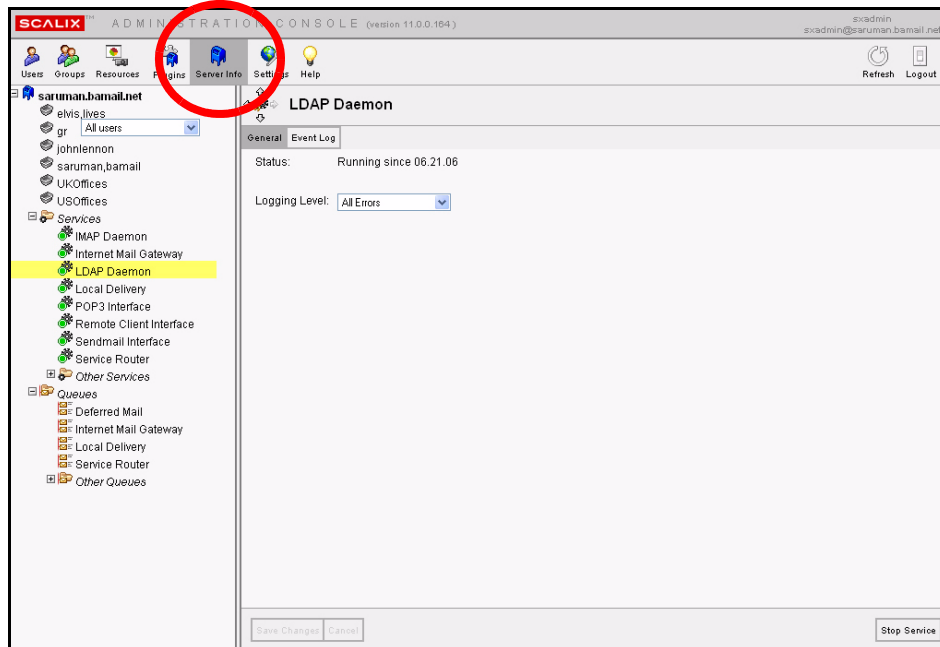
- “Introducing Server Info” on page 73
- “Stopping and Starting Services” on page 75
- “Monitoring Scalix Services and Daemons” on page 76
- “Monitoring the Active Users” on page 77
- “Monitoring the Message Store” on page 77
- “Monitoring Message Queues” on page 78
- “Review the Installation Summary” on page 79

Introducing Server Info

The Server Info feature allows you to manage server processes to ensure smooth and proper function of the Scalix system. The elements you can manage from the Management Console via the Server Info screens are:


- **Mailnodes:** Organizational aids in Scalix systems that you can create to help you organize your user base. Your Scalix server has one mailnode by default, but you can add more and then sort your users and groups into the various nodes as you prefer.
- **Services:** Processes that you can start and stop while the Scalix Server is operating. The contents pane displays the most commonly-used services on top. To see the less frequently-used services, click the plus sign to expand the tree.
- **Daemons:** Processes that Scalix starts automatically when you start the Scalix Server and they must operate continuously while the Scalix Server is running. (You’ll find all daemons listed under “Services”.)
- **Queues:** Pass messages and requests to Scalix services involved in message processing. Some services have associated queues, others do not.

There are a number of server-specific tasks you can perform with the assistance of the Server Info features in the Management Console.

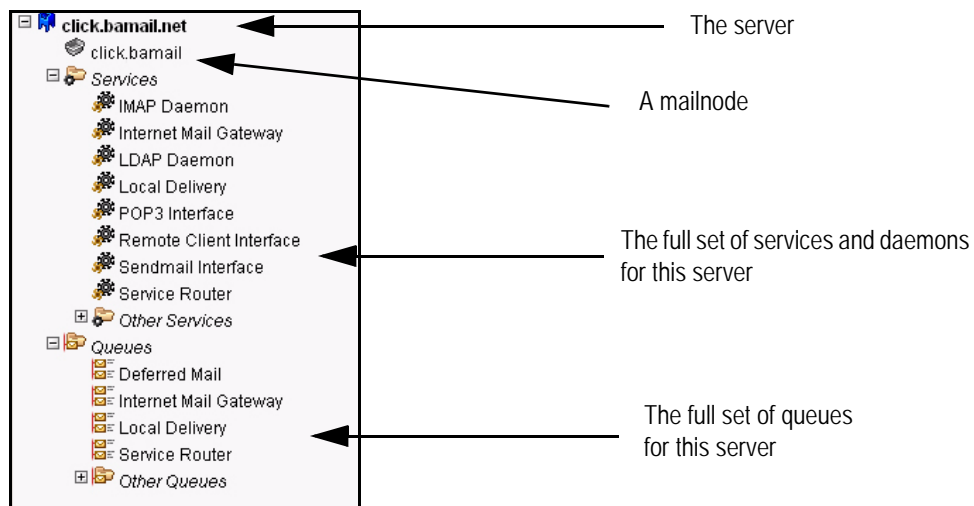


To view server information contents:

- 1 In the Management Console toolbar, click **Server Info**.



- 2 A hierarchy of server-specific information appears in the contents pane.



This includes all related mailnodes, services, daemons and queues.

- 3 To view the status of a server component (service, queue, etc.), click the item in the contents pane and look at the related options in the display pane to the right.

Stopping and Starting Services

You can use the Management Console to stop and/or start any of the services (daemons) or queues. The Management Console enables you to do this with individual resources or to globally stop/restart all the resources.

Note that services are sorted into two lists—*Key* and *Other*. Both comprise all of the Scalix services, but the most important services are listed as “Key” for your convenience.

To stop and start services:

- 1 In the Management Console toolbar, click **Server Info**.
- 2 In the contents pane, select the server whose services you want to stop or start.
- 3 In the display pane, click the **Services** tab.
- 4 This tab lists only the key services by default, but you can browse all other Scalix services at this time by clicking the button by **Display all services**.

Note that each service has either a green light or red light, to indicate its active status.

- 5 After locating the specific service entry, do one of the following:
 - Click **Stop** (A Stop button is visible if a service is running.)
 - Click **Start** (A Start button is visible if a service has been stopped.)

<input checked="" type="radio"/> Display key services <input type="radio"/> Display all services		
Stop	 IMAP Daemon	Running since 12:15:33
Stop	 Internet Mail Gateway	Running since 12:15:34 with 0 messages
Stop	 LDAP Daemon	Running since 12:15:33
Stop	 Local Delivery	Running since 12:15:34 with 0 messages
Start	 POP3 Interface	Stopped at 15:19:36
Stop	 Remote Client Interface	Running since 12:15:34 with 0 users
Stop	 Sendmail Interface	Running since 12:15:34 with 0 messages
Stop	 Service Router	Running since 12:15:34 with 0 messages

Alternative Stop/Start Procedure

Another way to stop or start services:

- 1 In the Management Console toolbar, click **Server Info**.
- 2 In the contents pane, select the server whose services you want to stop or start.
- 3 In the display pane, click the **Services** tab.
- 4 Click the **Stop Service** button in the lower right corner of the General tab.

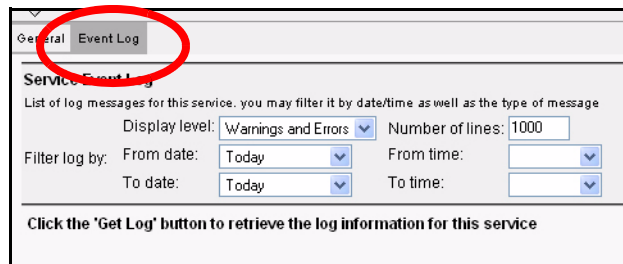
- 5 When the service has been stopped, a status update message appears in the **General** tab, reporting that the service was stopped, and noting the current date and time.
- 6 To restart the service, click the now-active **Start Service** button.
- 7 An updated **Status** message appears in the General tab, noting that this service has been running since (current date and time).

Monitoring Scalix Services and Daemons

The operational status of all Scalix system daemons and services can be individually monitored with the Management Console. After using the Management Console to customize event logging levels on a per-service basis, you can view varying increments of system activity.

To monitor services and daemons on Scalix:

- 1 In the Management Console toolbar, click **Server Info**.
- 2 In the contents pane, select the service you want to stop or start. Remember that you can expand the services tree to view the *Other Services* sub-branch.
- 3 In the display pane, click the **General** tab.
- 4 Open the **Logging Level** menu and select the level you want for monitoring from the Management Console.
- 5 At the bottom of the display pane, click the now-active **Save Changes** button.
- 6 Click the **Event Log** tab to view its contents.



- 7 Use the **Service Event Log** options (menus) to customize your log settings:
 - Open the **Display Level** menu and choose the type of activity you want to monitor with the Management Console.
 - Adjust the **Number of Lines** you want to review in the Management Console.
 - Open the **Date/Time** menus to focus on the interval you want to review.
- 8 Click the **Get Log** button.
- 9 The Management Console polls Scalix and displays the requested data (if any) in the area below the Service Event Log options.
- 10 You can refine your settings and re-poll the server as needed.

Note

The ability to save the event log or any customized samples as data files is limited to the Scalix CLI.

Monitoring the Active Users

You can use the Management Console to monitor the currently connected client users and assess the processes generated by their connection.

To monitor services and daemons on Scalix:

- 1 In the Management Console toolbar, click **Server Info**.
- 2 In the contents pane, select the server on which the users you want to monitor are located.
- 3 In the display pane, click the **Active Users** tab.
This tab lists all connected users who are now logged in to Scalix. The following summary is provided:
 - Each user's name
 - The kind of connection they've made (IMAP, MAPI or POP)
 - The ID assigned to each user-generated process
- 4 You can refresh the list manually or start an automatic refresh/update process.
 - Click the checkbox by **Enable Automatic Refresh**. The default refresh rate is once a minute.
 - Click the **Refresh** button—whenever you want to manually update this tab's contents.

Monitoring the Message Store

You can use the Management Console to obtain a high-level overview of activity in the Scalix server message store. The Management Console, however, does not allow you to perform any direct management of the Message Store in Scalix.

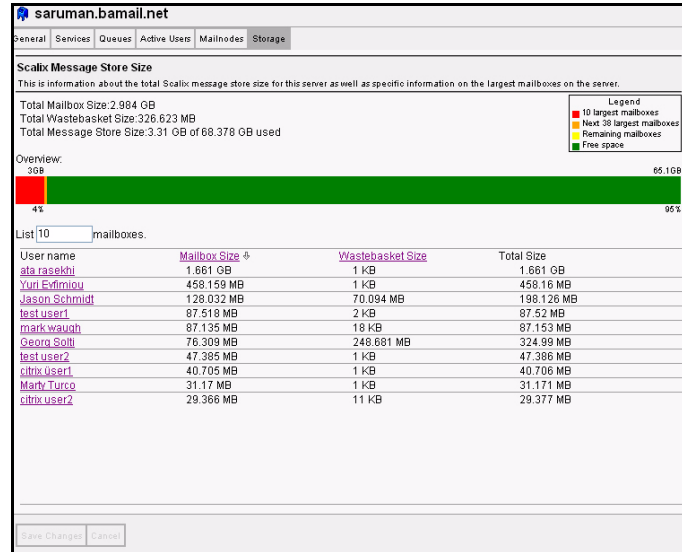
To monitor services and daemons on Scalix:

- 1 In the Management Console toolbar, click **Server Info**.
- 2 In the contents pane, select the server whose message store you want to monitor.
- 3 In the display pane, click the **Storage** tab.

The Storage tab summarizes these capacity statistics:

- **Total mailbox size:** The cumulative total of all current mailboxes, including folder items, inbox items and calendar items.
- **Total wastebasket size:** The cumulative total of items in the Trash that have not been expressly deleted.
- **Total Message Store size:** The actual capacity of this server's Message Store.

A broad colored bar shows (at a glance) the total capacity (in one solid color) and the used portion (in a contrasting color).



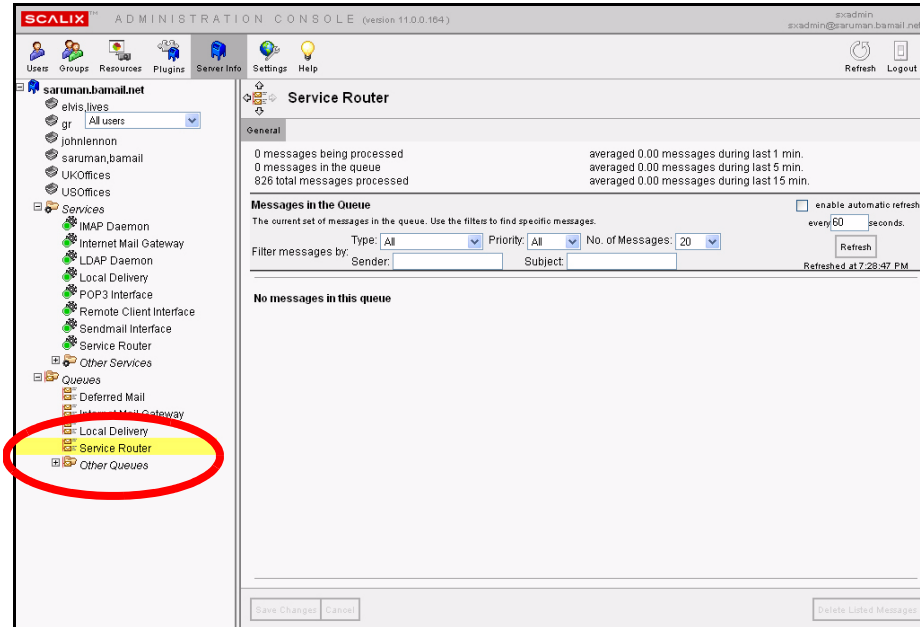
- 4 You can also browse the most active (highest-capacity) mailboxes, and see (in kilobytes) how much space each mailbox takes (including undeleted trash).
 - The default is “10” mailboxes, but you can increase this number to your preference.

Monitoring Message Queues

Using the Management Console, you can individually monitor each message queue. If needed, you can use the Management Console to purge messages from a queue.

To monitor message queues:

- 1 In the Management Console toolbar, click **Server Info**.
- 2 In the contents pane, select the queue you want to monitor. If the queue you want to monitor is not visible, expand the tree to view the *Other Queues* sub-branch.
- 3 The queue’s information appears in the display pane, including the most recent additions to this queue.
- 4 The **General** tab (which appears by default) summarizes the following information about this queue:
 - The number of messages processed in the last 60 seconds
 - The number of messages actually in the queue (over a five minute period)
 - The total number of messages processed in the queue in the last sixty minutes.



- 5 You can customize the message list settings:
 - Use the menus to set the **Type**, **Priority** and **Number** of messages listed below.
 - Filter the messages by entering the text of a particular subject or sender.
 - If you choose, you can permanently store these settings for later reuse by clicking **Save Changes** (at the bottom of this tab).
- 6 Click the **Refresh** button.
- 7 The message list updates according to your specifications.
- 8 If needed, you can delete all listed messages by clicking **Delete all listed messages** in the lower right corner of this tab.

Review the Installation Summary

You can use the Management Console to review installation-specific information about your Scalix server and its primary components. This is useful when you are troubleshooting problems.

To review the installation summary:

- 1 In the Management Console toolbar, click **Server Info**.
- 2 In the contents pane, select the server whose installation you want to review.
- 3 In the display pane, click the **General** tab.
- 4 Review the following summaries of the Scalix installation:
 - The version number of each Scalix component.
 - The release number
 - The date this component was installed (or upgraded)

Using Management Plugins

This chapter covers Management Plugins: What they are and how to use them.

Contents

This chapter includes the following information:

- “About Management Plugins” on page 80
- “Running Management Plugins” on page 81
- “Writing Management Plugins” on page 81
- “Execution Environment” on page 83
- “Output Format” on page 83
- “Language Considerations” on page 84
- “Deployment Script” on page 85
- “Sample and Template Scripts” on page 86

About Management Plugins

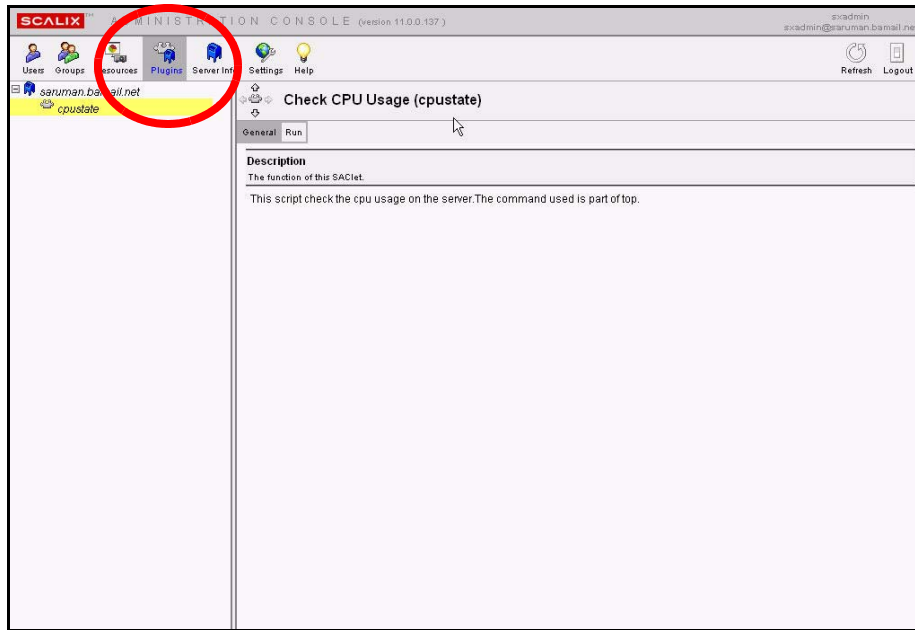
Management Plugins allow you to save and run from the Management Console many frequently-done tasks such as checking CPU or disk usage, scanning logs, listing public folders, checking the message store and more. They extend the functionality of the Management Console and provide more ease of use by launching simple, one-way scripts from the GUI.

For security reasons, only Scalix Admins, also known as full administrators, can run these plugins and even then, those administrators must have special plugin permissions.

There are several other restrictions on Management Plugins:

- Plugins cannot download files or executables from the Internet
- Plugins cannot remove or delete anything from the Scalix system’s file system
- Plugins cannot remove or delete bulk users and groups
- Plugins cannot consume large amounts of system resources, run repeatedly or for long periods of time

Scalix provides several template plugins in the Administration Resource Kit and encourages system administrators to write their own.



Running Management Plugins

As with other functions in the Scalix Management Console, the contents pane lists all available management plugins in a hierarchy, with individual plugins appearing underneath the servers on which they run. There also is an **All Servers** node with plugins that have global impacts.

To run management plugins:

- 1 In the Management Console toolbar, click **Plugins**.
- 2 In the contents pane, select the plugin you want to run.
- 3 The **General** tab of the display pane explains the plugin's function.
- 4 To run the plugin, click the **Run** tab.
- 5 Fill in the fields as needed. Those fields vary depending on the script being used.
- 6 If results are called for, they appear in a pop-up window titled, **Plugin Results**.
- 7 To refresh the results, click **Enable automatic refresh** in the upper right-hand corner, fill in the time interval desired, then click **Refresh**.

Writing Management Plugins

The goal of the framework created by the Management Plugins is to facilitate a simple mechanism to run custom scripts. You can write your own plugin, but they should be simple and low in resource usage. This framework is not intended for complicated programs that consume large amounts of host resources.

When writing your own management plugins, there are several other things to keep in mind:

- **Exit Status:** Each plugin developed must exit with non-zero status upon failure. The exit status should not be 1, as failed shell scripts exit with 1.
- **Error Message:** In case of problems, a comprehensible error message should be returned on a single line. The error must be written to stderr. For example a failed plugin error message may look like this: script_name: Failed to locate the user, <user name>, on the server directory on host <host name>.
- **Naming Conventions:** Management plugins should not have script or language extensions such as ".py" or ".sh", or ".pl". In addition, plugin names should not have any shell interpretable characters such as "*", "!", "\$", ":", "(", ")", "~" or other unconventional characters such as "()", "[]", "{}". Underscore and dashes are okay, but preferably the script or plugin should be a single word, less than 20 characters. For example, check-queues, monitordisk, and docollocate_entries.
- **Usage or Help Information:** All management plugins must have a -- help option. This help option must provide the following information about the script.
 - **Version:** This is a place holder for backward compatibility of parameter parsing.
 - **Friendly Name:** Each plugin should have a unique name, known as the "friendly name." For example, a checktemperature script can have a friendly name of "Check Temperature", and the output should look like:

NAME:

Check Temperature

This must be a single line of output with token NAME:, and on the next line, followed by the name

- **Description:** This can be a multi-line output that describes the script's function. Its format should be as follows:

DESCRIPTION:

This script checks the room temperature usage for the SAN appliance in the remote data center at the Exodus Facilities in Santa Clara.

- **Parameters List:** This final line should have the token `PARAMETERS:` and a six-column, tab-separated output on the next line, describing the attributes of each argument or option. On each line of output, each column must have the following attributes:

column 1: option flags

column 2: symbolic or human readable name for the flag

column 3: Default value if any for the flag

column 4: Type of value: number, boolean, or string

column 5: Type of Argument: single, multi, or none.

column 6: Short Description of the flag or option.

For example, these values, as per the above specification, must be printed or echoed by the script onto stdout, and must have the six columns in the following tab separated manner:

PARAMETERS:

```
-t temperature 78 number single "Temperature of the Data Center"
-- help Help "" "" "" "Usage Information for checktmp"
```

- **Output Type:** This token represents the type of output generated by the plugin. The two expected types are text/html or text/plain with text/plain being the default.

Altogether the output of the above script when executed with "checktmp --help" is shown below. New tokens begin in the first column and are separated from the previous set by a new line.

VERSION:

1.0

NAME:

Check Temperature.

DESCRIPTION:

This script checks the room temperature of the data center in Virginia Exodus collocated servers.

PARAMETERS:

```
-t temperature 78 int single "Temperature of the Data Center"
-- help Help "" "" "" "Usage Information for checktmp"
```

OUTPUT-TYPE:

text/html

Execution Environment

The choice of language for plugin development is up to the administrator, as long it has an execution or run time environment on the production machine. If using Perl, Python, or Java, the production server (or the target host) must have the appropriate runtime environment.

Note that all management plugins are forked and executed by the remote execution service, so can create further create processes. All resources consumed during the execution of these plugins must be reaped, especially child processes created, or threads created or sockets and temporary files used. Where possible, judicious usage of threads must be observed for potential deadlock issues.

Management plugins must be owned by the root super user, and must have Unix 700 permission set. Avoid using any setuid programs.

Output Format

The output format or the results from the run of a plugin must each be on a separate line. The output displays on the console, as is, from the run.

Deploying Management Plugins

Once a plug-in is written according to the guidelines listed above, you can deploy it.

All plugins must be deployed by a system administrator with root access on the target or production server. That is, the Management Console does not provide an upload facility for security reasons.

The scripts and all associated files must be deployed onto the directory `~/plugin/`. This directory is instance specific.

A default deployment script (`sxcfgplugin.py`) is provided for the system administrator to bootstrap or deploy plugins and initialize respective ACL items for specified user(s). The user authorized to execute the script must exist (or its directory entry) on the target host where the script is deployed.

All deployed scripts must have 700 Unix permission and the owner and group must be root.

All plugins should be reviewed by the site security administrator to ensure that no malicious code is embedded.

To deploy a Plug-In:

- 1 Copy the plugin to the deployment directory


```
cp <plugin_name> ~/plugin/
```
- 2 To create an ACL for a resource-type request:


```
omaddacl -t plugin -l <plugin_name>
```
- 3 Provide execute permission for the specified user, after which only full administrators with 'execute' permissions have access to the plugins within the console.


```
omaddacl -t plugin -l <plugin_name> -n ORN_user_admin -c "execute"
```

or

```
omaddacl -t plugin -l <plugin_name> -a ORN_PDL -c "execute"
```

Language Considerations

Python

If you deploy a python script, remove any `.py` extension, and preferably compile it. The first line of any python script must have `#!/usr/bin/python`, which will negate the need to precede its invocation from within RES with the prefix python runtime environment.

Perl

Similar guidelines as above should be observed for Perl, short of compilation.

Shell

All `.sh` or `.csh`, or `.ksh` extensions or prefixes should be stripped, and respective `#!/path_to_shell` should be the first line of the code.

C/C++

All C/C++ programs deployed as management plugins must be compiled.

Java

All Java programs deployed as management plugins must be compiled and a shell wrapper provided around it to invoke the Java runtime environment with its desired java options. The remote execution server will only invoke the shell wrapper, supplying it with the received argument list from the Ubermanager, via the console.

Deployment Script

To deploy plugins, use the deployment script, `sxcfgplugin.py`. It has several options:

```
-- add [options]
-- delete [options]
-- list [options]
-- deploy [options]
-- undeploy [options]
-- help
```

Using the 'add' subcommand to allow new users to be added to the plugin ACL:

```
-- add -l { <plugin_name> | all } -u user_authid -i
{<instance_name> | all }
```

Using Delete to remove users :

```
-- delete -l {<plugin_name> | all } - user_authid -i
{<instance_name> | all }
```

Listing all the plugins deployed for an 'instance' or listing all plugins for which the -u user_authid has execute access:

```
-- list -i {<instance_name> | all } [-u authid]
```

Deploy this plugin:

```
-- deploy { -D <source_directory> | -l <plugin_path_name> } { -i
<instance_name> | all } [-u user_authid]
```

If -D is specified, take all plugins under that directory and deploy them, or -l plugin_path_name will deploy only a single plugin for all instances or a single specified instance.

The -u user_authid option will create the appropriate ACLs as well. This is the equivalent of doing an -- add operation, after -- deploy.

Undeploying the plugin from the specified instance or all instances :

```
-- undeploy -l <plugin_name> -i {<instance_name> | all }
```

This involves removing all the ACLs associated with the ACL and removing the file.

Sample and Template Scripts

To see some sample and template scripts, go to http://www.scalix.com/wiki/index.php?title=Administration_Plugins.

Advanced Administration Tasks

About This Section

The remaining chapters in this guide involve more advanced administration tasks that are done only on the command line. These include such procedures as backup and recovery, managing public folders and troubleshooting.

This Section's Contents Include:

Included in this section are the following topics:

- "Running Backups and Recovery" on page 88
- "Managing Public Folders" on page 101
- "Access Control Lists" on page 114
- "Creating a Redirect Account" on page 128
- "Changing Hostnames and IP Addresses" on page 130
- "Recovering Deleted Items" on page 132
- "Setting Message Delivery Rules on the Router" on page 136
- "Working with SIS" on page 148
- "Configuration Options" on page 151
- "Scalix Command Line Reference Guide" on page 230

Running Backups and Recovery

This chapter covers backup and recovery strategies and procedures to protect your data and system configurations from possible loss or corruption.

Contents

This chapter includes the following information:

- “Concepts and Strategies” on page 88
- “Full Backup and Restore” on page 89
- “Disaster Recovery” on page 98
- “Export/Import Backup” on page 99
- “Single User Recovery” on page 100

Concepts and Strategies

As with any messaging system, you should back up your Scalix data on a regular basis - daily at the very least. Backups should encompass all contents of the instance’s home directory (normally `var/opt/scalix` in a single-server or typical installation scenario), including all sub-directories and their files. That way, they include not only email messages and their folders, but also calendar items, public folders, directories and system configurations. You should back up all servers, and you can back up to tape, to the same machine, or to a different one.

Because the Scalix system is a series of flat Linux folders, the backup procedure can take the form of a simple snapshot or copy procedure. But because many of those files are inter-related and dynamically reference one another, files and pointers (reference information) are continually created and deleted during messaging transactions, making for a constantly changing environment. As a result, capturing a full and accurate snapshot requires temporarily suspending the system to ensure a complete and consistent copy.

There are several different methods for backup and recovery, each with its own strengths and limitations. No single solution meets all needs and some combination of all may be the best approach for your organization. The options include:

- Full Backup and Restore: Most comprehensive with smallest size, but can be cumbersome for restoring a single user’s data and cannot restore individual items
- Disaster Recovery: Restores (and replaces) everything on the system, including messages, calendar items, contacts, routings, system settings and more

- Mailbox Export and Import: Best for restoring a single user's data or individual items, but results in a larger backup

Some possible strategies to consider:

- Do regular full backups for safety
- Do occasional full backups with regular incremental backups. For example: Back up the full system weekly, but every night, record only the changes (the increment) since the last full backup
- Do regular full backups for the system as a whole and complement that with export/import backups for key executives to enable single-user restore in the most important cases
- Do regular full backups and infrequent export/import backups on the system as a whole

Some other decisions to make before beginning a backup procedure:

- Do you want to back up to tape, to a different partition on the same machine, or to a different server?
- Do you want to stop the system entirely before taking a snapshot or temporarily suspend write activity?
- Do you want to do a full backup each time, or record only the increment from a baseline?

Full Backup and Restore

In the full backup and restore scenario, it is best to use LVM to take a snapshot of the entire folder structure, including the message store, public folders, directories and system synchronizations on all servers and store the data on tape in a safe location. For full backup, there are two options:

- Back up the entire system every time: The most comprehensive and failsafe method but takes longer and results in larger file size.
- Back up the entire system the first time to establish a baseline, then use the synchronization command to record only the changes that have occurred since that baseline: This is less comprehensive, but faster and results in smaller backup size.

If and when you need to restore, you have several choices:

- Full Recovery: Restore the entire system to the original servers, including all folders.
- Partial Recovery: Restore the entire system to a secondary server (or servers) then extract the part(s) you need and copy them in to the live system.

After deciding on the basic approach to backup and restoration of data, there are choices about how to stop the system long enough to get an accurate snapshot. Because the Scalix database consists of files representing many different components that dynamically reference each other (such as mailboxes, folders, messages and attachments), files and pointers (reference information) are continually created and deleted during messaging transactions, making for a constantly changing environment. Therefore, any backup actions taken while the system is active can result in an incomplete and inconsistent copy.

To solve this problem, there are two options when doing a full backup:

- Shut down Scalix (omshut), and then copy, tar, tar-gz or rsync the contents of the instance's home directory to another location, then restart (omrc).
- Temporarily suspend write activity to Scalix (omsuspend), create a snapshot of the instance home directory, release the suspension, and then copy, tar, tar-gz or rsync the contents of the snapshot to another location. A five-second suspension of activity normally suffices.

Best Practices for Full Backups

- Some best practices for backups:
- Take advantage of snapshot capabilities. Logical Volume Manager (LVM), provided with Linux, provides snapshot functionality.
- After an initial backup, you can use an rsync procedure which backs up only the changes.
- When using the rsync method, use another Linux host if possible. This provides a redundant spare if needed, a message recovery server and multiple daily copies.
- Where applicable, use compression in the form of a tar -zxvf command.
- Always back up all of the instance home directory, including all subdirectories and their files.
- When backing up to a different server, the permissions must be the same on both servers.
- When using rsync, use the -H switch to ensure that hard links are retained. For example:

```
rsync -azvH
```

where a=archive, z=compress data during transport, v=verbose and H=retain hardlinks. If hard links are lost, the size of the message store can grow significantly.


Using LVM to do Full Backups of Scalix

If you intend to use LVM to make a snapshot of the system for backup purposes, all Scalix servers must be configured with specific mount points. If you did not do that when installing the server, do it now.

The following diagram displays one potential configuration on RedHat. In this configuration, there were 30 GB of storage available at installation and 20 are now allocated for LVM. However, the mount points are not established until further configuration is done.

Note

Actual configuration of production servers typically involve more disks and overall storage.



To configure a Scalix server with the proper mount points for LVM:

- Do an initial look at the disk configuration using the fdisk command.


```
[root@showtime root]# fdisk -l
```

Disk /dev/hdc: 30.0 GB, 30005821440 bytes
255 heads, 63 sectors/track, 3648 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/hdc1	*	1	16	128488+	83	Linux
/dev/hdc2		17	147	1052257+	82	Linux swap
/dev/hdc3		148	2696	20474842+	8e	Linux LVM
/dev/hdc4		2697	3648	7646940	f	Win95 Ext'd (LBA)
/dev/hdc5		2697	3345	5213061	83	Linux
- Execute the command vgscan. This scans all disks for volume groups and builds the files /etc/lvmtab and /etc/lvmtab.d/*. These files are the databases for all other LVM commands.


```
[root@showtime root]# vgscan
```

vgscan--reading all physical volumes (this may take a while...)
vgscan--"/etc/lvmtab" and "/etc/lvmtab.d" successfully created
vgscan--WARNING: This program does not do a VGDA backup of your volume group
- Execute the command pvcreate, which initializes a disk or partition for use by LVM.


```
[root@showtime root]#
```

pvcreate /dev/hdc3 pvcreate -- physical volume "/dev/hdc3" successfully created

- 4 Create a volume group with the command `vgcreate`.

```
[root@showtime root]#
vgcreate vgscalix /dev/hdc3 vgcreate -- INFO: using default physical extent size 4 MB
vgcreate -- INFO: maximum logical volume size is 255.99 Gigabyte
vgcreate -- doing automatic backup of volume group "vgscalix"
vgcreate -- volume group "vgscalix" successfully created and activated
```

- 5 If it isn't already there, create the subdirectory `Scalix` under the directory `/var/opt`. During the installation of Scalix, it will correctly recognize this directory, where it will then install and create all appropriate files and subdirectories.

```
[root@showtime opt]# mkdir scalix
```

- 6 Execute the command `lvcreate` to create a logical volume in an existing volume group.

Note that from our original 20 GB, we have now created a 10 GB logical volume `lvscalix` in the `vgscalix` volume group.

```
[root@showtime root]# lvcreate -L10000M -nlvscalix vgscalix
lvcreate -- doing automatic backup of "vgscalix"
lvcreate -- logical volume "/dev/vgscalix/lvscalix" successfully created
```

- 7 Run the `mkfs` command. This creates a new file system on a specified device and initializes the volume label, file system label, and startup block.

```
[root@showtime root]# mkfs -t ext3 /dev/vgscalix/lvscalix
mke2fs 1.32 (09-Nov-2002)
Filesystem label =
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
1281696 inodes, 2560000 blocks
128000 blocks (5.00%) reserved for the super user
First data block=0
79 block groups
32768 blocks per group, 32768 fragments per group
16224 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632
Writing inode tables: done
```

Creating journal (8192 blocks): done

Writing superblocks and filesystem accounting information: done

- 8 This filesystem is automatically checked every 28 mounts or 180 days, whichever comes first. Use tune2fs -c or -i to override that.

```
[root@showtime root]# mount /dev/vgscalix/lvscalix ~/s
```

```
[root@showtime root]# vi /etc/fstab (add the line at the bottom)
```

```

LABEL=/          /          ext3      default ts      1 1
LABEL=/boot      /boot      ext2      default ts      1 2
none            /dev/pts   devpts    gi d=5,mode=620 0 0
none            /proc      proc      default ts      0 0
none            /dev/shm   tmpfs     default ts      0 0
/dev/hdc2        swap       swap      default ts      0 0
/dev/cdrom /mnt/cdrom    udf,iso9660 noauto,owner,kudzu,ro 0 0 /
dev/vgscalix/lvscalix ~/s ext3      default ts      1 3
[root@showtime root]# mount ~/s
```

Performing the Full Backup

Every backup is different. The specifics vary by the server setup, the method, the size of the data files and more. So no single script covers all scenarios. However, your Scalix software distribution contains a directory titled "admin_resource_kit", which contains template backup scripts that can help you get started. These scripts also are included in a larger tar/gzip file, which is typically prefixed with "sxbackup".

One example of a backup script is outlined below. This script takes daily snapshots of the data volume in the message store and then mounts it for backup. It can be invoked by a cron job or backup software.

An example of a script for daily message store backups is:

- 1 Setup paths and variables.


```

LVGRP=/dev/vgscalix
SXLV=lvscalix
BULV=sxbackup
LVSIZE=20G
MNTPT=/mnt/sxbackup
LOGRCPT=
DEVICE=/backup

SXBIN=/opt/scalix/bin
LVBIN=/sbin
```

```

BIN=/usr/sbin
LOG=/tmp/sxbackup.log
DAY=`date +%F`

touch $LOG

if [ -z "$LVGRP" ];
then
    echo "The Logical Volume Group Variable (LVGRP) has not been set"
    exit
fi

if [ -z "$SXLV" ];
then
    echo "The Scalix Volume Variable (SXLV) has not been set"
    exit
fi

if [ -z "$BULV" ];
then
    echo "The Backup Volume Variable (BULV) has not been set"
    exit
fi

if [ -z "$LVSIZE" ];
then
    echo "The Backup Volume Size Variable (LVSIZE) has not been set"
    exit
fi

if [ -z "$MNTPT" ];
then
    echo "The Mount Point Variable (MNTPT) has not been set"

```

```

    exit
fi

if [ -z "$LOGRCPT" ];
then
    echo "The Log Recipient Variable (LOGRCPT) has not been set"
    exit
fi

if [ -z "$DEVICE" ];
then
    echo "The Tarball Directory Variable ($DEVICE) has not been set"
    exit
fi

#
#
sxback_begin ()
{
    date > $LOG

```

```

    echo "Taking snapshot & Mounting Scalix Data Volume" | tee -a $LOG

```

- 2 Verify the mount point exists. If it doesn't, create it.

```

    if ! test -d "$MNTPT"
    then
        mkdir $MNTPT >> $LOG 2>&1

```

- 3 Verify the Mount Point was created, if not exit

```

    if ! test -d "$MNTPT"
    then
        echo "!! Error creating the $MNTPT directory" | tee -a $LOG
        exit
    fi
fi

```

- 4 Suspend scalix and sync the disks.

```

echo "Suspending Writes to Scalix"|tee -a $LOG
$XBIN/omsuspend -s 10&
/bin/sync

echo "Creating the Logical Volume $LVGRP/$$XLV"|tee -a $LOG
# create the snapshot volume and mount it
$XBIN/lvcreate -L $LVSIZE -s -n $BULV $LVGRP/$$XLV >> $LOG
2>&1

echo "Enable Writes to Scalix"|tee -a $LOG
# release the suspend
$XBIN/omsuspend -r

echo "Mount $LVGRP/$BULV to the Mount Point $MNTPT"|tee -a $LOG
# mount the snapshot volume
mount $LVGRP/$BULV $MNTPT >> $LOG 2>&1

```

- 5 This next section uses tar. We recommended you use something more robust in case you need to do a single user restore. If you do that, comment this section out and have the backup software execute.

- 6 Do a "sxsnapshot -begin" before backing up and "sxsnapshot -end" upon completion of the backup.

- 7 If you want to back up directly to tape, set DEVICE to something like this:

```
#DEVICE=/dev/rmt0
```

- 8 Back up to tape device.

```
echo "tar to device $DEVICE. $DAY. tar the contents of
$MNTPT..."|tee -a $LOG
```

```
tar cf $DEVICE. $DAY. tar $MNTPT >> $LOG 2>&1
```

- 9 If the device is a file system, then compress it.

```
echo "gzip'ing the $DEVICE. $DAY. tar file..."|tee -a $LOG
```

```
gzip -f $DEVICE. $DAY. tar >> $LOG 2>&1
```

- 10 This ends the commented section. If you did use tar, resume from here.

```
}
```

```
sxback_end ()
```

```
{
```



```

#Unmount backup file system
echo |tee -a $LOG
echo ..... |tee
-a $LOG

date >> $LOG

echo "Umouting Scalix Backup Volume"|tee -a $LOG

umount $MNTPT >> $LOG 2>&1

echo "Removing the snapshot logical volume"|tee -a $LOG

$LVBIN/lvremove -f $LVGRP/$BULV >> $LOG 2>&1

```

- 11 Send the result to the Administrator. Configure a .forward file in the root home directory and forward all mail to the Error Manager. Identify the Error Manager with the "omshowenu" command.

```

echo "Mailing the report to $LOGRCPT"|tee -a $LOG

mail -s "Scalix Backup Results" $LOGRCPT < $LOG

}

usage ()
{
    echo "USAGE: $0 [ -begin | -end ]"
    exit 1
}

if [ $# -ne 1 ]; then usage; fi

case $1 in
    -begin)    sxback_begin ;;
    -end)      sxback_end   ;;
    *)        usage        ;;
esac

```

Backing Up Changes Only

After performing an initial baseline backup, you can use a variation on the synchronization command to back up only the changes. This is faster and results in a smaller backup file size.

The general syntax when using rsync is:

```
#rsync -options source target
```

Alert

Do not mix up source and target options, which can cause critical problems.

Alert

Deletions do not replicate with synchronization, so when using this method for backup, some deletions may not reflect in the backup tape. To make sure that they do, use the `--delete` command.

Using Synchronization to Back Up to the Same Server

You can use the `synchronization` command to back up only the changes in a same-server environment, then store that as an augmentation to the baseline full backup.

To use synchronization to record the changes on a single server:

- 1 Run the following command.

```
#rsync -avz --delete /var/opt/ /backup
```

This recursively copies all files from the directory `/var/opt` to the directory `/backup`. The files transfer in "archive" mode, which ensures that symbolic links, devices, attributes, permissions, ownerships etc. are preserved. In addition, compression is used to reduce the size of data portions.

In this example the `/backup` directory then contains a `/scalix` directory (with all subdirectories), and perhaps a `/jakarata-tomcat-5.0.2x` directory (with all subdirectories). This does mean the `/or /backup` partition must have as much space available as is stored in the contents of `/var/opt/`, which is typically rare. Run it initially, then run it again, notice the second time that only a few files copy. Write a message into a mailbox on Scalix, run it again, you'll notice now more files copy from `/user` and `/data` directories.

Using Synchronization to Back Up to a Different Server

You also can use the `synchronization` command to back up to a different server.

To use synchronization to record changes and store them to a different server:

- 1 Run the following command.

```
#rsync -avz --delete /var/opt/ backup.company.local : backup
```

This does the same as above, only it copies to the `/backup` directory on the host "backup.company.local". The backup.company.local host is nfs mounted from the Scalix server. On the backup machine you can then set a nightly cron job to build a day-of-week .tgz file to another area, which is backed up to tape weekly.

Disaster Recovery

Full disaster recovery brings back the entire system by building an entire new server. That includes all mailbox data, public folders, directories and configurations.

For more on disaster recovery, contact Professional Services.

Export/Import Backup

An alternative to the full backup procedure is to use the export/import method by which you can store and restore individual users' mailboxes much more efficiently. Using this method, you back up individual mailboxes or public folders - in their entirety - as separate files. That includes all attachments, referenced items, etc.

The advantages to this method are that the export file is easier to handle. You can copy it, move it around, import it, and more.

The disadvantage is that the sum file size of all individual mailbox backups can be as much as two times larger than the sum total of a full backup procedure.

The explanation is in efficiencies the Scalix system uses to store items. Consider the example of an attachment sent out to 10 people: In Scalix, the attachment does not replicate in all 10 mailboxes. Rather, all 10 mailboxes contain pointers to one single copy of the attachment, which is housed in the message store on the Scalix server. So when you export individual mailboxes and all of their contents, that attachment is repeated 10 times over - once with each box. That increases the size of the backup significantly. In large enterprise systems, this can be prohibitive.

TIP

The export/import procedure is also helpful for migrating users from one server to another.

Exporting for Backup

Exporting copies the contents of a user's mailbox or a series of public folders into one, single archive file. This can include email messages, calendar items, contacts, public folder contents and user settings.

Alert 2

For data consistency, be sure the user is logged out when running an export. If the user is not available to log out on his own, use the -force option to do it yourself.

To export a mailbox:

- 1 Make sure the user is signed off. If he/she is not, use the command -F <user name> to force them off.
- 2 On the server on which the user's mailbox is stored, run the export command, specifying the name of the user whose mailbox you want to back up.

```
sxmbosexp --u "User Name" -a /backup/<mail boxname>.mbox
```

Where <mailbox name> takes the form of first initial and last name (jsmith).

Importing from Backup

Once you've exported individual mailboxes or public folders, you can import them to restore data. This is especially helpful if one of your users deletes a large amount of data or experiences corruption that requires a reset to a previous point in time. It is also useful for migrating a complete mailbox.

To import a mailbox:

- 1 Make sure the user is signed off. If he/she is not, use the command -F <user name> to force them off.

- 2 On the server on which the user's mailbox is stored, run the import command, specifying the name of the archive file you want to import. It must be an archive file created through the export procedure outlined above or the command `omcpoutu`. By default, the data is restored to a mailbox with the same name as the original stored data, but with the `--user` option, you can target any existing mailbox.

```
sxmboximp -a /backup/<mailbox name>.mbox
```

Where `<mailbox name>` takes the form of first initial and last name (jsmith).

TIP

You can do selective imports. Some possibilities are: To import all folders with the exception of one, use the command `sxmboximp -a /backup/<mailbox name>.mbox --exclude F-<folder ID>`. (To get the ID number for the folder you don't want to import, use `sxmboxlist`.) To import two folders, use the command `sxmboximp --archivefile /backup/<mailbox name>.two --folder F-<folder ID> --folder F-<folder ID>`. To restore a single public folder, use the command `sxmboximp -a /backup/public.folders -f F-<folder ID>`.

Single User Recovery

If one of your users experiences corruption in some part of his or her message store, you can restore just that mailbox. There are two methods for single-user recovery:

- Full backup method
- Export/import method

Single-User Restore after a Full Backup

If you only have a full backup on record and want to restore a single user's mailbox, you can do so.

To restore a single user's mailbox:

- 1 Restore the full backup to a secondary server using whatever tool you used to archive the backup (such as `tar`).
- 2 Extract the specific mailbox data from the secondary server using the `sxmboxexp` command.
- 3 Copy it in to the live server using the `sxmboximp` command.

Single-User Restore After an Export Backup

If you have an export of the user's mailbox using either of the commands `sxmboxexp` or `omcpoutu`, you can do an import procedure to restore a single user's data. For more on the import procedure, see "Importing from Backup" on page 99.

Managing Public Folders

This chapter covers public folders: How to configure, access, synchronize and maintain them. It also outlines how to set permissions and assign them email addresses.

Contents

This chapter includes the following information:

- “Public Folder Overview” on page 101
- “Creating Public Folders” on page 102
- “Listing Public Folders” on page 102
- “Permissions for Public Folders” on page 103
- “Maintaining Public Folders” on page 105
- “Assigning Email Addresses to Public Folders” on page 106
- “Forwarding Public Folder Items” on page 111
- “Synchronizing Public Folders” on page 107

Note

Only Premium-level users of the Scalix Small Business Edition or Enterprise Edition can use public folders. For more information, see “About Scalix Product Editions”.

Public Folder Overview

Public folders are like electronic bulletin boards - they enable sharing of information such as documents, email messages, or calendars with other users. You can organize them by common interests, team projects, departments, or any other need. They also are useful for sharing meetings, appointments, or contacts.

Within the Scalix system, public folders are shared areas in the message store. Users can add items to them by cutting and pasting, dragging and dropping, sending mail messages to the Public Folder, and so on.

Public folders can contain the following items:

- Email messages
- Documents such as spreadsheets, text or word-processing files
- Calendars

- Contacts
- Other (“nested”) public folders

Public folders can be set up with permissions to restrict access to certain users or classes of user. In addition, expiry dates can be set so that short-lived information is automatically deleted after a specified period.

Note

Scalix commands and directories use the term “Bulletin Board” and the abbreviation “bb” to refer to public folders.

For more on the various public folder commands, see “Public Folder Commands” on page 112.

Creating Public Folders

You can create public folders from the client or the command line.

To create folders from the client, see the associated online help system.

To create folders on the command line:

- 1 To create a top-level folder, run the following command.

```
omaddbb -s <“Folder Name”>
```

Where an example of <“Folder Name”> would be `omaddbb -s “Engineering Offsite”`

- 2 Set the permissions level for this folder. For more on permissions, see the next section.
- 3 To create a subfolder under that one, use the following command. This would be a second-level folder.

```
omaddbb -m “TOP” -s <“Folder Name”>
```

- 4 To create another subfolder under the second level, use the following command. This would be a third-level folder. Note that the delimiter between the top level and the next level is the > angle bracket that appears after the word, “TOP.”

```
omaddbb -m “TOP><Folder Name>” -s <“Folder Name”>
```

Where the first folder name is the name of the parent and the second is the name of the new folder you are creating.

Note

When adding public folders, you can set such as expiration dates for the contents. For more on expiration dates, see “Maintaining Public Folders” on page 105 or the man page.

Listing Public Folders

You also can list public folders.

To list public folders:

- 1 Run one of the following command.

For top-level folders only: `omlistbbs`

For all levels of folders: `oml i stbbs -d 0`

Where the available options are:

Table 1: Options for the Public Folder List Command

Option	Function
-m	The name of the folder (match), within quotes, from which to start the listing
-d	The number of levels (depth) to list. The default is 1. To list all levels below the current one, specify a depth of 0
-s	Display the size of the public folder in KB
-S	Display whether a permission file is associated with the folder. A '+' character at the end of the folder's line says one is

Permissions for Public Folders

Access to public folder functionality through Scalix depends on the permissions of the clients. The level of access any single user is granted determines whether they can see a folder, read its contents, change those contents, add to them, or delete from them.

Anybody who sets up a folder hierarchy can assign other users access to manage its folders and subfolders. And those managers (aka administrators) can, in turn, assign permission levels for the folders and subfolders they own.

By default, subfolders take on the permissions of their parent but those settings can be changed.

The four basic levels of permission are:

- Read
- Write
- Edit
- Delete

These various levels of permission are combined into “roles” that are set either on the command line or the client.

The command line roles are:

- **Administrator:** Has read, write, edit and delete permissions on all folders
- **Local user:** Has read edit and delete permissions only on folders that are stored on their server
- **Default:** Has read permissions on all folders

The roles used in SWA and Outlook are:

- **None:** Cannot read, write, edit or delete
- **Contributor:** Can create items and see folders, but cannot edit or delete
- **Reviewer:** Can read items and see folders, but cannot edit or delete

- **Non-editing author:** Can create and read items as well as see folders, but cannot edit or delete
- **Author:** Can create and read items as well as see folders. Can edit and delete only those items they created themselves
- **Publishing author:** Can create and read items, see folders, and create subfolders. Also can edit and delete only those items they created themselves
- **Editor:** Can create and read items, and see folders. Can edit and delete both those items they created themselves and any items in their principal's mailbox
- **Publishing Editor:** Can create and read items, see folders, and create subfolders. Also edit and delete both those items they created themselves and any items in their principal's mailbox
- **Owner:** Can create and read items, create subfolders, and acts as the owner of folders. Also edit and delete both those items they created themselves and any items in their principal's mailbox
- **Custom:** Any combination of creating, reading, and viewing either messages, folders or calendar items

Setting Public Folder Permissions

By default, the owner of a public folder can read, write, edit and delete items while everybody else can only read them. You can change that default, though. There are two ways to change permissions:

- On the client
- On the command line

For more on how to set permissions through the various clients, see the online help system associated with each one.

To set permissions on the command line:

- 1 Run the following command.

```
omaddacl n -t bulletin -l <"Folder Name"> -c update
```

Where some of the available options are:

Table 2: Options for Permissions Commands

Option	Function
-t	The resource type, which in this case is bulletin
-l	The name of the resource, which in this case is the name of the folder for which you want to set permissions or the path. An example of the path would be "Top <Folder Name>"

Table 2: Options for Permissions Commands

Option	Function
-c	The capabilities you want to set, which can be: s - see r - read u - update d - delete m - modify

Other commands that relate to permissions are shown in the table below. All take the same options outlined above. For more on options you can use, see each command's man page.

Table 3: Commands for Setting Permissions in Public Folders

Option	Function
omaddacl	Adds capabilities for users to an Access Control List (permissions) for a specific resource. The capabilities are a combination of those given to the user based on the O/R address and any groups to which they belong. To use this command, you must have root or configuration capability.
omshowacl	Displays (shows) the contents of an Access Control List (permissions), showing the capabilities given to specific users for specific folders. To use this command, you must have root or configuration capability.
omdelacl	Deletes an Access Control List (permissions). To use this command, you must have root or configuration capability.
omchkacln	Displays (checks) the capabilities of a user in an Access Control List (permissions).

Alert	The creator of a public folder gets full read, write, edit, modify and delete capabilities by default.
--------------	--

Maintaining Public Folders

To prevent public folders from growing too large and unwieldy or irrelevant, you can delete items individually or in bulk, remove outdated items, or configure the system to automatically remove items after a certain period of time has elapsed.

There are many ways to do this:

- **Delete individually or in bulk:** Deletes items from a public folder or its subfolders
- **Expiry Date:** Deletes all contents of the folder automatically on a designated date
- **Expiry Delay:** Deletes all contents of the folder automatically after a specified period of time has elapsed since the last modification to that folder.

To delete items from a public folder either individually or in bulk:

- 1 Use the following command.

```
ommai ntbb -a
```

Where the available options are:

Table 4: Options for the Public Folder Maintenance Command

Option	Function
-a	Delete all items from all public folders
-m	Delete items from the public folder by the name <"Folder Name">
-e	Delete items added to the public folder more than <number of> days ago
-A	Age items in that public folder
-R	Perform the delete action (-a or -e) to all subfolders of this folder

Assigning Email Addresses to Public Folders

If you want users to email items to public folders (instead of, or in addition to the drag and drop method), you can assign email addresses to those folders.

Note	A new Management Plugin accomplishes this task in a more streamlined manner. If you choose to use the plugin instead, you can find it on the Scalix Wiki at http://www.scalix.com/wiki/index.php?title=Scalix_Wiki .
------	---

To assign a folder a mailing address:

- 1 On the command line, add an email entry to the system directory for that folder. Assuming the folder is called "Top Level", the command would be:

```
omaddent -e "S=+BB/CN=Top Level /OU1=mail node/DDT1=BB-NAME/DDV1=Top Level /IA=top.level@domain.com"
```

- 2 To add an entry for a lower-level folder, for example Top Level>Second Level, the command would be:

```
omaddent -e "S=+BB/CN=Second Level /OU1=mail node/DDT1=BB-NAME/DDV1=Top Level >Second Level /IA=second@domain.com"
```

Where the folder separator is a ">" character.

Note	If the public folder name contains non-ASCII characters such as Japanese characters or German unlaufs, DDV1-TX must be used instead of just DDV1. In this scenarios, the command in Step 1 would be: <code>omaddent -e "S=+BB/CN=Top Level/OU1=mailnode/DDT1=BB-NAME/DDV1-TX=Ümläutfolder/IA=top.level@domain.com"</code>
------	---

Synchronizing Public Folders

Public folder synchronization is the process of automatically updating public folders and their contents from one system to another. This process ensures that when you add an item to one public folder, the same item is also added to all equivalent public folders in the network.

Synchronization is accomplished through synchronization agreements. These agreements define the rules of each exchange. Each agreement defines whether items are imported or exported, and the public folders to which the agreement applies. All items within the hierarchy of a specified public folder are included in the agreement.

Synchronization is performed by exchanging mail messages between two public folder servers (BB servers). Each of these messages adds one item to a public folder, where items can include messages, calendar items, contacts, etc. The “sending” server is the public folder system that exports the BB server, the “receiving” server is the one that imports.

Alert

Before synchronizing public folders on two Scalix servers, you must have set up routes between the two. For more on how to establish routes between servers, see the *Scalix Server Setup and Configuration Guide*.

Synchronization Concepts

Public Folder synchronization is used to replicate information across a number of servers in your network.

For example, a corporation in New York has a server (A) that has a Public Folder called Sales. This contains attachments (mail messages, text files, distribution lists, and so on) that relate to the Sales function of the corporation. To give the branch office in Los Angeles access to the information on the Sales Public Folder, you set up an export synchronization agreement to export the Sales Public Folder to the Los Angeles server (B). A corresponding import agreement must also be set up on Server B.

This creates a Public Folder on Server B named Sales that contains exactly the same attachments as the Sales Public Folder on Server A. At regular intervals, updates to the Public Folder on Server A are automatically sent to the Public Folder on Server B to ensure that the two public folders are synchronized.

The items on Sales on Server A are master items, because they were originally created on Server A. The same items on Sales on Server B are secondary items, because they are copies of the master. Whether an item is a master or a secondary is important when deleting or modifying items.

Users on Server B can also attach items to their version of Sales. However, these are not replicated on Server A until you create the appropriate export synchronization agreement on B and a corresponding import agreement on A. When you do this, the two public folders become identical, except that master items on A are secondary items on B, and secondary items on A are master items on B.

You specify whether object files attached to messages or items are included when the message or item is exchanged during synchronization. This is set using the `BBS_SEND_OBJECT_FILES` option in the general configuration file (`~/sys/general.cfg`). See “Configuration Options” on page 151 for more information.

Synchronization Agreement Guidelines

Matching agreements must exist on the exporting and importing systems before items can be exchanged. Typically, a number of agreements are specified on each system, with each agreement specifying the exchange for several public folders.

Alert

The deletion of public folders is not replicated across servers during synchronization. In the example above, if you delete the Sales public folder on Server A, the Sales public folder on Server B is not deleted. You have to remove it manually. For the contents of public folders, deletion of items from the master server replicates, but deletion of items from the secondary does not.

Follow the following guidelines when setting up synchronization agreements:

- Always activate the importing synchronization agreement before activating the corresponding remote exporting agreement. This prevents exported items from arriving at the importing system before it can accept them.
- With a two-way import/export agreement, activate the corresponding agreement on the remote system at the same time.
- If the primary mailnode on a system changes, all agreements (both import and export) must reflect the new mailnode name.
- Use the wildcard character (*) when specifying the subjects of top-level public folders to import or export. This avoids adding individual agreements for each folder. However, beware that indiscriminate use can lead to significantly increased network traffic.

Wildcard characters represent zero or more characters. One or more wildcard characters can be placed anywhere in the subject string.

Note

If using a multibyte character set, wildcards can be placed only at the beginning and end of the subject string. Also, the output from the Scalix commands displays multi-byte characters as asterisks, so users cannot distinguish between a subject containing wildcards and one containing multibyte characters.

Synchronization Prerequisites

The following minimum-level access capabilities are required to import items:

- The originator of the message (the exporting server) has use access on the importing server
- The originator of the message (the exporting server) has read access to the public folder area, and attach and delete access to the folder being synchronized.
- If the top-level public folder does not already exist, then the originator of the message must have attach and delete access to the public folder area.

Synchronization and Permissions

After a synchronization:

- Synchronization messages (type OMSYNC) are automatically given the capabilities of the admin and default groups in addition to any permissions explicitly granted to the O/R Address pattern of the originator..

Note

An originator recipient address (O/R address) is an attribute list that distinguishes one user, or distribution list, from another and defines the user's point of access to the message handling system or the distribution list's location.


- The minimum access permissions required to export items are the combined permissions of the standard groups local, default and admin plus any permissions explicitly granted to the O/R Address pattern of the originator (OMSYNC +BB/local_primary_mailnode) must give read access to the public folder area and the public folders being synchronized.
- The public folder permissions themselves are not exported when public folders are synchronized. New top-level public folders created as a result of a synchronization agreement use the default settings. New subfolders created from synchronization agreements inherit their permissions from their parent.

Synchronization Topologies

Public folders can be synchronized in two ways:


- Chains, where each server passes new items to the next in the chain
- Hubs, where every server receives updates from one central server.

In a chain, a user adds an item to a public folder on Server A. This item is replicated to the equivalent folder on Server B, and then to Servers C and D.

Bulletin board servers


In a hub, all changes are communicated to Server E, which then propagates them to the other servers.

Bulletin board servers



A hub can be extended to a linked hub (Server J), as shown below.

Bulletin board servers



Other Synchronization Topologies

There are a number of other synchronization topologies. Some of these, however, are complex to administer, could cause network congestion, and might not give the required results.

For example, grids (where all systems are updated by each other) and loops (linked chains) can result in duplication of synchronization messages, and lead to a loss of synchronization.

A Sample Synchronization Agreement

The basic commands used for synchronization are `omaddbbsa`, `omdelbbsa`, `ommodbbsa`, and `omlistbbsa`.

Because there are so many unique variables in setting up a synchronization agreement, no one procedure will work for every situation. This is an example of the most basic, two-way synchronization agreement.

To run a simple, two-way public folder synchronization agreement:

- 1 Set up an import agreement on one server and an export agreement on the other by running the following commands, each on its own server:

On Server A: `omaddbbsa -i y -m "OMSYNC +BB/serverB, mai l node" -s "BB subject" -t "010000 00:00"`

On Server B: `omaddbbsa -e y -m "OMSYNC +BB/serverA, mai l node" -s "BB subject"`

Where the `-t` option specifies when the agreement should begin. The date format is `yymmdd hh:mm`.

To do this the other way around:

On Server A: `omaddbbsa -e y -m "OMSYNC +BB/serverA, mai l node" -s "BB subject"`

On Server B: `omaddbbsa -i y -m "OMSYNC +BB/serverB, mai l node" -s "BB subject" -t "010000 00:00"`

- 2 In the file `~/sys/general.cfg`, change the default interval from one hour to a smaller time period such as one minute (the example below shows one minute).
`BBS_CUST_CHECK_TIME=1`
- 3 Restart the synchronization service and enable auditing to see the messages transfer between machines.
`omconfaud bbs 15`
`omoff -d 0 -w bbs; omon bbs`
- 4 To check that mail is flowing correctly, review the messages in the directory `~/logs/audit`.
- 5 Repeat as needed until all directories have synchronization agreements.

Forwarding Public Folder Items

If you have a Scalix public folder hierarchy co-existing with a legacy system such as Microsoft Exchange or Lotus Notes, you can configure Scalix to forward any public folder postings to the legacy system's public folder hierarchy. This sends a copy of every message posted to Scalix to the similarly-named folder on Exchange or Lotus.

This forwarding arrangement creates a directory entry associated with the public folder and creates an SMTP address to receive inbound mail. After creating that directory entry, you must provide a "forward to external address" option in the Scalix public folder synchronization mechanism.

Alert

Make sure you migrate Exchange public folder data to Scalix before enabling auto forwarding. If you don't, you see duplicate messages in Scalix folders.

Note

This procedure provides functionality equivalent to the Exchange public folder forwarding rule. That is, only new public folder messages are synchronized. Modified and deleted messages are not.

To forward public folder postings to a legacy system:

- 1 Run the following command.

```
omaddbbsa -f <Forwarding Address> -s <Public Folder Name>
```

Posting to Public Folders by Email

You can enable a public folder to receive mail on an Internet address so that users from outside can post to it.

To post to a public folder by email:

- 1 Run the following command.

```
# omaddent -e "S=+BB/CN=ABC/OU1=sxhost1/DDT1=BB-NAME/DDV1=KCI /
IA=abcpost@xyzcorp.com"
```

where ABC is the name of the public folder and sxhost1 is the mailnode. This enables people to send to abcpost@xyzcorp.com and it goes directly onto the folder.

Note

If the public folder name contains non-ASCII characters such as Japanese characters or German unlaufs, DDV1-TX must be used instead of just DDV1.

Hiding Public Folders

If needed, you can hide public folders from view in IMAP clients such as SWA and Thunderbird by making a small change in the general configuration file. You can do this on a per-user level.

To hide public folders from view:

- 1 Go to the file ~/scalix/sys/general.cfg.
- 2 In the value IMAP_PUBLIC_FOLDERS, change the value to FALSE to hide public folders from view.
- 3 Stop and restart the IMAP and Remote Client Interface processes.

```
omoff -d0 -w imap rci
```

```
omon imap rci
```

Public Folder Commands

The following table lists commands associated with public folders.

Table 5: Public Folder Commands and their Descriptions

Command	Description
omaddbb	Add a top-level Public Folder
omdelbb	Delete a top-level Public Folder

Table 5: Public Folder Commands and their Descriptions

Command	Description
omlistbbs	List top-level public folders
ommaintbb	Maintain top-level public folders
ommodbb	Modify the subject of a top-level Public Folder
omshowbb	Show details of a top-level Public Folder
omaddbbsa	Add a Public Folder Synchronization agreement
omdelbbsa	Delete a Public Folder Synchronization agreement
omlistbbsa	List Public Folder Synchronization agreements
ommodbbsa	Modify a Public Folder Synchronization agreement

Access Control Lists

This chapter describes Access Control Lists (ACLs), which control user access to Scalix resources such as public folders, directories, scripts and services. The following topics are detailed:

- "About ACLs" on page 114
- "ACL Commands" on page 115
- "Creating ACLs" on page 115
- "ACL Address Patterns" on page 118
- "Combining Users and Permissions" on page 119

About ACLs

In Scalix, you can create Access Control Lists (ACLs) to limit permissions for services, scripts, public folders and directories. ACLs list users and the permissions they have for any given resource.

While ACLs for services, scripts and directories are created and removed explicitly via commands such as `omaddacl` and `omdelacl`, ACLs for public folders and directories are created and deleted automatically when a folder or directory is either created or deleted.

In Scalix ACLs, users can be listed by standard groupings or OR addresses. The standard groupings are: local administrators (`admin`), local users (`local`), and everyone/default (`default`). Additionally, specific users can be added to a list by their O/R address or grouped more efficiently using "wildcards" in place of specific address attributes. Each entry in an ACL refers either to a standard group or to an O/R address pattern.

Every resource has configuration permission; that is, the permission to modify the ACL. This permission is always given to the standard group "local administrators" (and the root user) when the ACL is created. It can be changed later.

The ACL configuration file is named `acl.cfg` and is located in the directory `~/sys`. The ACLs themselves are located in directories under `~/acl`.

ACL Commands

The following table lists commands associated with Access Control Lists.

Table 1: ACL Commands and their Descriptions

Command	Description
omaddacl	Add an Access Control List
omaddacln	Add permissions for users to an Access Control List
omchkacln	Check permissions of a user in an Access Control List
omdelacl	Delete an Access Control List
omdelacln	Delete permissions for users from an Access Control List
ommodacln	Modify permissions for users in an Access Control List. Permissions can be removed using the ommodacln command and a dash (-) in front of the permission you want to remove.
omshowacl	Show the contents of an Access Control List

Creating ACLs

The methods for creating service, script, public folder and directory ACLs differ in subtle ways. There are a few commands that apply to all:

You can create an ACL using the command:

```
omaddacl -t <type> -l <name>
```

Where

- <type> is the "type" of resource such as service, script, public folder or directory
- <name> is the name you give to this list.

An ACL is identified by its "type" and "name".

The following table lists the resource "types" available:

Table 2: Access Control List Resource Types

Resource	Type	Abbreviations	Value
Services	service	svc	s
Request Server Scripts	request server	req	r
Public Folders	bulletin	bb	b
Directories	directory	dir	d

The resource types are defined in the file named acl.cfg, which is located in ~/sys/.

The name of an ACL depends on its resource type. The following table lists how ACL names are determined:

Table 3: Naming Conventions for Access Control Lists

Resource	Name determined by
Services	the queue name of the service.
Request Server Scripts	the file name of the request as listed in the <code>/opt/scalix/req</code> directory.
Public Folders	the temporary or absolute reference number of the public folder as listed by the <code>omlistbbs</code> command. (Use 0 for the public folder area itself.)
Directories	the name of the Directory as listed by the <code>omlistdirs</code> command.

Service ACLs

You can add a user to a service ACL using the command:

```
omaddacl n -t service -l queue_name
```

Directory ACLs

You can add a user to a directory ACL using the command:

```
omaddacl n -t directory -l Dir_name
```

Public Folder ACLs

Public folder ACLs are handled differently. In Outlook, most public folder access settings can be handled through the Outlook client. Elsewhere, they can be set on the command line.

Whether in Outlook or another mail client, ACLs have “implicit” permissions unless otherwise specified. That is, they inherit their permissions from the parent folder unless stated otherwise by creating new “explicit” permissions. Explicit permissions added later overwrite the implicit permissions.

When you delete a public folder’s permissions, its ACL changes from explicit to implicit and the ACL settings change from those in the ACL file to those of its parent.

You can add a user to a public folder ACL using the command:

```
omaddacl n -t bulletin -l BB_ref
```

The following table lists the levels of access that can be given to users or groups of users.

Table 4: Public Folder Permissions and their Descriptions

Permis- sion	Short- hand	Description
Owner	O*	Grants all permissions in the folder. This user can create, read, modify, and delete all items, including e-mail messages, appointments, contacts, tasks, posted items, and documents and files. Also can create subfolders. The folder owner also can change permission levels that others have for the folder. (Does not apply to delegates.)
Contact	C	Grants the user folder contact status. Folder contacts receive automated notifications from the folder.
Create	c	Grants the user permission to post items in the folder.
Read	r	Grants the user permission to open any item in the folder.
Folder	f	Grants the user permission to see the folder.
Edit All	E	Grants the user permission to edit any public folders, whether his own or owned by somebody else
Edit Own	d	Grants the user permission to edit only his own public folders
Delete All	D	Grants the user permission to delete any public folders, whether his own or owned by somebody else
Delete Own	d	Grants the user permission to delete only his own public folders
Visible	v	Makes the folder visible

* The shorthands 'O', 'C', etc. are only used for setting the UAL_FLDR_ACL_DEFAULT option in the file general.cfg. The command omaddacln only accepts the full words or those defined in the acl.cfg file. For more on the UAL_FLDR_ACL_DEFAULT setting, see the chapter titled "Configuration Options" on page 151.

Note that, if an item attached to a public folder is itself a public folder, access to it is determined by its own ACL, and not the ACL of its parent folder.

Default ACLs

The following table lists the default ACLs for a top-level public folder.

Table 5: Default Public Folder ACLs and their Permissions

User	Permission
Scalix Administrators	Create, read, subfolder, edit own, delete all, owner, contact, visible
Local	None
Default	Create, read, edit own, delete own, visible

You can change the default level of access that “other users” are granted by setting the general configuration option `UAL_FLDR_ACL_DEFAULT`. When you create a nested public folder, it has a default ACL that is copied from its parent public folder.

ACL Address Patterns

You can use originator recipient addresses (O/R addresses) to identify individual users in ACLs. The following rules apply when specifying O/R address patterns in ACLs:

- The O/R address attributes by which a user is identified in an ACL are restricted to the mnemonic address form and are hierarchically ordered.

Table 6: O/R Addresses

<ul style="list-style-type: none"> • Country Name • Administration Domain Name • Private Domain Name • Organization Name • Organization Unit Name 1 	<ul style="list-style-type: none"> • Organization Unit Name 2 • Organization Unit Name 3 • Organization Unit Name 4 • Personal Name
--	---

- An attribute value must be fully specified, partly represented with a wildcard, wholly represented with a wildcard, or left blank.

If an attribute value contains a wildcard, either wholly represent all less-significant attributes with a wildcard or leave them blank.

If an Organizational Unit Name is left blank, leave all less significant Organizational Unit Names blank.

The Organization Name, Organizational Unit Name, and Personal Name attributes are specified in either printable strings or teletex strings, or both. If both forms are specified, and one form of the attribute value is represented by a wildcard, then the other form must be represented also with a wildcard to the same extent.

Matching Addresses to O/R Address Patterns in ACLs

The following rules are used when matching the O/R address of a user to an O/R address pattern in an ACL entry:

- Match characters regardless of whether they are uppercase or lowercase.
- Ignore address attributes that are not used by the mnemonic address form.
- If an address pattern specifies both printable and teletex strings for an attribute, and the address being matched contains one form only, then the other form in the address pattern is ignored.

Match each attribute:

- A specified attribute matches if each character compares "one for one".
- An attribute partly represented with a wildcard matches if each character in the specified part of the attribute compares one for one.
- A blank attribute matches if the attribute is also blank in the address of the user.

- An attribute wholly represented by a wildcard matches anything.
- If the address of the user matches an address pattern, the user is granted the capabilities specified for that address pattern.

Examples of O/R Address Patterns in ACLs

The O/R address for a user named "John Doe" is:

```
G=John/S=Doe/CN=John Doe
OU1=paris/OU2=sales/OU3=mis
O=pinewood/P=forester/A=atl as/C=fr
```

It matches the following address patterns:

```
*/CN=*/OU1=paris/OU2=*/OU3=*/O=pinewood/P=forester/A=atl as/C=fr
*/CN=*/OU1=*/OU2=*/OU3=*/OU4=*/O=*/P=forester/A=atl as/C=fr
*/CN=*/OU1=*/OU2=*/OU3=*/OU4=*/O=*/P=for*/A=atl as/C=fr
G=John/S=Doe/CN=John Doe/OU1=*/OU2=*/OU3=*/OU4=*/O=*/P=*/A=*/C=*
```

The address for "John Doe" does not match the following valid address patterns:

```
*/CN=*/OU1=paris/OU2=*/O=pinewood/P=forester/A=atl as/C=fr
*/CN=*/OU1=paris/OU2=sales/OU3=mis/P=forester/A=atl as/C=fr
G-TX=John/S-TX=Doe/CN-TX=*/OU1=*/OU2=*/OU3=*/OU4=*/O=*/P=*/A=*/C=*
*/OU1=*/OU2=*/OU3=*/OU4=*/O=*/P=for*/A=atl as/C=fr
```

The first address pattern does not specify an OU3. The address for John Doe does have an OU3 equal to mis. Similarly, the Organization Name in the second address pattern is blank. In the third address pattern, only the teletex string version of the Personal Name is specified rather than the printable string version. In the fourth address pattern, the Common Name is blank.

The following address patterns are not valid:

```
*/CN=*/OU1=paris/OU2=sales/OU3=mis/O=*/P=forester/A=atl as/C=fr
*/CN=*/OU1=paris/OU3=mis/O=pinewood/P=forester/A=atl as/C=fr
*/CN=*/OU1=paris/OU2=sales/OU3=mis/O=pinewood/P=forester/A=atl as/
C=fr
```

In the first address pattern, the Organization Name is represented with a wildcard. Because of the lesser significance of the Organization Unit Names, these should also be represented with wildcards. In the second address pattern, OU2 is blank. Therefore, OU3 must also be blank. In the third address pattern, the wildcard character is incorrectly positioned in the Organization Name. The wildcard character must be at the end of the partial value.

Combining Users and Permissions

If a specific user is a member of more than one group, the permissions given to each group are combined for that user. The following example shows a Directory ACL:

Table 7: Combining Users and Permissions

User	Permission
John Doe/ny,hq,mis	update

Table 7: Combining Users and Permissions

User	Permission
<code>*/ny,*</code>	<code>modifyself</code>
<code>*/*,*,*,*/**/**/*</code>	<code>read</code>
Local Administrators	<code>config</code>
Local Users	<code>none</code>

The `omshowacl` command outputs data in positional format.

The local user John Doe is explicitly given update permission. The user is also part of the group `*/*,*,*,*/**/**/*` and so has read permission.

If John Doe is an Administrator for this system, then John Doe also has config permission.

Note that John Doe/ny,hq,mis does not match `*/ny,*` because a wildcard has not been specified for the third Organizational Unit Name.

When a user attempts to use a resource, Scalix searches the ACL entries of the resource as follows:

- 1 The group entries in the following order:
 - Default
 - Local users
 - Local administrators
- 2 The address entries for an address pattern match.

This search order is used for all ACL entries. Because the group entries are checked in ascending order, a local administrator has the combined permissions of a default user and a local user, and a local user also has the permissions of a default user. If the group entries give sufficient permissions, the address entries are not searched. If Scalix does not find sufficient permissions within the group entries, it starts checking the address entries. If Scalix does not find an address match in the address entries, permission is denied.

Working with Scalix Directories

This chapter explains about Scalix's internal directories. If you need information about how to integrate with external directories such as Microsoft Active Directory or LDAP, see the *Scalix Server Setup Guide*.

Contents

This chapter includes the following information:

- "Directory Overview" on page 121
- "Structure and Functions of Shared Directories" on page 122
- "Introduction to Directory Commands" on page 123
- "Listing Fields In The SYSTEM Directory" on page 124
- "Creating New Directories" on page 124
- "Adding and Modifying Directory Entries" on page 125
- "Searching Directories" on page 125
- "Using the Client Directory Access Server" on page 126

Directory Overview

Scalix directories are made up of entries that identify users (or entities) and their attributes.

There are two types of directories in Scalix:

- **Shared Directories:** Can be accessed by a number of users.
- **Personal Directories:** Can only be used by a specific user.

Scalix's internal services and gateways as well as its external clients use directories to resolve incomplete addresses or help route email if needed. Clients can also use directories to obtain additional information such as telephone numbers, job titles or postal addresses of users.

Scalix must have a default shared directory, which most often is either provided by an LDAP or Microsoft Active Directory. The Scalix router uses this directory when it resolves addresses. This directory also contains the master list of all users in the network.

You can add other shared and personal directories to Scalix to fulfill the requirements of specific users or groups of users. Access to shared directories is controlled through the use of Access Control Lists

The attributes that make up directory entries are defined in the attribute definition file. This file contains internal attribute tags, syntaxes, and lengths of all attributes defined for use on the Scalix system. You can customize Scalix directories by adding fields that are automatically generated each time you add an entry to the directory.

There can be any number of Scalix directories. Each one is in a database structure in a separate directory under `~/dir`. Scalix provides a number of diagnostic utilities for monitoring the use and integrity of directory databases.

The Scalix directory system is not only used for the storage and retrieval of names and addresses, but also contains Public Distribution Lists (PDLs). You can manage access to PDLs using Access Control Lists. Also, using the Scalix LDAP server, you can access a Scalix directory using the LDAP protocol.

The internal attribute tags defined in the attribute definition file are mapped to language dependent tags and descriptions in the localized attribute file. The localized attribute file provides local language display of attribute tags and descriptions.

Structure and Functions of Shared Directories

One of the functions of the Scalix directory system is the storage and retrieval of names and addresses. This function is shared across the following directories:

- The SYSTEM directory is the default directory for name and address storage and retrieval. Clients access this directory when a user enters a full or partial name in the "To" field and the client returns a list of possible names. In addition, the system accesses this directory to verify recipient mailnodes if additional information is needed to route a message.
- The USERLIST directory verifies local mailnode and Scalix recipient information. This directory contains information about all the local addresses and valid recipients on the Scalix Server. Each time you add a mailnode or local recipient, an entry is automatically added in the USERLIST directory. A user entry includes items such as name, mailnode, Scalix user ID, Linux user ID, privileges, and Scalix password (including when the password was last changed).

You can access the USERLIST directory like other Scalix directories. Because the USERLIST directory is "hidden", you must specify `-t h` (type=hidden) if you want to search the directory. For example: `omsearch -d USERLIST -t h -e s=*`

Caution

Scalix recommends that you do not manually edit the USERLIST directory. Use commands such as `ommodu` and `ommodmn` to modify recipient and mailnode information. However, this directory is hidden (the `omlistdirs` command does not display the existence of this directory).

- The FREEBUSY directory allows MAPI users to share calendar information. This directory is created automatically when you install Scalix. The local calendar information of each user is periodically added to the shared FREEBUSY directory on the local Scalix Server. All users who want to share calendar information need to have an entry in this directory.

When a user attempts to schedule an event involving more than one user in the client calendar application, Scalix queries the appropriate FREEBUSY directories on the Scalix Server to determine the availability of users and, if necessary, other resources such as facilities.

Directory Entries

A directory contains addressing information (including the Internet address) and can store additional information such as a telephone number, office location, and company name.

You can add customized attributes and these are entered or displayed using a TAG=value pair. The TAG identifies the type of attribute, and implicitly defines the syntax and size of the value that the attribute can have.

The dir.attrs file defines attribute types and the diratt.loc file (one for each language installed on the Scalix server) provides language dependent tags and descriptions for each attribute type. New attribute types are defined by editing these files. You can display attribute types available on Scalix by entering the omshowatt command.

Directory entries are added, deleted, and modified using the omaddent, omdelete, and ommodent commands

Introduction to Directory Commands

Use the following commands to manage Scalix directories:

Table 1: Commands for Working with Scalix Directories

Command	Description
omaddent	Add one or more entries to a directory
omdelent	Delete one or more entries from a directory
omdiropt	Optimize a directory
omdoptall	Optimize all directories
omfmtent	Format directory and address attributes
omlistdirs	List directories
ommoddir	Modify a directory
ommodent	Modify a directory entry
omremdir	Delete a directory
omsearch	Search a directory
omshowatt	Show available attribute types

Listing Fields In The *SYSTEM* Directory

To view the fields defined for the *SYSTEM* directory, enter this command:

```
omshowatt
```

The following information appears, with the output sorted by Field Name, Data Type, Length, and Descriptive Name—as illustrated here.

Table 2: Fields in the System Directory

Field Names	Data Types	Lengths	Descriptive Names
S	KX	40	X. 400 Surname
C	X	3	X. 400 Country Code
DIT-RDN		U/L	DIT Relative Distinguished Name
DIT-RDN-ID	KPS	40	DIT RDN Global Identifier
DIT-PARENT-ID	KV	40	DIT DN Parent Identifier
DIT-OWNER	V	U/L	DIT node owners
DIT-SEARCH-REF	V	U/L	DIT search referrals
DIT-MODIFY-REF		U/L	DIT modify referral
FULL-NAME		64	Personal Full Name
LAST-NAME	K	64	Personal Last Name
MIDDLE-NAME	V	20	Personal Middle Name
FIRST-NAME	K	20	Personal First Name
ALI AS	KMSV	20	Personal Aliases
...			

Creating New Directories

To create a new directory, enter:

```
omnewdir -d directory name -t
```

Adding the *p* parameter at the end of the command creates a “private” directory that is only accessible by the user who creates the directory.

New directories that you create have default attributes.

Adding and Modifying Directory Entries

To add a directory entry to the SYSTEM directory:

- 1 Enter the following command.

```
omaddent -e "s=lastname/g=firstname/ou1=mailnode of unix gateway/  
ou2=value/ou3=value/cn=display name/ia=user@domain.com"
```

To modify an entry (for example, with a new field named Job Title with an associated value):

- 1 Enter the following command.

```
ommodent -e s=lastname -n Job Title="Accountant"
```

Searching Directories

An entry in a directory is retrieved using a search filter. The filter is compared with the entries in the directory. All entries in the directory that match the filter are returned.

Use the `omsearch` command to search a directory. For example, the following command returns all entries in the system default directory containing a surname of Smith:

```
omsearch -e S=smith
```

A filter is made up of "filter items" and other filters. A filter item is an attribute tag, an attribute value, and a matching operator. For example, `S=ExampleEntry`. The following table lists available matching operators:

Table 3: Matching Operators for Directory Searching

Operator	Description
=	Equal to the value of the filter item.
-	Approximately equal to the value of the filter item (supported with filter items using ASCII string type syntaxes).
>	Greater than or equal to the value of the filter item (supported with filter items using integer type syntaxes; for example <code>INTEGER</code> and <code>DATE</code>).
<	Less than or equal to the value of the filter item (supported with filter items using integer type syntaxes; for example <code>INTEGER</code> and <code>DATE</code>).
&	AND.
	OR.
!	NOT.

Use parentheses () to nest filters within other filters. For example, to search the system default directory to find all entries that have a surname of "Wolf", and a given name of "Mike" or "Mikey" but not "OU1=sales", the following filter could be used:

```
omsearch -e "S=Wol f & (G=Mi ke | G=Mi key) &! OU1=sal es"
```

Wildcards (*) can also be used in filter items with string type syntaxes to represent whole or partial attribute values (this applies to string syntaxes only).

To list all surnames, enter:

```
omsearch -e s=*
```

Or enter the following to search for entries containing a specific string of characters:

```
omsearch -e s=*wo*/ou1=*sal *
```

To display only a specific attribute value, use the -m parameter. Enter:

```
omsearch -e s=* -m OU1=
```

Only the OU1=Sales attribute displays for all the surnames returned by omsearch.

Using the Client Directory Access Server

The Client Directory Access (CDA) Server builds access tables for Scalix directories to provide sorted lists of directory entries.

The Outlook client used with Scalix requires sorted entries in the Address Books. This enables “typedown” functionality when selecting addresses in the interface. The CDA Server is used to provide the sorted lists of directory entries.

The commands omaddcda, ommodcda, omdelecda, and omshowcda, add, modify, delete, and show directories are configured for processing by the CDA Server, respectively. The omexeccda command forces the immediate processing of a directory without waiting for the next periodic rebuild of its access tables.

Options you set in the general.cfg file determine how the CDA Server operates. See “Configuration Options” on page 151 for more information.

The CDA Server periodically checks its configuration settings (by default, once every 5 minutes in ~/sys/cda.cfg) and if the processing of a directory is required, the CDA accesses the directory, extracts the required information, sorts the entries (on surname, given name, and initials by default), and stores the entries in access tables within the ~/cda directory.

As the access tables for a directory are created periodically, modifications to a directory are not immediately reflected in the access tables. Changes to directory entries become visible to a client as follows:

- Added entry: The new entry is visible only after the CDA Server (or omexeccda) rebuilds the access tables and the client closes and opens the directory.
- Deleted entry: The substitute text <Deleted Entry> appears until the CDA Server (or omexeccda) rebuilds the access tables and the client closes and opens the directory. Then neither the entry nor the substitute text appears in the directory.
- Modified entry: The change is immediately visible, but the sort order can be incorrect until the CDA Server (or omexeccda) rebuilds the access tables and the client closes and opens the directory.

To force the server to process a directory immediately, you can use the omexeccda command.

You can configure the time interval between the processing of a directory by the CDA Server. The interval is configured by the omaddcda and ommodcda commands. By default, the interval is 24 hours. The practical minimum interval depends on the time taken for the CDA Server to build the access tables. This in turn depends on a number of factors such as

system size, system resources, system loading, and directory size. If the interval is set too low, the CDA Server might continuously processes the directory. To verify the amount of time a directory takes to process, use the command `omshowcda -d Dir_name`.

To optimize the rebuilding of the access tables, configure the CDA Server to check the directory change log. Do this by setting the option `CDA_USE_CHANGE_LOG=TRUE` in the `general.cfg` file. See “Configuration Options” on page 151 for more information.

For the Outlook client user, Public Distribution List directory entries might appear not to exist if the user does not have the required privileges defined by the Access Control List (ACL) for the PDL. For example, when accessing the Address Book, any PDL for which the user does not have read privileges is replaced by the text “<Deleted Entry>”.

Starting the CDA Server

To start the CDA Server, enter:

```
omon -s cda
```

CDA Command Summary

The following table lists commands associated with the CDA.

Table 4: CDA Commands

Command	Description
<code>omaddcda</code>	Add a directory to the CDA Server configuration. You can specify the directories to be processed, configure the interval at which the directories should be reprocessed, the fields to be used to sort directory entries, and how often the CDA server re-reads its configuration details.
<code>omdelcda</code>	Delete a directory from the CDA Server configuration.
<code>omexeccda</code>	Force the CDA Server to process a directory immediately.
<code>ommodcda</code>	Modify the CDA Server configuration for a directory.
<code>omshowcda</code>	Show the CDA Server configuration for a directory.

Creating a Redirect Account

This chapter covers the process of creating a redirect account (or a catch-all alias) to hold messages that aren't delivered to other mailboxes.

Contents

This chapter includes the following information:

- “About the Redirect Account” on page 128
- “Creating the Redirect Account” on page 129

About the Redirect Account

If needed, you can create a redirect or “catch-all user” account that receives all email not caught by other, more specific addresses.

The format for configuring the redirect account is:

CATCH PATTERN RECIPIENT where PATTERN can be:

- user*
for any unknown address starting with a known string.
- @domain
for any unknown address in a known domain.
- user*@domain
for any unknown address starting with a known string but in a known domain.

and RECIPIENT is:

- The SMTP address for the catch-all account.

The catchall address can be a Scalix or a unix user but it is subject to any relay rules if the catchall address is outside the local domain.

If you don't set up a redirect account, non-deliverable mail is bounced and non-delivery reports are sent to the Error Notification User (enu). So, if you type:

omshowenu

it will show which user is currently defined as the error notification user. Unless you've changed it, this is the sxadmin user. You can change the enu using the omconfenu command (see the man pages for details).

Creating the Redirect Account

To create a redirect account:

- 1 Go to the configuration file at `~/sys/smtpd.cfg`.
- 2 Create a mailbox to receive redirected mail. To create that mailbox, add one of the three potential lines listed below:
 - Any unmatched user beginning with a known-string in any domain:
`CATCH user* catchall@domain.com`
 - Any unmatched user in a known domain:
`CATCH @domain.com catchall@domain.com`
 - Any unmatched user beginning with a known-string in a known domain:
`CATCH user*@domain.com catchall@domain.com`
- 3 Stop and restart the SMTP Relay by typing:


```
omoff -d0 smtpd
omon smtpd
```

Changing Hostnames and IP Addresses

This chapter covers the procedure for changing a machine's hostname or IP address after installation has been completed.

Contents

This chapter includes the following information:

- “Overview” on page 130
- “Changing a Hostname” on page 130
- “Changing an IP Address” on page 131

Overview

In general, we do not recommend changing a server's hostname after Scalix is installed. If, however, you absolutely must change it, you can do so.

The method for changing hostnames varies with the Linux distribution. But the basic idea is to change */etc/hosts* so that the server resolves its own IP and changes the files that the distribution uses to store the hostname. For example, */etc/default/hostname*, */etc/sysconfig/network* or similar.

Changing a Hostname

To change a hostname:

- 1 Run the following command.

```
sxmodfqdn -o <oldfqdn> -n <newfqdn>
```
- 2 Change the hostname configuration by going to the file */etc/opt/scalix/instance.cfg* and changing the old hostname to the new one in the lines:

```
OMNAME=<change old hostname to new one>  
OMHOSTNAME=<change old hostname to new one>
```

- 3 Change the hostname configuration in `~`. In this case, there are several files you need to change. Use `vi` or a `vi` macro to change the old hostname to the new one in:

```
~/caa/scalix.res/config/ubermanager.properties
~/platform/platform.properties
~/webmail/swa.properties
~/res/config/res.properties
~/tomcat/conf/server.xml
~/mobile/mobile.properties
~/platform/platform.properties
```

And in all the files found in these directories:

```
~/tomcat/connector/ajp
~/tomcat/connector/jk
```

- 4 Reboot the server.
- 5 Change all hostname references on Apache VirtualHost declarations or any other servers that communicate via hostname identifiers.

Changing an IP Address

Scalix bases most of its calls on fully qualified domain names rather than IP addresses. So the primary step in changing an IP address is updating the operating system's DNS tab. There is one additional change needed, though, to update the PostgreSQL database white list.

To change a Scalix server's IP address:

- 1 Change the IP address in the operating system's DNS tab.
- 2 Then update the IP address in Scalix by going to the following file.


```
/opt/scalix-postgres/bin
```
- 3 Run the following script to specify the new IP address. This is a space-separated list of all IP addresses allowed to access the database. Because it overrides (it's not additive), you must re-type all IP addresses accessing the database, including any additional machines running SWA, SIS, etc.


```
./sxpsql -whitelist <IP address 1> <IP address 2> <IP address 3> etc
```
- 4 In addition, change the IP address in the following file.


```
~/sis/sis.properties
```
- 5 Restart Tomcat and the PostgreSQL database.


```
/etc/init.d/scalix-postgres restart
/etc/init.d/scalix-tomcat restart
```

Recovering Deleted Items

This chapter covers the procedure for recovering items that have been deleted from the Deleted Items folder.

Contents

This chapter includes the following information:

- “Overview” on page 132
- “Recovering Deleted Items” on page 133
- “Changing the Default Hold Period on a Per User Basis” on page 133
- “Changing the Default Hold Period on a System-Wide Basis” on page 134
- “Disabling the Recovery Folder Feature” on page 134

Overview

If a user accidentally does a hard delete of an important email, or if a message or calendar item auto-expires from the Deleted Items folder, you can recover them. For a configurable period of time, the Scalix Recovery Folder holds on to emails removed from the Deleted Items folder. The default is seven days, but you can change that length of time on either a per-user or system-wide basis.

The Recovery Folder, which is created automatically during profile creation, normally is not visible to users. But it can be made temporarily visible as a regular mailbox folder in the client UI for the purposes of recovering an item that has been deleted. When visible, the items in the Recovery Folder are not modifiable until they are copied or moved to the live mailbox. Once the message is recovered, the Recovery Folder can be hidden from view again.

Items within the Recovery Folder do not count towards a user's mailbox limits.

Items cannot be recovered through the SmartCache interface. If needed, use another client such as a non-SmartCache Outlook, SWA or other IMAP client to recover lost messages.

There is a point where messages and calendar items do not exist even in the Recovery Folder. There are two ways this can happen:

- At client sign-off, the system tidies both the Deleted Items and Recovery folder, removing items that exceed the configured expiry period. The expiry period is configurable, as explained below.

- The recovery folder can be emptied manually using the `omtidyu` or `omtidyallu` commands. The Recovery Folder is treated as a separate area (identified by the letter 'r'). This area 'r' (the Recovery Folder) can be tidied by using either an age or the user's configured Recovery folder expiry period (cf. the Deleted Items folder). For example:

Recovering Deleted Items

If a user deletes an item that he or she later needs to recover, make the Recovery Folder temporarily visible so the user can retrieve the item and return it to the live mailbox. Then make the Recovery Folder invisible again.

To make a Recovery Folder visible to the user:

- 1 Run the following command with the recovery option.
`ommodu -o <username> --recovery Y`
- 2 When the user has finished recovering the item, change the visibility setting back by running the following command.
`ommodu -o <username> --recovery N`

Changing the Default Hold Period on a Per User Basis

By default, Scalix holds items in the Recovery Folder for seven days. You can change that on either a per-user or system-wide basis.

The explicit per-user option explained here overrides any system-wide option explained in the next section.

To make a per-user change to the default period that items are held:

- 1 Go to the following configuration file.
`~/sys/user.cfg`
- 2 Change the value
`RECOVERY_FOLDER_EXPIRY_TIME=<time_period>`
where `<time_period>` is the amount of time that deleted items remain in the "Scalix Recovered Items" folder before being finally removed from the system.
Some example settings for this option are:
 - `4d12h` (4 days and 12 hours)
 - `240h` (240 hours)
- 3 Restart the client.

Changing the Default Hold Period on a System-Wide Basis

By default, Scalix holds items in the Recovery Folder for seven days. You can change that on either a per-user or system-wide basis.

The explicit per-user option explained above overrides any system-wide option changed here.

To make a system-wide change to the default period that items are held in the Recovery Folder:

- 1 Go to the following file.

~/sys/general.cfg

- 2 Change the following value.

RECOVERY_FOLDER_EXPIRY_TIME=<time_period>

where <time_period> is the amount of time that deleted items remain in the "Scalix Recovered Items" folder before being finally removed from the system.

Some example settings for this option are:

- 4d12h (4 days and 12 hours)
- 240h (240 hours)

- 3 Restart the client.

Disabling the Recovery Folder Feature

If the recovery folder expiry time is explicitly configured to zero, the recovery feature is disabled. This means deleted items are not moved to the Recovery Folder and users do not have the option of retrieving messages that they have deleted by mistake.

To turn off the Recovery Folder feature on either a per-user or system wise basis:

- 1 Go to the following file.

~/sys/user.cfg (to disable a single user)

~/sys/general.cfg (to disable the feature on a system-wide basis)

- 2 Change the following value.

RECOVERY_FOLDER_EXPIRY_TIME=7d

to

RECOVERY_FOLDER_EXPIRY_TIME=0

Emptying the Recovery Folder

If needed, you can empty the recovery folder, deleting all of its items. But once this is done, none of the items can be restored.

- The recovery folder can be emptied manually using the `omtidyun` or `omtidyalldu` commands. When using these commands, the Recovery Folder is treated as a separate area (identified by the letter 'r'). This area 'r' (the Recovery Folder) can be tidied by using either an age or the user's configured Recovery folder expiry period (cf. the Deleted Items folder).

To empty the Recovery Folder:

- 1 Run the following commands.
 - # Tidy my Recovery Folder of items more than 2 days old
`omtidyu -B -n "Richard Hall" -d -T r -a 2`
 - # Tidy all Recovery folders using configured expiry times
`omtidyalldu -d -T r -a c`
 - # Clear all Recovery folders of all items
`omtidyalldu -d -T r -a 0`

Setting Message Delivery Rules on the Router

This chapter describes how to use the service router to set message delivery rules.

Contents

This chapter includes the following information:

- “About Message Delivery Rules” on page 136
- “About the Service Router” on page 136
- “Configuring Message Delivery Rules” on page 138
- “Examples of Message Delivery Rules” on page 144
- “Listing Deferred Mail” on page 147

About Message Delivery Rules

You can set message delivery rules for any given route so that special routing instructions are handled at the server level. Some reasons you might use router rules are:

- Rejecting mail from unknown or certain known senders before it reaches the client
- Deferring delivery of low priority mail until off-peak hours to preserve bandwidth
- Blocking highly sensitive messages or preventing them from leaving the company
- Preventing delivery status notification messages (DSN) from going out to the Internet or locally to other servers
- Filtering on the basis of sensitivity levels

These rules are set in the Service Router, which detects them and acts appropriately.

About the Service Router

All messages arriving on the system or generated by the system pass through the Service Router, which determines which services are used to deliver a message on to its next “hop” (The next Scalix interface or gateway, the local delivery service or any number of other

queues). For this reason, the Service Router is a good place to set rules that determine how, when, where and whether messages are delivered.

When the router receives a message, it performs several essential tasks, including:

- Checking the file types and performing any coercions
- Updating the recipient list to match the latest routing information
- Checking and applying rules or filters for each route
- Adding routes for recipients
- Checking that message aren't looping
- Checking that the sender has permission to use each delivery service
- Attaching messages to each delivery service queue

Messages are passed to the Service Router through its input queue, which is named ROUTER. The service name is router and the process name is service.router.

If a message is addressed to several recipients, several services might be required to deliver the message on to its next hop.

The Deferred Mail Manager

If the action of a service router rule is to defer delivery of a message, it submits the message to the Deferred Mail Manager queue (DMM.)

The DMM process (named defer.manager) monitors the DMM queue, and picks up any new messages submitted to it by the Service Router. These messages and the deferral period defined in their rules are stored in the deferred message list located at (/var/opt/scalix/msglists/DEFER.SR).

When the deferral period is reached, the DMM delivers the messages. The Deferred Mail Manager does not resubmit messages to the Service Router, but routes the messages itself

Routing Commands

Use these commands to add, delete, modify, and list routes.

Table 1: Routing Commands

Command	Description
omaddrt	Add a route
omdelrt	Delete a route
ommodrt	Modify a route
omshowrt	List routes and show how an address is routed

Configuring Message Delivery Rules

A router rule is a text file that you create. There are no restrictions on the filename you give it. And rules can be combined to create rulesets, which consist of a series of text lines.

Some guidelines for writing rules are:

- Lines that are blank or start with a hash character (#) are regarded as comment lines.
- If an argument contains white space, enclose the argument in double quotation marks (").
- If an argument contains a double quotation mark (") or a backslash character (/), precede the character with a backslash character.
- Each rule in a ruleset must be defined on a single line.
- A rule contains a number of attributes specified as TAG=value pairs, which define the criteria to be matched in a message.

Although some rules within a ruleset are associated to the sender of a message (for example the OMLIMIT-EXCEEDED rule), the rules apply only to messages sent by that user to recipients associated with routes that have the same ruleset.

To configure a Service Router rule:

- 1 Create a text file with the rule inside it. It should look something like:

```
<Message_Del i very_Attri bute> <Acti on>
```

For example:

```
TYPE=ACK ACTION=DI SCARD
```

- 2 Store it in the following directory:

```
~/rul es
```

For example, a rule with a filename of "ack-discard" would be stored as:

```
~/rul es/ack-di scard
```

- 3 Associate the rule or ruleset with the route using either the omaddrt command (for new routes) or ommodrt (for existing routes).

```
ommodrt -m <route_name> -d <fi le_name>
```

For example, to associate a ruleset named "ack-discard" with an existing route named "remote,sales", enter the following command:

```
ommodrt -m "remote, sales" -d ack-di scard
```

- 4 Restart the service router.

```
omoff -d 0 -w sr
```

```
omon -w sr
```

Note

When you specify a ruleset to a Scalix command, Scalix checks the contents of the ruleset and reports any syntax errors. If you change any rules in a ruleset, you are recommended to run the ommodrt command to verify the syntax of the new rule is correct.

Default Rules

There are two reserved ruleset names:

- **ALL-ROUTES:** If this ruleset exists, it applies to all routes, except routes for which you configure a specific ruleset. This file is used by SAC and should not be changed.
- **ALL-ROUTES.VIR:** This ruleset enables virus protection for the Scalix system. If this ruleset exists, Scalix executes the ALL-ROUTES.VIR ruleset before all other rulesets.

Note

Rulesets apply to a message only if you associate the recipient of a message with a route for which you configure a ruleset.

Rule Attributes

Rules must have attributes such as the type of message to filter, the period during which the rule applies or the action it should take.

Any files specified as input to ruleset attributes must be stored in the `~/rules` directory.

The following table lists the three categories of ruleset attributes:

Table 2: Message Delivery Rule Attributes and their Descriptions

Category	Description
message_filter	Defines the part of the message that the router checks when processing it for rules. If no message filter attributes are defined, then all message filter attributes are considered a match.
day_time	Defines the period during which the Deferred Mail Manager should perform the specified action. These attributes apply only when the DEFER action is defined in the rule.
action_info	Defines the action to be performed or the information to be supplied when the values of the <i>message_filter</i> and <i>day_time</i> attributes are matched.

Message Filter Attributes

These attributes define the parts of a message that cause the specified action to be performed. They are optional and if none are specified, the Service Router assumes a match for all attributes.

The following table lists the message_filter attributes:

Table 3: Message Filter Attributes and their Descriptions

Attribute	Description
BCC-COUNT	The number of BCC addressees that can be contained in a message. This is matched when the number of BCC addresses in the message is greater than or equal to the value specified for this tag. For example, BCC-COUNT=10 causes all messages that have 10 or more BCC addresses to be deferred or rejected, depending on the defined action.

Table 3: Message Filter Attributes and their Descriptions

Attribute	Description
DL-COUNT	<p>The number of addressees that can be contained in the primary Distribution List of a message. This is matched when the number of addresses in the Distribution List is greater than or equal to the value specified for this tag. For example, <code>DL-COUNT=100</code> causes all messages that have 100 or more addresses in the primary Distribution List to be deferred or rejected, depending on the defined action.</p> <p>Note that a Distribution List can contain one or more Public Distribution Lists. Each such PDL is initially counted as just one address by the Service Router. However, when it is expanded by the Local Delivery service and passed back to the Service Router, each individual address in the PDL will add to the number of addresses in the primary Distribution List, and so can cause the message to be deferred or rejected.</p>
OMLIMIT-EXCEEDED	<p>A percentage indicating how full a message store component is in relation to its configured limit.</p> <p>A number of message store size limits can be configured for users: overall message store, In Tray, Pending Tray, and so on. See the man page for <code>omlimit</code> for more information.</p> <p>Any NOTIFY action associated with this attribute is executed if the sender of a message has not already been notified within the last day. You can change the default value of 1 day by configuring the <code>OMLIMIT_MIN_WARN_INTERVAL</code> parameter in the <code>general.cfg</code> file.</p> <p>This filter is matched when the sender has a message store component which is at the specified percentage of its configured limit. For example, a value of 100 would match all messages from senders who had exceeded a limit by any amount; a value of 110 would match all messages from users who had exceeded a limit by 10 percent or more. A value less than 100 can be used to match messages from senders who are near to, but not yet at, one of their limits. For example, a value of 90 would match messages from senders who were at 90 percent or more of a limit.</p>
ORIGINATOR	<p>pattern</p> <p>A Scalix Address pattern to match against the originator of the message.</p> <p><code>!filename[:charset]</code></p> <p>A separate file containing one or more Scalix Address patterns to match against the originator of the message.</p> <p>The optional charset attribute specifies that the text of the file uses a character set other than IA5. The appropriate character set conversion is performed.</p> <p>Address patterns must observe the format and rules for Access Control List (ACL) address patterns.</p>
PRIORITY	<p>HIGH</p> <p>An urgent message</p> <p>MEDIUM</p> <p>A normal message</p> <p>LOW</p> <p>A non-urgent message</p>
RECIPIENT-SERVICE-LEVEL	<p>The service level of the recipient of a message.</p> <p>Service levels are assigned to users solely to enable receipt and delivery rules to be constructed. A value of 0 would check for those recipients for which a service level has not been created.</p>

Table 3: Message Filter Attributes and their Descriptions

Attribute	Description
SENDER-SERVICE-LEVEL	The service level of the sender of a message. Service levels are assigned to users solely to enable receipt and delivery rules to be constructed. A value of 0 would check for those senders for which a service level has not been sent.
SENSITIVITY	0 = Normal 1 = Personal 2 = Private 3 = Company Confidential
SIZE	The size in KBs of a message (this is matched when the message size is greater than or equal to the value specified for this tag).
SUBJECT	<p>pattern The string of text to match against the subject of the message. Wildcard characters (*) can be used. The entire subject line of the message is compared with the <i>pattern</i> for a match. This comparison is case sensitive.</p> <p>!filename[:charset] A separate file containing the string of text to match against the subject of the message. The optional charset attribute specifies that the text of the file uses a character set other than IA5. The appropriate character set conversion is performed. The entire subject line of the message is compared with the specified string of text for a match. This comparison is case sensitive.</p> <p>@script[:charset] A separate script or program containing predefined protocols syntax to communicate with the Service Router and the Deferred Mail Manager to match the contents of the subject of a message. <i>script</i> must be a readable and executable file. The optional <i>charset</i> attribute specifies that the script (or program) uses a character set other than IA5 for the display of the subject. The appropriate character set conversion is performed, and the script (or program) checks for the contents of the subject in the display character set. The script returns a status of matched, or not matched, based on the predefined protocols syntax. A sample script is provided in the section "Example Script for SUBJECT Attribute" on page 145.</p>
TYPE=ACK	Specifies the type of message is ACK or "acknowledgement," which is used to block delivery status notifications from going out to the Internet or locally to other servers.
VIRUS-FOUND	Specifies whether a message contains a virus. A value of 0 indicates that a virus was not found in the message; a value of 1 indicates that a virus was found. Include this attribute in a rule to enable virus scanning, but not virus cleaning. This attribute can only be used in the ruleset ALL-ROUTES.VIR, and applies to all routes. It cannot be applied selectively to specified routes.
VIRUS-UNCLEANED	Specifies whether a message that contains a virus could not be cleaned. A value of 0 indicates that the message was checked and was successfully cleaned of any viruses that were found. A value of 1 indicates that the message contained a virus that could not be cleaned. Include this attribute in a rule to enable virus cleaning. This attribute can only be used in the ruleset ALL-ROUTES.VIR, and applies to all routes. It cannot be applied selectively to specified routes.

Day/Time Attributes

These attributes define the period during which the Deferred Mail Manager should defer messages when all other rule attributes are matched and ACTION=DEFER is specified. See the section “Listing Deferred Mail” on page 147 for more information.

The following table lists the day_time attributes.

Table 4: Day and Time Attributes and their Descriptions

Attribute	Description
DAY	The day of the week on which to defer messages. This attribute is specified with an integer where 0=Sunday, 1=Monday, 2=Tuesday, etc. You can specify a range of days (for example, 1 – 5 for Monday through Friday) If no value is specified, all week is assumed.
TIME	<p>The time of day the DEFER action should be performed. This attribute is the local time specified in HH:MM or HH format, using the 24-hour clock (00:00 or 00 is midnight). This attribute must be specified.</p> <p>You can specify a duration of time during which deferred messages can be delivered using the following syntax:</p> <p>time_interval controls the start and stop times during which actions should be performed specified in HH:MM–HH:MM format or in HH–HH format.</p> <p>You also can specify a time at which to begin storing a batch of deferred messages and an interval of time during which to deliver the deferred messages using the following syntax:</p> <p>batch_spec controls the time to start batching messages and the interval during which to deliver the batch of deferred messages specified in HH:MM@HH:MM format or in HH@HH format.</p>

Action/Information Attributes

These attributes define the action to be taken or the information to be provided when the message_filter and date_time attributes are matched.

The following table lists the `action_info` attributes.

Table 5: Action Attributes and their Descriptions

Attribute	Description
<code>ACTION</code>	<p>An <code>ACTION</code> attribute must be specified for a rule. For actions that do not automatically return a Non-Delivery Notification, you can specify a message to be returned to the user with the <code>NOTIFY</code> tag.</p> <ul style="list-style-type: none"> • <code>ALLOW</code> Route the message immediately. • <code>DEFER</code> Defer delivery of the message during the period specified by the <i>day</i> and <i>time</i> attributes. • <code>DISCARD</code> Discard the message without returning a Non-Delivery Notification to the originator. • <code>REJECT</code> Do not route the message and return a Non-Delivery Notification to the originator. • <code>RETURN</code> Do not route the message and return a Non-Delivery Notification and the original message to the originator.
<code>NAME</code>	<p>The name of a rule within a ruleset. This name is used in any notification message generated and written to the Scalix Event Log and is useful in determining which rule is being applied when a message is routed. This tag is optional.</p>
<code>NDN-INFO</code>	<p>Enables you to replace the standard text string in the supplementary information field with the specified text when Non-Delivery Notification is required. This tag is optional. <code>NDN-INFO</code> is one of:</p> <ul style="list-style-type: none"> • <i>text</i> The text to be included in a Non-Delivery Notification. • <i>!filename[:charset]</i> A separate file containing the text to be included in a Non-Delivery Notification. <p>The optional <i>charset</i> attribute specifies that the text of the file uses a character set other than IA5. The text is converted to IA5.</p>

Table 5: Action Attributes and their Descriptions

Attribute	Description
NOTIFY	<p>Enables you to include supplementary text in the standard notification message returned to the originator when the defined action has been performed on the message. This tag is optional. When <code>NOTIFY</code> is specified, the supplementary text is imported as a text part and attached in front of the standard notification message text. <code>NOTIFY</code> is one of:</p> <ul style="list-style-type: none"> • <code>No value</code> The standard notification message text containing details of the ruleset applied and the recipients affected is used if <code>NOTIFY</code> is specified with no value. • <code>text</code> The supplementary text to be included in a message to the originator. The maximum size of <code>text</code> is 255 bytes. • <code>!filename[:charset]</code> A separate file containing the supplementary text to be included in a message to the originator. The optional charset attribute specifies that the text of the file uses a character set other than IA5. The appropriate character set conversion is performed. There is no restriction on the size of text contained in filename. <p><code>NOTIFY</code> is typically used with the <code>ACTION=DEFER</code> or <code>ACTION=DISCARD</code> attributes, since Non-Delivery Notification is not returned to the originator of the message.</p>

Examples of Message Delivery Rules

Some sample rules are:

- Deferring Messages Based on Originator

Delivery of all low priority messages sent from the lab mailnode is deferred between 11 a.m. and 6 p.m. (that is, these messages are delivered after 6 p.m. and before 11 a.m.).

```
PRI OR I TY=LOW ORI GI NATOR="*/CN=*/I ab" ACTI ON=DEFER TI ME=11: 00-18: 00
```

- Delivering Messages Based on Subject

Messages with a subject starting with the text "PUBLIC:", of normal priority and of a size 1 MB or greater are not routed, and a Non-Delivery Notification is returned to the message originator. Any other messages are routed without delay according to the normal Service Router routing process.

```
PRI OR I TY=MEDI UM SI ZE=1000 SUBJECT="PUBLI C: *" ACTI ON=REJECT
```

- Delivering Messages Based on Priority

Any high priority messages are delivered Sunday through Saturday without delay according to the normal Service Router routing process. Any medium priority messages

are submitted to the Deferred Mail Manager for deferral between 9 a.m. and 5 p.m., Monday through Friday. Any low priority messages are submitted to the Deferred Mail Manager for deferral between 6 a.m. and 9 p.m., Monday through Friday.

```
PRI OR I TY=HI GH ACTI ON=ALLOW
```

```
PRI OR I TY=MEDI UM ACTI ON=DEFER TI ME=09: 00-17: 00 DAY=1-5
```

```
PRI OR I TY=LOW ACTI ON=DEFER TI ME=06: 00-21: 00 DAY=1-5
```

- Rejecting or Deferring Message Delivery Based On Size

Any messages of a size 1Mbyte or greater are rejected, and a Non-Delivery Report is sent to the message originator. The delivery of any messages 100 KBs in size (but less than 1Mbyte) is deferred between 10 p.m. and 6 a.m.

```
NAME=Rej ect_1Mb SI ZE=1000 ACTI ON=REJECT
```

```
NAME=Defer_100Kb SI ZE=100 ACTI ON=DEFER TI ME=22: 00-06: 00
```

- Providing Supplementary Text for a Non-Delivery Notification

Messages over 1Mbyte in size are rejected, and a Non-Delivery Notification containing the text "Message too large" is returned to the message originator.

```
SI ZE=1000 ACTI ON=REJECT NDN- I NF0="Message too large"
```

- Notifying Deferred Message Delivery

The delivery of any low priority messages is deferred between 9 a.m. and 6 p.m. on Friday. A notice of deferred delivery is sent to the originator, with the text contained as a separate file in the same directory (~/rules/):

```
PRI OR I TY=LOW ACTI ON=DEFER DAY=5 TI ME=09:00-18:00 NOTI FY="!delaynote"
```

The text of the deferred delivery confirmation could be:

```
I tem 2 (del aynote)
```

To speed the delivery of normal and urgent messages during normal office hours (9am-6pm) on Fridays, any low priority messages are not delivered during these hours.

- Deferring Delivery of a Batch of Messages

Any low priority message received from 8 a.m. are held in a batch and delivered at 1-hour intervals, that is, at 9 a.m., 10 a.m., 11 a.m., and so on.

```
PRI OR I TY=LOW ACTI ON=DEFER TI ME=8: 00@1: 00
```

Example Script for SUBJECT Attribute

Below is an example of a script that can be used with the SUBJECT tag. This script (/opt/scalix/examples/general/subject) is supplied with Scalix. You can copy this script into /var/opt/scalix/rules and modify it as required. You must use the protocol syntax as defined in the comments at the top of the supplied script. You must also ensure that the script is both readable and executable.

```
#!/bi n/sh
```

```
#####
#Scal i x Router Subj ect Mappi ng Protocol s
#####
```

```

# PROTOCOLS SYNTAX:
# The following table outlines the possible commands sent by
# Scalix and the expected replies sent by the Mapper. Note:
# 1) each command/reply must end with a new line (n) character
# 2) the Mapper must NOT buffer its output, each reply must be
#    flushed
# 3) the Mapper must reply to each command
# COMMAND          REPLY          REPLY COMMENTS
# =====
# <start>          220<SP><text>    Mapper must output this when
#                                     starts up
# HELO<SP><text>    250<SP><text>    Mapper accepts Scalix session
# SUBJECT: <text>  251<SP><text>    Subject does not match require-
#                                     ment
# SUBJECT: <text>  252<SP><text>    Subject matches requirement
# QUIT<SP><text>   221<SP><text>    Mapper terminates session
# <others>        500<SP><text>    Unexpected command/syntax
#####
# handle "<start>"
# return ready status
rep="220 Subject Mapper Ready"
echo "$rep"
# loop to process commands
Quit="FALSE"
while read cmd
do
    case "$cmd" in
        "HELO"*)
            # handle "HELO<SP><text>"
            # return ok status
            rep="250 Ok"
            ;;
        "SUBJECT: PUBLIC: "*)
            # handle "SUBJECT: <text>"
            # subject matches requirement, strip off "SUBJECT: "
            subject='echo $cmd | sed -e "s/SUBJECT: //" '
            rep="252 $subject"
            ;;
        "SUBJECT: "*)
            # handle "SUBJECT: <text>"
            # subject does not match requirement, return reason
            rep="251 Subject Does Not Match"
            ;;
        "QUIT"*)
            # handle "QUIT<SP><text>"
            # return status, set flag to exit loop
            rep="221 Subject Mapper Close"; Quit="TRUE"
            ;;
        *)
            # handle "<others>"

```

```

        # return error status
        rep="500 Unrecognized Command or Syntax Error"
        ;;
    esac

    # must reply to each command
    echo "$rep"
    if [ "$Quit" != "XTRUE" ]
    then
        continue
    else
        break
    fi
done
exit 0

#####
# End of script
#####

```

Listing Deferred Mail

To view a list of deferred messages:

- 1 Enter this command:

```
omstat -d
```

You can attempt to force the delivery of a deferred message by using the `omresub` command.

Working with SIS

This chapter explains some of the intricacies for working with the Search and Index Service (SIS), including how to rebuild the index, how to configure it, the document types it handles and more.

Contents

This chapter includes the following information:

- “Overview” on page 148
- “Document Types Handled” on page 149
- “Re-Creating the Index” on page 149
- “Localizing the Search and Index Service” on page 150
- “Identifying an Individual Subdirectory” on page 150
- “Upgrading the Search and Index Service” on page 150

For information on creating indexes for existing users when upgrading to Scalix 11, see the *Scalix Installation Guide* in the chapter titled, “*Upgrading Scalix.*”

Overview

The Scalix Search and Index Service (SIS) is a set of flat files that enable the indexing and subsequent search for email messages, calendar items, and more. It consists of subdirectories designated for all individual users, who are identified by SIS URLs.

By default, SIS indexes every 10 minutes or after 200 new items have been added to the index, whichever comes first. That means new messages or calendar items do not show up in search results until after that period has elapsed. That interval can be changed -- even to the point of real-time indexing -- but higher-frequency of indexing can create a drain on other system resources.

SIS can be localized to many different languages, but can only use one language at a time.

And it has plugin capabilities so you can write your own message analyzers. For more on that, contact professional services.

The index need not be backed up because it can be re-created at any time.

Document Types Handled

The Scalix search index only parses the following document types:

- Microsoft Word, PowerPoint and Excel
- PDF
- HTML

Creating Users' Indexes

When new users are added to the system, they get an index by default that populates automatically as emails arrive. In these cases, no more action is needed.

When existing users are upgraded to 11.0 from a previous version, though, only messages that arrive after upgrade are automatically indexed. This is because previous versions of Scalix did not include the service. To add messages that predate the update to the index, you have to run the `sxmki ndex` command.

To create a user's index:

- 1 Signed on to a client as the user, perform a search. Or signed on as the administrator, run the following command.

```
sxmki ndex <user_name>
```

To index all messages for all users:

- 1 Signed on as the administrator, run the following command.

```
sxmki ndex
```

Disabling Indexing

If needed, you can disable indexing on a per-user basis, but if you do, that user's search function does not work.

To disable indexing:

- 1 On the command line, run the following command.

```
ommodu -o <user_name> --i ndex none
```

- 2 If you later re-enable indexing, you must run the `sxmki ndex` command again to index past messages.

```
sxmki ndex <user-name>
```

Re-Creating the Index

The index may need manual re-creating from time to time if it becomes corrupted, out of sync or out of date. This process can put a heavy load on the servers, so is best undertaken during off hours.

To recreate the search index:

- 3 Restart SIS Tomcat and make sure it's initialized.
- 4 Run one of the following commands:

To re-create all indexes at once

```
sxmki ndex
```

To re-create indexes in recovery mode, which removes deleted messages.

```
sxmki ndex -r
```

To re-create one user at a time

```
sxmki ndex {<username>} -r 0
```

Where <username> is the user's common name.

To re-create more than one index at a time, provide a list of names in the user-name variable.

- 5 Allow the indexes to create, which may take quite some time depending on the number of users and the size of their mailboxes.

Localizing the Search and Index Service

The Scalix Search Index Service can be configured to process text for any language. To work with different languages, it uses stemming rules for that specific language, which break down words by removing suffixes and endings just as they do with the English language. For example, a search for the English word "singing" will match the word "sing".

For more on how to localize SIS, see the *Scalix Server Setup Guide*.

Identifying an Individual Subdirectory

If needed, you can identify a user's SIS URL, which is the subdirectory of the index that belongs to that user.

To identify which subdirectory belongs to a particular user:

- 1 Run the following command.

```
omshowu -n <cn>
```

- 2 In the output, look for the following value.

```
SIS_URL
```

Upgrading the Search and Index Service

For information about upgrading SIS, see the Scalix Installation Guide.

Configuration Options

This chapter describes options you can modify in configuration files to customize the Scalix Server. If you do not need to customize any configuration options, skip this chapter.

Contents

This chapter includes the following information:

- “Configuration Files” on page 151
- “System-Wide Configuration Options” on page 153
- “Client-specific Configuration Options” on page 210
- “User-Specific Configuration Options” on page 218

Configuration Files

The general configuration files contain options that affect the behavior of the Scalix system. You can modify system-wide, client-specific, and user-specific configuration options.

Scalix includes a number of hard-coded default options. You can change these options by placing TAG=value pairs in one or more of the following configuration files listed in the following table in the ~/scalix/sys/ directory.

Table 1: Configuration Files Names and their Descriptions

File Name	Description
general.cfg	System-wide configuration options affecting the server
client.cfg/ fqdn	Client-specific configurations
user.cfg/scalix-uid	Per-user configurations
domain.cfg/domainname	Domain-specific configurations
lang.cfg/language	Language-specific configurations

Note that any values that contain underscores (_) or spaces should be specified within double quotes, for example:

```
EXAMPLE_OPTION="one_two three"
```

Some options can be set in more than one of the above files. In this case, note that user-specific options override client-specific options, and client-specific options override general options.

System-Wide Configuration Options

Table 2: System-Wide Options and their Descriptions

Option	Description
<p>ARCHIVE=TRUE</p> <p>NOTE: Because all of these settings include modifications to the general.cfg file, you must restart the service router and archiver for the changes to take effect.</p>	<p>Enables the archiving of all messages that traverse the Scalix Server. Messages are archived to the <code>~scalix/archive</code> directory. This parameter includes the following parameters:</p> <p>FALSE Disables the Archiver.</p> <ul style="list-style-type: none"> <code>arch:/path/</code> <p>This archives messages to the directory you specify. Within this directory, the Archiver automatically creates subdirectories based on the date that messages traverse the Scalix Server. For example, if a message arrives at 11:55 on June 10th, 2004, the message is archived to: <code>/path/2004-06-10/14:55+0000.12345.1</code> where <code>+0000</code> is the local time zone offset from GMT. and <code>12345</code> is the PID of the Archiver. <code>.1</code> indicates the number of messages that arrived during that second.</p> <p>NOTE: The Archiver process operates as the user <code>"scalix"</code>. If you want the Archiver to archive messages to the <code>/home</code> directory, you must configure the permissions for the <code>/home</code> directory to allow the <code>scalix</code> user write access (to this directory).</p> <ul style="list-style-type: none"> <code>inet:host.example.com inet:host.example.com:2000</code> <p>This allows connection to the host on port 25 or on a port you specify (<code>host.example.com:2000</code>). This creates an SMTP session and enables you to use third-party archiving systems.</p> <ul style="list-style-type: none"> <code>ARCHIVE=bcc:archive@example.com</code> <p>This forwards to a designated "bcc" mailbox created solely for archiving purposes a blind (bcc) copy of every message that is sent. You must create this mailbox before adding this parameter. We recommend it be on a separate box because archive files use significant memory space.</p> <ul style="list-style-type: none"> <code>file:/path/archive_file_name</code> <p>This writes all messages to a single file. You cannot use this option with auxiliary processes.</p> <ul style="list-style-type: none"> <code>fork:/bin/archive_script.sh</code> <p>The archiver forks the script and communicates with SMTP using stdin and stdout to the script.</p>
ARCH_TNEF_ENCODE=TRUE	<p>Sets TNEF as the message format for archived messages. By default, the Archiver converts messages to MIME format and consequently loses some MAPI information (if applicable).</p>

Audit Log Options

Table 3: Audit Log Options and their Descriptions

Option	Description
AUD_88_NAMES=TRUE	Sets audit logging on for the X.400 1988 address attributes. By default, only the 1984 attributes are logged.
AUD_LOG_UX_NAME=FALSE	By default, users are identified in the Audit Log by their Scalix IDs. Set this option to TRUE if you want users to be identified by their Linux user names.

Auto Actions Options

Table 4: Auto Action Options and their Descriptions

Option	Description
AA_DEFAULT_LOGGING_ON=TRUE	Sets the default setting for logging of automatic actions to on for all users.
AA_GLOBAL_LOGGING_OFF=TRUE	Stops all logging of automatic actions even if configured in a client. AA_GLOBAL_LOGGING_OFF overrides AA_DEFAULT_LOGGING_ON if it is set.
AA_MAXCFG_LOG_SIZE= <i>size_in_bytes</i>	Sets, for all users, the maximum size of their automatic action log file. The maximum size is 65536 bytes.
FLT_ESC_NO_CONV=TRUE	If set to TRUE , a serious error is reported when the character set for a filter and that for a string being filtered are not of the same kind while filtering for automatic actions. If set to FALSE , this is reported as a failed match.

Client Directory Access Server Options

Table 5: Client Directory Access Server Options and their Descriptions

Option	Description
CDA_CHECKTIME= <i>minutes</i>	Set the time interval, in minutes, at which the CDA Server checks its configuration for Directories that need processing. The default is 5 minutes.
CDA_USE_CHANGE_LOG=TRUE	Set this option to optimize the rebuilding of Directory access tables by the CDA Server. By default, the CDA Server will rebuild the access tables periodically. If CDA_USE_CHANGE_LOG=TRUE is set, the CDA Server first checks the change log for the Directory, and only rebuilds the access tables if the change log shows that the Directory has been modified. If the Directory does not have a change log, the CDA Server will configure one.

Daemon Options

Table 6: Daemon Options and their Descriptions

Option	Description
DADM_DAEMON_TIME_TO_EXIT= <i>seconds</i>	The number of seconds Scalix will wait for a daemon to exit before sending a <code>SIGKILL</code> signal to stop the daemon process. The default is 30 seconds; if daemon processes are being stopped too quickly, increase this number. Use this option with caution as the <code>SIGKILL</code> signal does not allow the daemon process to tidy up before it stops.

Directory Options

Table 7: Directory Options and their Descriptions

Option	Description
CU_LOG_OLD_DIR_CMDS=TRUE	If set, the old Directory commands <code>omadddir</code> and <code>ommoddir</code> are used when logging changes to the Directory update file (<code>om_record</code>), instead of the new <code>omaddent</code> , <code>omdelent</code> , and <code>ommodent</code> commands.
DA_IGNORE_INDEXES= <i>attribute,attribute,...</i>	Under normal circumstances, Scalix will fail to locate a match for any Scalix attribute which is keyed but for which the index does not exist. Use this option to specify those Scalix attributes that are keyed but for which Scalix should search sequentially, rather than attempt to use the indexes for the attributes. <i>attribute,attribute,...</i> is a comma-separated list of Scalix internal attribute names. Do not insert spaces after the commas. This option can be used for newly keyed attributes for which the indexes have not yet been built.
DIR_IA_UNIQUECHECK_OFF= <i>directory-list</i>	Use this option to specify those directories in which uniqueness checking of the Internet address attributes is turned off. The value of this option is a comma-separated list of Scalix directories or <code>ALL</code> .
DIR_IA_UNIQUECHECK_ON= <i>directory-list</i>	Use this option to specify those directories in which uniqueness checking of the Internet address attributes is turned on. The value of this option is a comma-separated list of Scalix directories or <code>ALL</code> .
VI_NON_SUPP_ATTS_UNIQUE=TRUE	By default, the combined values of the O/R Address attributes (shown by an X in the <code>omshowatt</code> command) in a Directory entry must be unique. If they are not, the entry cannot be added or modified. If <code>VI_NON_SUPP_ATTS_UNIQUE=TRUE</code> is set, this rule is relaxed and the combined values of all attributes, other than supplementary attributes, are used to determine the uniqueness of the entry rather than just those of the O/R Address attributes. (The same effect can be obtained for an individual session by setting and exporting the environment variable <code>VI_NON_SUPP_ATTS_UNIQUE=TRUE</code> .)

Table 7: Directory Options and their Descriptions

Option	Description
VI_SORTED_VISTA_DATABASE= FALSE	Determines whether the matching entries resulting from a directory search are returned sorted or in random order. When this option is <code>FALSE</code> (the default), the matching entries are returned in random order. Set this option to <code>TRUE</code> to cause the matching entries to be sorted by key. If you set this option to <code>TRUE</code> , you must rebuild all the directories (using <code>diropt</code>) for it to take effect.

Directory Relay Server Options

Table 8: Directory Relay Server Options and their Descriptions

Option	Description
<code>DRS_HOST_RETRY_TIMEOUT=seconds</code>	If the remote directory access mechanism fails to contact a remote host, it avoids retrying the connection for the number of seconds specified by this option. If a subsequent request for the same host occurs within the specified period, an error is returned immediately. The default value is 60.
<code>DRS_MAX_CHILDREN=number_of_child_processes</code>	Specifies the maximum number of processes the Directory Relay Server can support at once. Each process has a separate bind to the Directory Information Base, and some X.500 implementations can support only a limited number of binds at once. Therefore, specify the maximum number of processes with this in mind.
<code>DRS_RESERVED_CHILDREN=number_of_child_processes</code>	Defines the minimum number of child processes that the Directory Relay Server will maintain at once. Each child process has its own bind to the Directory Information Base. The greater the minimum number of child processes, the better the performance of Directory lookups. However, each process consumes resources. You should not normally set this to a value lower than 3: one process each for Local Delivery and the Service Router, and one to handle requests from the UAL and <code>omsearch</code> . The default value is 3.

Directory Synchronization Options

Table 9: Directory Synchronization Options and their Descriptions

Option	Description
<code>DR_NO_MOD_STRIP_PDL=TRUE</code>	Determines the amount of information included for an entry in the Directory change log when changes are made to PDL members within the exporting Directory using the following PDL commands: <code>omaddpdln</code> <code>omdelpdln</code> <code>ommodpdln</code> When set to <code>TRUE</code> , the full PDL entry, including PDL members, is logged in the "from" entry in the <code>MODIFY</code> record. Otherwise, only the X.400 attributes in the PDL entry is logged to save disk space and improve firectory synchronization performance.

Table 9: Directory Synchronization Options and their Descriptions

Option	Description
DS_ADD_UPDATE_LOCAL_ENTRY=TRUE	<p>Determines the behavior of a Directory synchronization search. If the ADD operation returns a matching entry and DS_ADD_UPDATE_LOCAL_ENTRY is set to TRUE, a MODIFY operation is performed.</p> <p>During the MODIFY operation, the matched entry in the importing Directory is modified with the corresponding entry in the exporting Directory.</p> <p>If the MODIFY operation is successful, a MODIFY record is added to the Directory change log.</p> <p>Using this option can modify entries in the importing Directory which are genuinely different from those in the exporting Directory.</p>
DS_CUST_IMP_ADD_UPDATE_LOCAL_ENTRY=TRUE	<p>Determines the result of a Directory synchronization search. During Directory synchronization, an entry in the exporting Directory will match an entry in the importing Directory if the search attributes are identical, even though other attributes of the entry may be different. In this case, if this option is set to TRUE, the entry in the importing directory is overwritten by the corresponding entry in the exporting directory.</p> <p>If the option is set to FALSE, the entry in the exporting Directory is treated as a duplicate of the corresponding entry in the importing directory, and so is not exported.</p> <p>Using this option can have the effect of modifying entries in the importing Directory which are genuinely different from those in the exporting Directory.</p>
DS_CUST_MSGQ_TIMEOUT= <i>seconds</i>	<p>The number of seconds the DS Server will wait on an empty message queue before checking if any of its timers have expired. Defaults to 1800 seconds. Use in conjunction with DS_CUST_PERIOD_TIMER_MINUTES when testing Directory synchronization.</p>
DS_CUST_PERIOD_TIMER_MINUTES=TRUE	<p>If set, the value of the period timer for Scalix-to-Scalix Directory synchronization is in minutes. The default is for the value of the period timer to be in hours. Use in conjunction with DS_CUST_SEND_REQ_NOW when testing Directory synchronization.</p>
DS_CUST_SEND_REQ_NOW=TRUE	<p>If set, Request messages are sent immediately the first timer check is performed after a restart of the DS Server if the period timer value (default 24 hours) is greater than the time from the current time to the start time. The default is to wait for the period timer to expire. Use in conjunction with DS_CUST_PERIOD_TIMER_MINUTES when testing Directory synchronization.</p>

Table 9: Directory Synchronization Options and their Descriptions

Option	Description
DS_SEND_SOURCE_LID=TRUE	<p>When set to TRUE, a unique identifier is propagated with each entry in each transaction during Directory synchronization. The value of the identifier is that of the LOCAL-UNIQUE-ID attribute of the entry in the exporting Directory.</p> <p>Setting this option to FALSE prevents this identifier being sent during Directory synchronization.</p> <p>The default setting for this option is TRUE.</p>

Table 10: Directory Synchronization Options and their Descriptions

Option	Description
DR_NO_MOD_STRIP_PDL=TRUE	<p>Determines the amount of information included for an entry in the Directory change log when changes are made to PDL members within the exporting Directory using the following PDL commands:</p> <ul style="list-style-type: none"> • omaddpdln • omdelpdln • ommodpdln <p>When set to TRUE, the full PDL entry, including PDL members, is logged in the "from" entry in the MODIFY record. Otherwise, only the X.400 attributes in the PDL entry is logged to save disk space and improve directory synchronization performance.</p>
DS_ADD_UPDATE_LOCAL_ENTRY=TRUE	<p>Determines the behavior of a Directory synchronization search. If the ADD operation returns a matching entry and DS_ADD_UPDATE_LOCAL_ENTRY is set to TRUE, a MODIFY operation is performed.</p> <p>During the MODIFY operation, the matched entry in the importing Directory is modified with the corresponding entry in the exporting Directory.</p> <p>If the MODIFY operation is successful, a MODIFY record is added to the Directory change log.</p> <p>Using this option can modify entries in the importing Directory which are genuinely different from those in the exporting Directory.</p>
DS_CUST_IMP_ADD_UPDATE_LOCAL_ENTRY=TRUE	<p>Determines the result of a Directory synchronization search. During Directory synchronization, an entry in the exporting Directory will match an entry in the importing Directory if the search attributes are identical, even though other attributes of the entry may be different.</p> <p>In this case, if this option is set to TRUE, the entry in the importing directory is overwritten by the corresponding entry in the exporting directory.</p> <p>If the option is set to FALSE, the entry in the exporting Directory is treated as a duplicate of the corresponding entry in the importing directory, and so is not exported.</p> <p>Using this option can have the effect of modifying entries in the importing Directory which are genuinely different from those in the exporting Directory.</p>

Table 10: Directory Synchronization Options and their Descriptions

Option	Description
DS_CUST_MSGQ_TIMEOUT= <i>seconds</i>	The number of seconds the DS Server will wait on an empty message queue before checking if any of its timers have expired. Defaults to 1800 seconds. Use in conjunction with DS_CUST_PERIOD_TIMER_MINUTES when testing Directory synchronization.
DS_CUST_PERIOD_TIMER_MINUTES=TRUE	If set, the value of the period timer for Scalix to Scalix Directory synchronization is in minutes. The default is for the value of the period timer to be in hours. Use in conjunction with DS_CUST_SEND_REQ_NOW when testing Directory synchronization.
DS_CUST_SEND_REQ_NOW=TRUE	If set, Request messages are sent immediately the first timer check is performed after a restart of the DS Server if the period timer value (default 24 hours) is greater than the time from the current time to the start time. The default is to wait for the period timer to expire. Use in conjunction with DS_CUST_PERIOD_TIMER_MINUTES when testing Directory synchronization.
DS_SEND_SOURCE_LID=TRUE	When set to TRUE, a unique identifier is propagated with each entry in each transaction during Directory synchronization. The value of the identifier is that of the LOCAL-UNIQUE-ID attribute of the entry in the exporting Directory. Setting this option to FALSE prevents this identifier being sent during Directory synchronization. The default setting for this option is TRUE.

IMAP Client Options

Table 11: IMAP Client Options and their Descriptions

Option	Description
IMAP_AUTOMATIC_MDN=FALSE	Determines whether the IMAP4 Server should generate Message Disposition Notification messages (MDNs) automatically. RFC2298 specifies that IMAP4 clients should generate MDNs where requested. However, some clients are unable to do this. Set this option to <code>TRUE</code> if you wish the IMAP4 Server to send MDNs automatically, without reference to the IMAP4 client.
IMAP_BB_FOLDER_PREFIX=#bb	Specifies the string that precedes all mailbox names for Public Folders. Use this option to enable users to distinguish easily between Public Folders and private folders. You can specify any ASCII character as a separator, although you are recommended not to choose 0-9, a-z, A-Z, +, comma, &, – (minus), tab, space, newline, %, or *. You must not use the same value for this option as for the <code>IMAP_FOLDER_PREFIX</code> option, or the <code>IMAP_FOLDER_SEPARATOR</code> or <code>IMAP_BB_FOLDER_SEPARATOR</code> options. You must enter the character itself, rather than the ASCII code for it.
IMAP_BB_FOLDER_SEPARATOR=/ /	Specifies the character used to separate public folder names in an IMAP mailbox specification. For example, with a separator of <code>/</code> and a Public Folder <code>beta</code> whose parent Public Folder is <code>alpha</code> , the corresponding IMAP mailbox name is <code>alpha/beta</code> . You can specify any ASCII character as a separator, although you are recommended not to choose 0-9, a-z, A-Z, +, comma, &, – (minus), tab, space, newline, %, or *. You must enter the character itself, rather than the ASCII code for it. You can need to use a Public Folder separator that is different from the default (<code>/</code>) if you wish to use this character in Public Folder names. For example, if there is an existing Public Folder called <code>Sales/Forecasts</code> , this Public Folder will not be seen through the IMAP Server since it could not be distinguished from a Public Folder called <code>Forecasts</code> within a Public Folder called <code>Sales</code> . There is a separate option to set the separator for private folders: see <code>IMAP_FOLDER_SEPARATOR</code> later in this section. See also the related options <code>IMAP_FOLDER_PREFIX</code> and <code>IMAP_BB_FOLDER_PREFIX</code> .

Table 11: IMAP Client Options and their Descriptions

Option	Description
IMAP_CAPABILITIES= <i>capabilities-list</i>	<p>Lists the capabilities that the IMAP Server advertises to the client. Each item in the list is separated by a space.</p> <p>The default list of capabilities is "IMAP4 IMAP4rev1 IDLE NAMESPACE AUTH=LOGIN".</p> <p>The capabilities that can be included in the list are:</p> <p>IMAP4: Support for the basic protocol as defined in RFC1730. Note that the IMAP4 commands defined in RFC1730 but absent in RFC2060 are still supported even if this capability is not advertised.</p> <p>IMAP4rev1: Support for the basic protocol as defined in RFC2060. The IMAP Server always advertises this capability.</p> <p>CHILDREN: Support for a means of indicating whether or not folders have child folders or not. This is not a standard extension.</p> <p>IDLE: Support for the IDLE extension as defined in RFC2177. This extension can provide a significant performance benefit for clients that can use it.</p> <p>LITERAL+: Support for non-synchronizing literals as defined in RFC2088. Use this capability with caution since it leaves the Server open to certain kinds of denial-of-service attacks.</p> <p>NAMESPACE: Support for the NAMESPACE command as defined in RFC2342. This command is used by certain clients to discover the namespace prefix for Public Folders so that these can be seen by the client user.</p>
IMAP_CONNECTION_LIMIT=0	<p>Specifies the maximum number of concurrent IMAP connections that the Server can support.</p> <p>If left at 0 (the default value), the IMAP4 Server will continue to accept all connections until machine resources are exhausted. This could adversely affect Scalix performance and eventually prevent other users from accessing the Scalix Server.</p>
IMAP_CONNRATE_LIMIT=0	<p>Specifies the maximum number of IMAP connection requests that the IMAP Server will accept in any one second.</p> <p>If left at 0 (the default), the IMAP Server will accept connection requests at a rate that is limited only by machine resources. This could adversely affect Scalix performance and eventually prevent other users from accessing the Scalix Server.</p> <p>If, for example, you set this value to 3, the IMAP Server is able to accept up to 180 connection requests per minute, and machine resources should still be sufficient to allow normal Scalix operation.</p>
IMAP_DELETE_SUBFOLDERS= FALSE	<p>Determines whether the IMAP4 Server permits the deletion of folders that contain subfolders.</p> <p>When this option is set to <code>FALSE</code> (the default), the IMAP4 Server will not allow a client to delete a folder that contains subfolders. This is in accordance with the IMAP4 protocol.</p> <p>However, some non-conforming clients will attempt to delete such folders. Set this option to <code>TRUE</code> if you wish to allow such attempts to succeed (and possibly enhance the usability of these clients).</p>

Table 11: IMAP Client Options and their Descriptions

Option	Description
IMAP_FOLDER_SEPARATOR=/	<p>Specifies the character used to separate private folder names in an IMAP mailbox specification. For example, with a separator of / and a folder named <code>beta</code> inside another folder named <code>alpha</code>, the corresponding IMAP mailbox name is <code>alpha/beta</code>.</p> <p>You can specify any ASCII character as a separator, although you are recommended not to choose 0-9, a-z, A-Z, +, comma, &, - (minus), tab, space, newline, %, or *. You must enter the character itself, rather than the ASCII code for it.</p> <p>You can need to use a folder separator that is different from the default (/) if you wish to use this character in folder names. For example, if an existing Scalix user has a folder called <code>Sales/Forecasts</code>, this folder will not be seen through the IMAP Server since it could not be distinguished from a folder called <code>Forecasts</code> within a folder called <code>Sales</code>.</p> <p>There is a separate option to set the separator for Public Folders: see <code>IMAP_BB_FOLDER_SEPARATOR</code> earlier in this section. See also the related options <code>IMAP_FOLDER_PREFIX</code> and <code>IMAP_BB_FOLDER_PREFIX</code>.</p>
IMAP_IDLE_TIMEOUT=30	<p>Specifies the number of minutes an IMAP connection can remain idle before the connection is closed by the IMAP Server.</p> <p>Specify a value of 0 to disable idle timeouts.</p> <p>Note that the IMAP protocol (RFC2060) specifies a minimum timeout of 30 minutes. Some clients can wait exactly 30 minutes between commands and so are liable to get logged out prematurely if this option is not set, or is set to its default value. For these clients, it is sometimes useful to set the idle timeout to 31 minutes.</p>
IMAP_IGNORE_SERVERNAME=FALSE	<p>Determines whether the IMAP Server uses the characters following the @ character in a username as the Server name for this user.</p> <p>When set to <code>FALSE</code> (the default), the name part of the username (up to and including the @ character) is stripped off and the remainder is used as the Server name to which the IMAP connection is relayed.</p> <p>Set this option to <code>TRUE</code> to prevent the IMAP connection being relayed to another Server.</p>
IMAP_LOGFILE=~/.tmp/imap.%h	<p>Specifies the name of the file to which IMAP events are logged, provided that logging is turned on using the <code>IMAP_LOGLEVEL</code> option.</p> <p>If the file you specify already exists, new events are appended to it.</p> <p>Note that at certain log levels log files can contain sensitive information, such as passwords.</p> <p>You can use the following tokens in the log file name:</p> <ul style="list-style-type: none"> <code>%p</code>: Expands to the PID of the IMAP Server process. One log file is created for each IMAP Server process. <code>%h</code>: Expands to the name of the client host. One log file is created for each client host that connects to the IMAP Server. <code>%u</code>: Expands to the Scalix UID. One log file is created for each Scalix user that connects to the IMAP Server.

Table 11: IMAP Client Options and their Descriptions

Option	Description
IMAP_LOGLEVEL=0	<p>Activates logging of IMAP commands and errors. The log file is specified by the <code>IMAP_LOGFILE</code> option.</p> <p>Set a value of 0 to disable logging, 1 to log basic commands/responses only, 2 to log unexpected UAL errors, and 8 to enable raw protocol logging. Note that, at log level 8, passwords are recorded in the log files. To avoid this, you can set a lower log level in the system-wide or per-client configuration file, and then set the log level to 8 in the user-specific configuration file. This log level will only take effect after authentication, and so passwords will not be recorded.</p> <p>If you require several kinds of logging information, add the numbers for the log levels you require.</p> <p>If the option <code>IMAP_UAL_TRACE_LEVEL</code> is not defined, then setting <code>IMAP_LOGLEVEL</code> to any value other than 0 enables UAL logging for the IMAP Server.</p>
IMAP_MAILSTORE_HOST= <i>hostname</i>	<p>Specifies the fully qualified domain name of the Scalix host to which the IMAP Server connects. Use this option when the IMAP4 Server does not reside on the same machine as the Scalix system that contains the relevant message store.</p>
IMAP_MDSENT_FLAG=\$MdnSent	<p>Determines the name of the flag that is set to indicate that a Message Disposition Notification (MDN) has been sent for this message. (See also the option <code>IMAP_AUTOMATIC_MDN</code>.)</p> <p>The name of this flag has not been standardized, and therefore different IMAP4 clients can use different flag names. Set this option to the name of the flag that your IMAP clients use.</p>
IMAP_MIN_SIZE_ESTIMATE=0	<p>Specifies if the client will compute message sizes, or will estimate them.</p> <p>Some clients report the message size when they list messages in the user's In Tray. To do this, they must render the message, which can be time-consuming, and cause a decrease in performance.</p> <p>To prevent the client from rendering messages above a certain size, specify this size in kilobytes. For example, to prevent the IMAP client from rendering all messages above 5 kilobytes, set this value to 5. Messages less than about 5 kilobytes are rendered and have their size reported accurately. Messages larger than about 5 kilobytes will have an estimate of their size reported.</p> <p>Note that some clients require the message size to be computed accurately. For these clients, you must set this option to 0 or leave it undefined.</p>

Table 11: IMAP Client Options and their Descriptions

Option	Description
IMAP_REMOTE_UAL_ENABLED=TRUE	<p>Specifies whether an IMAP client can use a remote UAL Server. Local connections have better performance.</p> <p>If set to <code>TRUE</code>, users can specify the name of a remote machine on which is running a UAL Server. The IMAP Server will then use this remote UAL Server.</p> <p>Users specify the use of a remote UAL Server by connecting as <i>username@hostname</i>, where <i>hostname</i> is the name of the remote machine to which they wish to connect.</p> <p>Set this option to <code>FALSE</code> if you wish to prevent users from connecting to a remote UAL Server.</p>
IMAP_SEARCH_TIMEOUT=0	<p>Specifies the number of seconds to wait before abandoning a search request.</p> <p>Specify a value of 0 to prevent search requests from timing out.</p>
IMAP_UAL_TRACE_LEVEL=0	<p>Activates tracing of IMAP Server information at the UAL Server. The trace files are placed in the <code>~/tmp</code> directory. If this directory cannot be found, they are placed in the <code>/tmp</code> directory. Possible values, and the corresponding file names, are shown in the following table.</p> <p><i>user-no</i> is the Scalix user number.</p> <p>If you require several different kinds of trace information, add the numbers for the levels you require and set the entry to the total.</p> <ul style="list-style-type: none"> • 0: No tracing. The default. • 1: Raw (unformatted) command/reply tracing (file name: <i>OMuser-noN.trc</i>). • 2: Symbolic command/reply tracing (file name: <i>OMuser-noC.trc</i>). • 4: Message Store file name mapping. No trace file. The subject of an item listed or displayed in the client is replaced by its corresponding Message Store file name. • 8: Full tracing of command/reply and file transfer data. This can be used to rerun a session (file name <i>OMuser-noU.log</i> and <i>OMuser-noU.fnnnn</i>). • 16: Raw (unformatted) command/reply tracing and file transfer data (<i>user-noN.trc</i>). <p>This option is similar to the <code>UAL_TRACE_LEVEL</code> option. However, the <code>UAL_TRACE_LEVEL</code> option is user-specific, and causes information on all UAL clients to be traced. The <code>IMAP_UAL_TRACE_LEVEL</code> option is IMAP-specific (that is, its trace files contain information on all users of a particular machine), but it traces only IMAP information.</p>
IMAP_PUBLIC_FOLDERS	<p>When set to <code>FALSE</code>, public folders are hidden in IMAP clients such as SWA and Thunderbird. This can be set per user, per domain or for the whole server. Note: This also affects SWA.</p>

Internationalization Options

Table 12: Internationalization Options and their Descriptions

Option	Description
UXO_MIME_TEXTFILE_CHARSETS=<cs1>,<cs2>,...	For internationalization of message subject and body, where <cs1> and <cs2> etc. is a comma-separated list of character sets that the text is converted to, with the system trying each one in order of left to right, to determine which provides the least degradation. If one provides no degradation, the search stops. Supports per domain text charset conversion, ahead of the steering file settings.
UXO_MIME_ADDRESS_CHARSETS=<cs1>,<cs2>,...	For internationalization of message headers (From, To, cc, bcc), where <cs1> and <cs2> etc. is a comma-separated list of character sets that the text is converted to, with the system trying each one in order of left to right, to determine which provides the least degradation. If one provides no degradation, the search stops. Supports per domain text charset conversion, ahead of the steering file settings.

Internet Addressing Options

Table 13: Internet Addressing Options and their Descriptions

Option	Description
INET_AUTOGEN_IA_ON_MODIFY=FALSE	Determines whether the commands <code>ommodmn</code> , <code>ommodu</code> , <code>ommod-pdl</code> , <code>ommodent</code> , and <code>omldapmodify</code> generate Internet addresses automatically, when automatic Internet address mapping is in operation. The default is <code>FALSE</code> .
INET_DISPLAY_IA_COMMENTS=TRUE	Determines whether the POP3 and IMAP4 interfaces display the comment, or display name, part of an Internet address in a message. The default is <code>TRUE</code> . Set the option to <code>FALSE</code> to prevent the comment part of the Internet address from being displayed.
INET_INLINE_FILE_MAX_SIZE= <i>bytes</i>	Determines which body parts of MIME messages generated by the Internet Mail Gateway are inline and which are attachments. Body parts whose size is greater than the value of this option are attachments, while other body parts will be inline. Set a value of 0 to cause all body parts to be inline. Set a value of -1 to cause all body parts to be attachments.

Table 13: Internet Addressing Options and their Descriptions

Option	Description
INET_INLINE_FNAME_ALLOWED=FALSE	<p>Determines whether MIME messages generated by the Internet Gateway or prepared for browsing by POP3 or IMAP4 clients can have <code>filename=</code> in inline body part <code>Content-Disposition</code> lines.</p> <p>If the option is set to <code>FALSE</code> (the default), Inline body parts cannot have <code>filename=</code> in the <code>Content-Disposition</code> line even if a candidate filename exists.</p> <p>Set this option to <code>TRUE</code> to allow inline body parts to have <code>filename=</code> in the <code>Content-Disposition</code> line, if a candidate filename has been selected.</p>
INET_NO_IA_IN_ORN=FALSE	<p>Determines whether the incoming Internet Mail Gateway saves the Internet address of the sender, each recipient and DL member in the Scalix message.</p> <p>When set to <code>FALSE</code> (the default), the addresses are saved in the message.</p> <p>Note that this option applies to the names of Internet mail users and not to the names of Scalix users.</p>
INET_NO_IA_COMMENTS=FALSE	<p>Determines whether comments present in Internet addresses are included without alteration in outgoing messages. The default is <code>FALSE</code>, causing such comments to be included.</p>
INET_USE_AUTO_IAM=TRUE	<p>Specifies whether Internet addresses are automatically created when configuring users, and -mapped at the Internet Mail Gateway and the POP3 and IMAP interfaces.</p>

Internet Mail Gateway Options

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
BRW_INLINE_PARTS=1	Specifies the number of parts in a multi-part item that are marked as inline if there is no other information about the part. The default is 1.
BRW_ITEMSUB_IS_FNAME=TRUE	For those items which do not have an Original Filename specified, determines whether the Subject of the item is used to generate a filename for the item. Leave this option at its default (TRUE) if you wish to use the item Subject as the filename (provided the value of BRW_T61_ITEMSUB_IS_FNAME has not caused a filename to be generated). Encoding is determined by BRW_MIME_FNAME_ENCODING. Note that if BRW_MIME_FNAME_ENCODING=D then this takes precedence over BRW_ITEMSUB_IS_FNAME. Set the option to FALSE to prevent the Subject from being used to generate the filename (although a filename can still be generated using one of the other fields).
BRW_MIME_EXPLICIT_ASCII=FALSE	Specifies the character set to be US-ASCII for plain text content types. The MIME standards specify the default character set to be US-ASCII for text/plain content types, and you should therefore leave this at its default (FALSE) unless your client cannot display such content unless the character set is explicitly defined.
BRW_MIME_FNAME_ENCODING=Q	Specifies the method for encoding MIME names and filenames used in the POP3 and IMAP4 interfaces. The possible values are: D: forces outgoing non-text filename to meet DOS filename conventions N: no encoding Q: quoted-printable encoding; this is the default B: base64 encoding
BRW_MIME_OMIT_DEF_CTENC_HDR=F or N	If set to T for TRUE or Y for YES, the Content-Encoding header is omitted if it is the default 7 bit. The default is F or N.
BRW_MIME_SPACE_OK_IN_FNAME=TRUE	Specifies that spaces are allowed in filenames based on the T.61 subject of a body part.
BRW_MIME_SUBJ_BENC_NONASCII=F	When BRW_MIME_SUBJECT_ENCODING is set to B for base64 encoding, you can set this option (using either T for TRUE or Y for YES), to encode only non-ASCII characters in MIME subjects using base64. The default is F or N.
BRW_MIME_SUBJ_NO_SPACE_SEPS=FALSE	If set (using TRUE or YES), a space separator between encoded and non-encoded data is not generated. This option can only be set when BRW_MIME_SUBJECT_ENCODING=B and BRW_MIME_SUBJ_BENC_NONASCII=T. The default is FALSE or NO. Note that setting this option generates messages in a form that is not strictly compatible with RFC1522.

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
BRW_MIME_SUBJECT_CHARSET=NULL	Specifies the POP3 and IMAP4 MIME subject character set when it is different from the body part text. This option is for RFC 1557, which uses ISO-2022-KR in the body part text and EUC-KR (equivalent to KSC5601) in the subject.
BRW_MIME_SUBJECT_ENCODING=Q, B or N	Specifies the method for encoding the message subjects of MIME messages retrieved from the POP3 and IMAP4 Servers. The methods available are: D: forces outgoing non-text filename to meet DOS filename conventions N: no encoding Q: quoted-printable encoding; this is the default B: base64 encoding
BRW_MIME_SUBJECT_FOLDING=F, or N	If set to T or Y, folds subject headers according to RFC 1522 rules (at 76 bytes after encoding). Multibyte characters are sometimes folded at slightly less, to avoid splitting characters and to handle escape sequences correctly. The default is F or N.
BRW_MIME_TEXTFILE_ENCODING=Q, B, N, or ?	Specifies the method for encoding the message texts of MIME messages retrieved from the POP3 and IMAP4 Servers. The methods available are: ?: use the relevant mapping in the mime.types file (the default) N: no encoding Q: quoted-printable encoding; this is the default B: base64 encoding
BRW_NAME_MAPPING=FALSE	If set, the originator's name and address will be mapped to the keyed INTERNET-ADDR attribute (number 167) in the Directory entry for that user. The Directory entry must contain the user name and domain name in the format expected by Sendmail. Routing set up within Scalix and Sendmail must correspond to the addresses used in mappings. Note that mappings occur only when there is an exact match between the name and address in the message and the Directory entry attribute INTERNET-ADDR. You can specify which Directory to use for name/address mappings using the UX_NAME_MAPPING_DIR option. You can specify a Directory entry attribute to use other than INTERNET-ADDR using the UX_NAME_MAPPING_ATTRIB option. The default is FALSE.
BRW_MIME_ADDRESS_CHARSETS	For internationalization of message headers (From, To, cc, bcc), where <cs1> and <cs2> etc. is a comma-separated list of character sets that the text is converted to, with the system trying each one in order of left to right, to determine which provides the least degradation. If one provides no degradation, the search stops. Supports per domain text charset conversion, ahead of the steering file settings.

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
BRW_MIME_TEXTFILE_CHARSET	For internationalization of message subject and body, where <cs1> and <cs2> etc. is a comma-separated list of character sets that the text is converted to, with the system trying each one in order of left to right, to determine which provides the least degradation. If one provides no degradation, the search stops. Supports per domain text charset conversion, ahead of the steering file settings.
BRW_NO_RETAIN_IF_CONVERTED=FALSE	Determines if a message containing an alternative filetype read through the POP3 or IMAP4 Server retains the original format file along with the converted plain text version. By default, this option is set to <code>FALSE</code> ; alternative filetypes are retained along with the converted plain text version. If this option is set to <code>TRUE</code> , the alternative file is discarded after it has been converted into a text file. The resultant MIME message is created using just the converted text file. <i>filetype</i> is a filetype (or a comma-separated list of filetypes) configured in Scalix to be discarded. For example, to discard an original RTF file after conversion to plain text, set this option to <code>BRW_NO_RETAIN_IF_CONVERTED=2130</code> .
BRW_T61_ITEMSUB_IS_FNAME=F	If set to <code>T</code> , the T61 item subject is used for the filename when browsing POP3 and IMAP4 mail messages. The encoding is determined by the setting of <code>BRW_MIME_FNAME_ENCODING</code> ; if <code>BRW_MIME_FNAME_ENCODING</code> is set to <code>D</code> , it takes precedence over this option. The default is <code>F</code> .
INET_USE_X400_ATTS_FOR_LOOKUP=TRUE	Determines whether name mapping, using Directory lookup, at the outgoing Internet Mail Gateway uses only X.400 attributes. When this option is <code>TRUE</code> (the default), the Outgoing Internet Mail Gateway will only use X.400 attributes when it performs Directory lookup to map OR addresses to Internet addresses.
MAX_MIME_BROWSERS=25	Specifies the maximum number of MIME browsers that the MIME Browser Controller can have in its pool. The default is 25. Specify a higher value to provide a faster response to IMAP4 and POP3 connection requests. Specify a lower value to conserve system resources.
MIME_CACHE_TARGET_SIZE=1	Specifies the target size in megabytes of the mime cache (<code>~Scalix/temp/mime_cache</code>). The cache can grow larger than this if everything in the cache is being used, but unused items are deleted to keep the size under control. The default is 1.
MIN_MIME_BROWSERS=0	Specifies the minimum number of MIME browsers that the MIME Browser Controller can have in its pool. The default is 0. When the Mime Browser Controller starts, it will start the number of browsers specified by this option.

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
UX_MIME_SUBJECT_CHARSET= <i>client_character_set</i>	Specifies the character set to be used for the incoming and outgoing MIME subject, when it is different from the body part text. <i>client-character-set</i> is the name of any client character set configured in Scalix; for example, KSC5601. This option is for RFC-1557, which uses ISO-2022-KR in the body part text and EUC-KR (equivalent to KSC5601) in the subject.
UX_NAME_MAPPING_ATTRIB= <i>attribute_tag</i>	Specifies a Directory entry attribute to use other than INTERNET-ADDR. <i>attribute_tag</i> is the internal (numeric) form of the Scalix attribute.
UX_NAME_MAPPING_DIR= <i>Directory_name</i>	If set, the Directory you specify here is used to map the originator or recipient's name and address to an Internet address specified in the INTERNET-ADDR attribute of the user's Directory entry, as the message passes through the Internet Mail Gateway. The Directory must be a shared Directory. If this option is not set, the default system Directory is used. See also the UX_NAME_MAPPING_DIR_PASSWD option.
UX_NAME_MAPPING_DIR_PASSWD= <i>password</i>	If you set the UX_NAME_MAPPING_DIR option, you can specify a Directory password using this option.
UX_NO_ROUTE_CHECK=TRUE	If set, the Internet Mail Gateway does not check O/R Addresses in the incoming messages. Normally, the Internet Mail Gateway checks if valid routes are configured for O/R Addresses in the ARPA heading information of each incoming message.
UX_PRE_5_20_COMPATIBILITY_MODE=TRUE	When set to TRUE (the default), this option causes Scalix to create a WINMAIL.DAT attachment for outgoing messages that contain MAPI properties, and to not decode WINMAIL.DAT attachments for incoming messages.
UX_USE_ARPA_SENDER=TRUE	If set, the incoming Internet Mail Gateway constructs the Scalix "From" address of incoming messages from the value of the Sender: token in the ARPA header rather than from the SMTP Mail From command.
UXI_AUTO_REPLY_BULK_MAIL=FALSE	Specifies whether the incoming Internet Mail Gateway should allow auto-replies to bulk mailing list messages. Bulk mailing list messages are those that contain one of the following lines in the ARPA header: Precedence: bulk Precedence: list Precedence: junk By default (this option set to FALSE), bulk mailing list messages are treated in a very similar way to auto-forwarded messages, and do not allow auto-replies. Set this option to TRUE to allow auto-replies to bulk mailing list messages.

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
UXI_DDATYPE_HPMEXT=TRUE	If set, then messages coming in through the Internet Mail Gateway that have their internet mail addresses copied into the DDA of the Scalix addresses will use a DDA with a type of HPMEXT1 rather than the default RFC-822.
UXI_DO_1327_SENDER_MAP=TRUE	When set, the ARPA Sender of an Internet message display as the Creator of the message. Any replies to this message will therefore be sent to the Sender. This is the behavior described in RFC 1327. Note that this can not be the required behavior; for example, it means that Replies to a mailing list can be directed to the Creator, rather than to the subscribers.
UXI_KEEP_ARPA_ADDRESS=TRUE	If set, the Internet Mail Gateway will preserve the ARPA address of an incoming message even if the header contains an ARPA-encoded Scalix address. Normally, the Internet Mail Gateway checks the ARPA heading information of each incoming message and, if there is an ARPA encoding of a Scalix address, the ARPA address is discarded and the Scalix address used instead.
UXI_KEEP_MIME_ARPA_HEADER=TRUE	If set, the Internet Mail Gateway will include the ARPA header of an incoming MIME message in the Scalix message.
UXI_KEEP_UUENC_ARPA_HEADER=TRUE	If set, the Internet Mail Gateway will include the ARPA header of an incoming UUENCODE message in the Scalix message.
UXI_MIME_CS_AUTODETECT	Scalix scans text MIME body parts which are marked as being in an ASCII or ISO8859_1 character set to check whether they have been incorrectly labeled and are actually another character set type that Scalix can recognize with a higher degree of certainty. This scan is enabled by default by setting UXI_MIME_CS_AUTODETECT=FALSE, but can be disabled to increase Internet Gateway performance.
UXI_NAME_MAPPING=TRUE	If set, the originator's name and address are mapped to the keyed INTERNET-ADDR attribute (number 167) in the Directory entry for that user. The Directory entry must contain the user name and domain name in the appropriate format. Routing set up within Scalix and Send-mail must correspond to the addresses used in mappings. Note that mappings occur only when there is an exact match between the name and address in the message and the Directory entry attribute INTERNET-ADDR. You can specify which Directory to use for name/address mappings using the UX_NAME_MAPPING_DIR option. You can specify a Directory entry attribute to use other than INTERNET-ADDR using the UX_NAME_MAPPING_ATTRIB option.
UXI_NO_CONVERT_REPORTS=FALSE	When set to FALSE, Internet acknowledgments and acknowledgment requests are converted to their Scalix equivalents. If you set this option to TRUE, they are not converted, but are passed into Scalix as Messages.

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
UXI_NO_INET_OBJFILES=FALSE	When set to <code>FALSE</code> , the contents of Internet headers are preserved in object files. Certain clients, such as the IMAP4 and POP3 clients, can make use of these object files. If set to <code>TRUE</code> , Scalix will not create these object files, and this header information will therefore be lost.
UXI_NO_UUDECODE_STRING= text string	If set, UUENCODEd parts of messages that are not of a recognized format, such as MIME or RFC1154, will not be decoded if the message body contains a text string that matches the one you specify here. If this option is not set, or no text string is supplied for it, such messages will have their UUENCODEd parts decoded into separate binary attachments.
UXI_PASSIVE_RECIPS_MAPI_ENABLED=FALSE	Determines whether passive recipients (that is, those recipients for which the Internet Mail Gateway does not have responsibility) appear to Outlook users with the "Send In RTF" flag set. When this option is set to <code>FALSE</code> (the default), then such recipients are assumed not to be MAPI-enabled, and the "Send In RTF" flag is not set. Set this option to <code>TRUE</code> if you wish to have the "Send In RTF" flag set for these recipients.
UXI_PRESERVE_MAPI_MSG_CLASS=FALSE	Specifies whether the MAPI message class of certain incoming messages is converted. To interoperate with Exchange, the Internet Mail Gateway must convert the MAPI message class of certain messages received from the Exchange Internet Mail connector for the Scalix MAPI Service Providers. The default is <code>FALSE</code> . See also the <code>UXO_PRESERVE_MAPI_MSG_CLASS</code> option.
UXI_SUPPRESS_ARPA_HEADER=TRUE	Suppresses, at a system level, the generation of the ARPA header for incoming messages from the Internet Mail Gateway. Note that MIME-encoded messages have ARPA headers suppressed by default. To enable MIME-encoded messages to have ARPA headers, you must set <code>UXI_KEEP_MIME_ARPA_HEADER=TRUE</code> .
UXI_TREAT_AS_MIME_SUBJECT=T or Y	If set (using either T for <code>TRUE</code> or Y for <code>YES</code>), incoming messages from the Internet Mail Gateway's UUENCODE or SHAR route (in other words, messages in which the ARPA headers did not contain a <code>Mime-Version: 1.0</code> tag) but which have MIME-conformant subjects, have their subjects decoded as if they came via the MIME route, and are subject to all other settings for MIME subjects.

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
UXI_UNIX_MAIL_CHARSET= <i>character_set</i>	Specifies the character set used in messages coming in through the Internet Mail Gateway. Only use this option to specify non-Latin character sets. For Latin character sets, Scalix assumes the same character set is being used as was used for outgoing messages; that is, the character set to which ISO8859/1 text was converted to in the <code>unix-out.str</code> or <code>mimeout.str</code> file. If the <code>UXI_UNIX_MAIL_CHARSET</code> option is set, the file <code>~/sys/unixin.str</code> or <code>mimein.str</code> can be used to specify conversions from this character set to a suitable interchange character set.
UXI_UUDECODE_ARPA_TOKEN= <i>string</i>	Specifies the token which, if present in the ARPA header of an incoming internet mail message, results in any UUENCODEd parts in the message being decoded. The default is to decode UUENCODEd parts in all messages. To prevent any messages containing UUENCODEd message parts from being decoded (except those that conform to RFC 1154), specify a null ("") string.
UXO_ADD_DELIM=TRUE	Specifies that a leading / is inserted in front of an O/R Address that is being used within internet mail. The / is inserted only if the O/R Address is format 2 (attribute format) and is enclosed in inverted commas. (This is done when attributes in the O/R Address contain characters that have a special meaning to internet mail.) Inserting the / ensures that Sendmail identifies the message as a message for Scalix.
UXO_CHECK_TYPES_OF_DDA= <i>DDA_type</i>	<p>Specifies the DDA types that are acceptable as valid internet addresses, to enable the UAL Client Interface and the Internet Mail Gateway to route the message to the correct destination for the recipient. <i>DDA_type</i> is one of the following:</p> <p>One or more valid DDA types, for example: RFC-822 HPMEXT1 HPMEXT2 HPMEXT3 HPMEXT4</p> <p>You can specify up to 10 DDA types; if you do specify more than one type, separate them with commas. For example: RFC-822,HPMEXT1,HPMEXT2,HPMEXT3,HPMEXT4</p> <p>FALSE No DDA type checking is performed. This behavior is as for Scalix systems before B.05.10.</p> <p>If your Scalix Directory contains DDAs with no type specified and they are not valid internet addresses, you are recommended to set this option to: UXO_CHECK_TYPES_OF_DDA=RFC-822,HPMEXT1,HPMEXT2,HPMEXT3,HPMEXT4</p> <p>If this option is not present, the default setting is: UXO_CHECK_TYPES_OF_DDA=,RFC-822,HPMEXT1,HPMEXT2,HPMEXT3,HPMEXT4</p> <p>The leading comma means that DDAs with no type specified are also acceptable as valid internet addresses.</p>

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
UXO_ITEMSUB_IS_FNAME=FALSE	<p>If set to TRUE, the item subject is used for the filename in outgoing MIME Internet Mail messages, if no Original filename is present and if UXO_T61_ITEMSUB_IS_FNAME has not caused a filename to be generated. The encoding is determined by the setting for UXO_MIME_FNAME_ENCODING; if UXO_MIME_FNAME_ENCODING is set to D, it takes precedence over this option.</p> <p>The default is FALSE; this setting prevents a filename from being generated from the item subject for the Content-Disposition header.</p> <p>See also:</p> <p>UXO_MIME_FNAME_ENCODING UXO_T61_ITEMSUB_IS_FNAME INET_INLINE_FNAME_ALLOWED</p>
UXO_MIME_FNAME_ENCODING=D	<p>Specifies how MIME names and filenames are encoded at the outgoing Internet Mail Gateway. The options are:</p> <p>D: forces outgoing non-text filename to meet DOS filename conventions</p> <p>N: no encoding</p> <p>Q: quoted-printable encoding; this is the default</p> <p>B: base64 encoding</p>
UXO_MIME_OMIT_DEF_CTENC_HDR=T or Y	<p>If set (using either T for TRUE or Y for YES), the Content-Encoding header is omitted if it is the default 7 bit.</p>
UXO_MIME_SPACE_OK_IN_FNAME=TRUE	<p>Specifies that spaces are allowed in filenames based on the T.61 subject of a body part.</p>
UXO_MIME_SUBJ_NO_SPACE_SEPS=TRUE	<p>If set (using TRUE or YES), a space separator between encoded and non-encoded data is not generated. This option can only be set when UXO_MIME_SUBJECT_ENCODING=B and UXO_MIME_SUBJ_BENC_NONASCII=T.</p> <p>Note that setting this option generates messages in a form that is not strictly compatible with RFC1522.</p>
UXO_MIME_SUBJECT_BENC_NONASCII=T	<p>When UXO_MIME_SUBJECT_ENCODING is set to B for base64 encoding, you can set this option (using either T for TRUE or Y for YES), to encode only non-ASCII characters in MIME subjects using base64.</p>
UXO_MIME_SUBJECT_ENCODING=Q, B or N	<p>Specifies the method for encoding MIME subjects of outgoing messages. The methods available are:</p> <p>N: no encoding</p> <p>Q: quoted-printable encoding; this is the default</p> <p>B: base64 encoding</p> <p>If UXO_MIME_FNAME_ENCODING is not set, this option is used for filename encoding, as well.</p>
UXO_MIME_SUBJECT_FOLDING=T, or Y	<p>Folds subject headers according to RFC 1522 rules (at 76 bytes after encoding). Multibyte characters are sometimes folded at slightly less, to avoid splitting characters and to handle escape sequences correctly. Enter either T for TRUE or Y for YES to set this option.</p>

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
UXO_MIME_TEXTFILE_ENCODING=Q, B or N	Specifies the method for encoding message text of outgoing MIME messages. The methods available are: N: no encoding Q: quoted-printable encoding; this is the default B: base64 encoding
UXO_NAME_MAPPING=TRUE	If set, the recipient's name and address are mapped to the keyed <code>INTERNET-ADDR</code> attribute (number 167) in the Directory entry for that user. The Directory entry must contain the user name and domain name in the format expected by Sendmail. Routing set up within Scalix and Sendmail must correspond to the addresses used in mappings. If the recipient already has an Internet mail name and address configured in the DDA fields in the message or in the entry retrieved from the directory, this is used in preference to the <code>INTERNET-ADDR</code> attribute value. You can specify which Directory to use for name/address mappings using the <code>UX_NAME_MAPPING_DIR</code> option. You can specify a Directory entry attribute to use other than <code>INTERNET-ADDR</code> using the <code>UX_NAME_MAPPING_ATTRIB</code> option.
UXO_NO_RETAIN_IF_CONVERTED=FALSE or <i>filetype</i>	Determines if a message containing an alternative filetype sent through the Internet Mail Gateway retains the original format file along with the converted plain text version. By default, this option is set to <code>FALSE</code> ; alternative filetypes are retained along with the converted plain text version. If this option is set to <code>TRUE</code> , the alternative file is discarded after it has been converted into a text file. The resultant MIME message is created using just the converted text file. <i>filetype</i> is a filetype (or a comma-separated list of filetypes) configured in Scalix to be discarded. For example, to discard an original RTF file after conversion to plain text, this option is set to <code>UXO_NO_RETAIN_IF_CONVERTED=2130</code> .
UXO_PRESERVE_MAPI_MSG_CLASS=FALSE	Specifies whether the MAPI message class of certain outgoing messages is converted. To interoperate with Exchange, the Internet Mail Gateway must convert the MAPI message class of certain messages destined for the Exchange Internet Mail connector for the Scalix MAPI Service Providers. The default is <code>FALSE</code> . See also the <code>UXI_PRESERVE_MAPI_MSG_CLASS</code> option.
UXO_SHAR_ARGS= <i>arguments</i>	Specifies the arguments used by the <code>shar</code> program when it is started by the Internet Mail Gateway. The default arguments are <code>-bc</code> . See also <code>UXO_SHAR_COMMAND</code> .
UXO_SHAR_COMMAND= <i>command</i>	Specifies the program used by the Internet Mail Gateway to create a shell archive package. The default is <code>shar</code> . See also <code>UXO_SHAR_ARGS</code> .

Table 14: Internet Mail Gateway Options and their Descriptions

Option	Description
UXO_T61_ITEMSUB_IS_FNAME=T	If set, the T61 item subject is used for the filename in outgoing MIME Internet Mail messages. The encoding is determined by the setting for UXO_MIME_FNAME_ENCODING; if UXO_MIME_FNAME_ENCODING is set to D, it takes precedence over this option.
UXO_TREAT_AS_MIME_SUBJECT=T or Y	If set (using either T for TRUE or Y for YES), messages going out via the Internet Mail Gateway's UUENCODE or SHAR route that have MIME-conformant subjects will have their subjects encoded as if going out via the MIME route, and are subject to all other settings for MIME subjects.
UXO_USE_SENDER_DDA=TRUE	Specifies whether a Domain Defined Attribute (DDA) is used directly when mapping the sender address in the outgoing Internet Mail Gateway. If set to TRUE, the Internet Mail Gateway maps DDAs in the sender address in the same way as it does for DDAs in the recipient address; that is, any internet address specified in the DDA is used directly. If this option is not present or is set to FALSE, the Internet Mail Gateway does not use the DDA directly when mapping the Internet address of the sender.

Item Structure Server Options

Table 15: Item Structure Server Options and their Descriptions

Option	Description
ISL_DISABLE_LOGGING=TRUE	Disables logging by the Item Structure Server, of structural changes made to the Message Store. This option is set when the Item Structure Database is not required, in order to save disk space. The option takes precedence over ISL_LOG_IF_OFF=TRUE.
ISL_LOG_IF_OFF=TRUE	Enables logging of structural changes made to the Message Store, when the Item Structure Server daemon is not running. Logging is performed directly to the Item Structure Server log files, which reduces performance. The option ISL_DISABLE_LOGGING=TRUE takes precedence over this option.

Local Delivery Service Options

Table 16: Local Delivery Service Options and their Descriptions

Option	Description
LD_ADD_ACKS_AS_TO	If set, the address returned in an acknowledgment that cannot be matched in the original distribution list of the message in the Outbox is added to the distribution list as a "To" record rather than as a "New Recipient" record.
LD_AUTOREPLY_CHECK_ON=TRUE	Each time a user configures auto-reply, Scalix creates a text file in the user's /g directory under ~/user/ which contains a list of addresses to which automatic replies have been sent since the current auto-reply session was created. With this option set, Local Delivery checks users' address list files against the address in each received message's transaction file. If a match is found, an automatic reply is not generated. This prevents more than one automatic reply from being generated for each unique sender address.
LD_AUTOREPLY_EXPIRY_TIME= <i>no_of_days</i>	Specifies the number of days an address can be present in the auto-reply address list file before it is removed.
LD_MAX_NEST_LEVEL= <i>depth</i>	Specifies the maximum level of nesting allowed in a message before further nested parts are flattened by the Local Delivery Service. A value of 0 means that all delivered messages are flattened. See SR_MAX_NEST_LEVEL for more information.
LD_READ_ACK_ON_AUTOPRINT=FALSE	If set, when a message is automatically printed, no "read" acknowledgment is returned to the originator. The default is to return a "read" acknowledgment when a message is automatically printed.
LD_TRACE_DISP_ACT=SHOW_ADMIN	Shows trace information on all messages received by anyone with Scalix administration permissions.

LDAP Server Options

Table 17: LDAP Server Options and their Descriptions

Option	Description
LDAP_MB_CN_IS_GS_IN_FILTER=FALSE	<p>This option only has effect when the LDAP session is multibyte, and you have not created an explicit Scalix attribute <code>COMMONNAME</code>. When the LDAP client sends a search filter that contains the LDAP <code>COMMONNAME</code> attribute, the LDAP Server uses a built-in parsing method to determine how to convert this to Scalix attributes. By default, it assumes that the <code>COMMONNAME</code> attribute contains the Scalix attributes <code>SURNAME</code>, <code>GIVENNAME</code> in that order. If your LDAP clients construct <code>COMMONNAMES</code> in the reverse order, you must set this option to <code>TRUE</code>. For example, if your LDAP client uses "Japanese-surname Japanese-givenname" as the <code>COMMONNAME</code>, leave this option at its default. The LDAP Server will correctly interpret this as <code>GIVENNAME=Japanese-givenname</code> and <code>SURNAME=Japanese-surname</code>. However, if your LDAP client uses Japanese-givenname Japanese-surname as the <code>COMMONNAME</code>, the LDAP Server will interpret this as <code>GIVENNAME=Japanese-surname</code> and <code>SURNAME=Japanese-givenname</code> unless you set this option to <code>TRUE</code>.</p>
LDAP_MB_CN_IS_GS_IN_SYNTH_OUT=FALSE	<p>This option only has effect when the LDAP session is multibyte, and you have not created an explicit Scalix attribute <code>COMMONNAME</code>. When the LDAP Server synthesizes the LDAP <code>COMMONNAME</code> attribute using the Scalix attributes <code>SURNAME</code> and <code>GIVENNAME</code>, it puts them in the order <code>SURNAME</code>, <code>GIVENNAME</code>. If you want <code>COMMONNAMES</code> to appear in the reverse order, you must set this option to <code>TRUE</code>. For example, when this option is left at its default value, the name "Japanese-givenname Japanese-surname" is returned from the LDAP Server as Japanese-surname Japanese-givenname. However, if you set this option to <code>TRUE</code>, the name is returned as Japanese-givenname Japanese-surname.</p>
LDAP_SEQUENTIAL_SEARCH=""	<p>If you do not set this option, the LDAP Server will not, in general, issue sequential searches of the Scalix directories. Instead, it will search using the indexes of keyed attributes, to keep search time to a minimum. However, it will issue sequential searches under certain circumstances, such as when the <code>DA_IGNORE_INDEXES</code> option is set to <code>TRUE</code>. Set this option to <code>TRUE</code> if you want the LDAP Server to issue sequential searches. This enables you to search for attributes that are not keyed, but searches could take a long time. Set this option to <code>FALSE</code> to prevent the LDAP Server from ever issuing sequential searches. This will keep search times to a minimum, but can prevent the LDAP Server from finding all entries that match a given filter.</p>

Table 17: LDAP Server Options and their Descriptions

Option	Description
OMLDAP_REMOVE_LEADING_WILDCARDS=TRUE	<p>If present and set to <code>TRUE</code>, leading wildcard characters (*) are stripped from substring filters when the LDAP Server searches a Scalix Directory for entries that match criteria specified by a search filter. This option causes filters of the form "(cn= *<i>name</i>*)" to be converted to the form "(cn=<i>name</i>*)". That is, the LDAP Server matches the filter "(cn=<i>name</i>*)" to all entries in the underlying Scalix Directory whose <code>SURNAME</code> or <code>COMMON-NAME</code> attributes start with <i>name</i>. This causes fewer system resources to be used when searching.</p> <p>If not present or set to <code>FALSE</code>, the leading wildcards are not stripped, and the LDAP Server searches for all <code>SURNAME</code> or <code>COMMON-NAME</code> attributes containing <i>name</i>.</p> <p>For example, if <code>OMLDAP_REMOVE_LEADING_WILDCARDS</code> is set to <code>TRUE</code>, "(cn=Marion Brand*)" would be matched to all entries whose <code>SURNAME</code> starts with "Brand" or whose <code>COMMON-NAME</code> starts with "Marion Brand". If this option is set to <code>FALSE</code>, "Ann-Marion Brandson" would be considered a match.</p>

Non-Delivery Notification Options

Table 18: Non-Delivery Notification Options and their Descriptions

Option	Description
NDN_EM_SERIOUS_ONLY=TRUE	Sends Non-Delivery Reports for serious errors to the Error Manager only. Sends Non-Delivery Reports for simple addressing problems to the originator only. If this option is not set, by default Non-Delivery Reports for simple addressing problems are sent to both the Error Manager and the originator.
NDN_NO_ALTERNATES=TRUE	If this option is set, alternate names are not placed in a Non-Delivery Notification if the original message contained an ambiguous O/R Name.

Notification Server Options

Table 19: Notification Server Options and their Descriptions

Option	Description
NS_INITIAL_MEM= <i>bytes</i>	Specifies the initial memory size of the Notification Server. Use this option to increase the initial memory size from 65536 (the default). This value is suitable for up to approximately 1200 configured and active users. You might want to increase this value if a larger number of users are configured such that, just after startup, the shared memory segment is repeatedly enlarged.

Offline Folder Synchronization Options (Outlook Clients)

Table 20: Offline Folder Synchronization Options and their Descriptions

Option	Description
OFS_LOG_AGE_LIMIT= <i>days</i>	When the age of a change log entry exceeds this value, it can be deleted when the change log file is compacted. Use this option to prevent change logs from growing too large when logging is activated but synchronization is not taking place for any reason. Take care when you set this value, since removal of any valid entries will cause the entire folder to be resynchronized. A value you set in <code>general.cfg</code> can be overridden on a per-user basis by setting a lower value for this option in the relevant user-specific configuration file.

Table 20: Offline Folder Synchronization Options and their Descriptions

Option	Description
OFS_LOG_SIZE_LIMIT= <i>kilobytes</i>	<p>Specifies, in kilobytes, the maximum size of the folder synchronization change log. Set a value between 20 and 10,000 KB. The default is 100 KB.</p> <p>When the size of a change log exceeds this value, the older entries can be deleted when the change log file is compacted.</p> <p>Use this option to prevent change logs from growing too large when logging is activated but synchronization is not taking place for any reason. Take care when you set this value, since removal of any valid entries will cause the entire folder to be resynchronized.</p> <p>A value you set in <code>general.cfg</code> can be overridden on a per-user basis by setting a lower value for this option in the relevant user-specific configuration file.</p>
OFS_WORK_FILE_DIR= <i>temp-directory</i>	<p>Specifies the location of the temporary files created during folder synchronization.</p> <p>Normally (when this option is not set), these temporary files are stored in <code>~/temp</code>. Specify a value for this option to cause all these files to be written to a different location. This allows you to use a high-speed (and possibly low-recovery) file system (for example, a RAM disk) to store these temporary files.</p> <p>The directory you specify must have:</p> <ul style="list-style-type: none"> permissions 771 a group of <code>scalix</code> an owner of <code>scalix</code> a path length of 225 characters or less <p>For example, to create a directory called <code>temp-ofs</code>, enter the following commands:</p> <pre>mkdir \$(omrealpath '~/temp-ofs') chown scalix:scalix \$(omrealpath '~/temp-ofs') chmod 771 \$(omrealpath '~/temp-ofs')</pre>

Omscan Options

Table 21: Omscan Options and their Descriptions

Option	Description
GS_DONT_SPLIT_FC=FALSE	When set to <code>TRUE</code> , <code>omscan</code> reports a single value for the size of a user's filing cabinet and waste basket combined, instead of separate figures for the filing cabinet and the waste basket.
SCN_KEEP_DATA_ORPHANS=FALSE	If set, files reported as orphans by <code>omscan</code> are not moved to the directory <code>~/orphans</code> but are deleted.
SCN_ORPHAN_DELETE_DELAY= <i>number_of_days</i>	Sets the age limit for orphan files before they are reported by <code>omscan</code> and moved to the directory <code>~/orphans</code> . The default is 1 day.
SCN_PREV_ORPHAN_DELETE_ DELAY= <i>number_of_days</i>	The number of days before a file in the directory <code>~/orphans</code> is deleted by the next run of <code>omscan -d</code> . The default is 30 days.
SCN_TEMP_DELETE_DELAY= <i>number_of_days</i>	Sets the age limit for files in <code>~/temp</code> before they are deleted by <code>omscan</code> . The default is 7 days.
SCN_TMP_DELETE_DELAY= <i>number_of_days</i>	Sets the age limit for files in <code>~/tmp</code> before they are deleted by <code>omscan</code> . the default is 7 days.

POP Server Options

Table 22: POP Server Options and their Descriptions

Option	Description
POP3_IGNORE_SERVERNAME= FALSE	Determines whether the POP3 Server uses the characters following the @ character in a username as the Server name for this user. When set to <code>FALSE</code> (the default), the name part of the username (up to and including the @ character) is stripped off and the remainder is used as the Server name to which the POP3 connection is relayed. Set this option to <code>TRUE</code> to prevent the POP3 connection being relayed to another Server.
POP3_MAILSTORE_HOST= <i>hostname</i>	Specifies the fully qualified domain name of the Scalix host to which the POP Server connects; for example, <code>omsrv1.acme.com-company.com</code> . Use this option when the POP Server does not reside on the same machine as the Scalix system that contains the relevant message store.
POP3_MAX_THREADS= <i>integer</i>	Restricts the number of threads that a single <code>pop3.server</code> process will use. By default, the value used is the maximum allowed from system resources (including system thread limits as well as limits in available file descriptors). Specify a value here to limit the number of threads used to a value less than the default. If you specify a value higher than the default, it will have no effect.

Table 22: POP Server Options and their Descriptions

Option	Description
POP3_RECORD_EMPTY_SIGNON=FALSE	<p>Determines whether POP3 user signons are recorded for empty In Trays.</p> <p>When a user with items in their In Tray signs on using POP3, the signon is recorded, and the last signon time for the user is updated. This causes the signon to take longer than if the signon is not recorded.</p> <p>When this option is <code>FALSE</code> (the default), if a user with an empty In Tray signs on using POP3, the signon is not recorded, and is faster. Set this option to <code>TRUE</code> to cause user signons to be recorded even when the user's In Tray is empty. This causes slower signons, but allows you to discover the last signon time for the user or to use <code>omstat -u</code> and <code>omstat -s</code> to report on POP3 users.</p>

Public Folder Server Options

Table 23: Public Folder Server Options and their Descriptions

Option	Description
BBS_ALLOW_LOCAL_SYNC=FALSE	Specifies whether public folder synchronization can occur between public folders on the same machine. By default, this option is <code>false</code> , and an attempt to synchronize public folders on the same machine results in a warning message in the log file.
BBS_CUST_CHECK_TIME= <i>minutes</i>	Sets the amount of time a public folder server spends in "import" mode before checking if any of its synchronization timers have expired. the default is 5 minutes. A public folder server toggles between import and export mode. It stays in import mode for the specified amount of time before checking the synchronization timers, and then if the timers have expired, it changes to export mode. when it has done all the exports, it changes back to import mode.
BBS_DELETE_MASTER_BY_SYNC=FALSE	Specifies that deletion of a Slave item from a Public Folder will not be propagated to its Master item. Set this option to <code>TRUE</code> if you want deletion of a Slave item to result in deletion of its Master item. (<code>BBS_PROPAGATE_SLAVE_DELETION</code> must also be set to <code>TRUE</code> for this to happen.)
BBS_DELLOG_RETENTION_PERIOD= <i>hours</i>	Specifies the Retention Period, in hours, for delete log files. The default is 24.
BBS_PROPAGATE_SLAVE_DELETION=FALSE	Specifies that deletion of a Slave item from a Public Folder will not be propagated to any other Slave copies or to the Master item. Set this option to <code>TRUE</code> if you want deletion of a Slave item to result in deletion of all other Slave copies of this item. (<code>BBS_DELETE_MASTER_BY_SYNC</code> must also be set to <code>TRUE</code> if you want deletion of a Slave item to result in deletion of its Master.)
BBS_SEND_OBJECT_FILES=TRUE	Specifies whether object files attached to messages or items are included when the message or item is exchanged during the Public Folder synchronization process. If set to <code>TRUE</code> , any object files attached to messages or items are included when the message is sent to another Public Folder Server. Note that this option affects only object files attached to messages or basic items on the Public Folder; object files attached directly to Public Folders are not synchronized.
BBS_SYNC_MESG_PRIORITY= <i>priority</i>	Specifies the priority with which Public Folder synchronization messages are sent. If Public Folder synchronization messages are slowing down other Scalix operations, you might want to use a lower priority for such messages. Alternatively, if Public Folder synchronization messages are taking too long, you might want to use a higher priority. Set the priority value to one of 0 (normal), 1 (nonurgent), or 2 (urgent).

Queue Options

Table 24: Queue Options and their Descriptions

Option	Description
QM_DONT_READ_MSG_AT_START=FALSE	When set to <code>TRUE</code> , this option specifies that all messages currently in the Scalix queues remain stationary when Scalix is restarted. New messages are processed normally. Set this option to <code>TRUE</code> if a major problem is encountered when the queue manager attempts to read queued messages from disk when Scalix starts up. This allows Scalix to start. Set the option back to <code>FALSE</code> and restart Scalix when the problem is resolved.
QM_FAILURE_DELAY_SEC= <i>seconds</i>	Specifies the number of seconds between messages being retried by the queue manager. When a message fails to be processed because the process that was handling the message died, the queue manager delays the message for this number of seconds before retrying the message. The default is 30. See also <code>QM_MAX_FAILURES</code> .
QM_MAX_FAILURES= <i>integer</i>	Sets the number of times the queue manager will attempt to retry a message before giving up and putting the message on the POISON queue. A failure occurs when the process which received the message dies before informing the queue manager that it has successfully dealt with the message. A value between 1 and 4 is normally suitable. A value of 0 will cause the queue manager to place the message on the POISON queue immediately on failure (that is, no retries). Too large a value can cause services to abort repeatedly. See also <code>QM_FAILURE_DELAY_SEC</code> .
Q_TIME_OUT= <i>seconds</i>	Sets the amount of time processes attempting to read a request from a Scalix queue will wait before timing-out. Setting this value low will ensure that processes remain swapped-in. The default is 30 seconds.

Recovery Folder Options

Table 25: Recovery Folder Options and their Descriptions

Option	Description
RECOVERY_FOLDER_EXPIRY_TIME=<time_period>	Where <time_period> is the amount of time that deleted items remain in the "Scalix Recovered Items" folder before being removed from the system. The default is seven days (7d). Example settings for this option are: 4d12h (4 days and 12 hours) or 240h (240 hours).

Search Server Options

Table 26: Search Server Options and their Descriptions

Option	Description
SE_DEFAULT_DELAY= <i>number_of_seconds</i>	Specifies the delay in seconds between runs of a persistent (that is, automatically repeated) background search of the Message Store. The Search Server checks for a specified delay in the following sequence: If a delay is specified in the search request transaction file, that delay is used. If no delay is specified in the search request transaction file (or it is set to zero), the delay specified in this SE_DEFAULT_DELAY option is used. If no delay is specified either in the search request transaction file or the SE_DEFAULT_DELAY option, a default delay of 300 seconds is used.
SE_MAX_CHILDREN= <i>max_number_of_child_processes</i>	Specifies the maximum number of child search processes that the Search Server can create. Each child process can execute only one search at a time. This option limits the number of background searches that can be performed simultaneously, but not the number of background searches that can be queued. The default number of child processes is 20.
SE_MAX_OVERDUE_TIME= <i>number_of_seconds</i>	Specifies the time after which an overdue persistent background search takes priority over one-off searches. The Search Server normally gives priority to one-off searches. When the time specified in this option is reached, the Search Server gives a persistent search priority over the one-off searches. This prevents persistent searches from being permanently blocked by a long queue of one-off searches. The default time is 300 seconds.

Service Router Options

Table 27: Service Router Options and their Descriptions

Option	Description
NDN_MSGFLAGS_OVERRIDE_RULE_ACTION_RETURN=FALSE	By default, messages created as a result of the RETURN action for a message delivery ruleset has the original message attached, even if the original message has a flag specifying that contents must not be included in non-delivery notifications. If you set NDN_MSGFLAGS_OVERRIDE_RULE_ACTION_RETURN to TRUE and the return of contents is not requested, the original message is not attached.

Table 27: Service Router Options and their Descriptions

Option	Description
RSL_BLANK_SUBJECT_BS_CHAR=FALSE	If the subject mapper is a shell script, a message subject containing a backslash () causes problems as the script interprets these as escape characters. The Service Router and Deferred Mail Manager replace any backslashes with an empty space. If the subject mapper is not a shell script, a backslash () can be preserved by setting this option to false.
SR_CONVERT_ISO7_FROM_UNIX=TRUE	If set, all textual body parts of messages coming in through the Internet Mail Gateway are converted to the ISO8859/1 character set, assuming the body parts contain IA5 characters with ISO-7 extensions. This option is only active if the SR_ISO7_HOST and SR_ISO7_language options are also present.
SR_CONVERT_ISO7_FROM_X400=TRUE	If set, all textual body parts of messages coming in through the X.400 Interface are converted to the ISO8859/1 character set, assuming the body parts contain IA5 characters with ISO-7 extensions. This option is only active if the SR_ISO7_HOST and SR_ISO7_language options are also present.
SR_CONVERT_ISO7_LANG=language	If set, activates the option SR_ISO7_language for messages passing through the Service Router. Only one instance of this option can be used and the language string must match a string in the ~/sys/LangMap file. See also UAL_ISO7_HOST.
SR_CONVERT_ONLY_IA5=TRUE	Used in conjunction with SR_CONVERT_ISO7_FROM_UNIX and SR_CONVERT_ISO7_FROM_X400. If set, only textual body parts with a character set of IA5 are assumed to contain ISO-7 extensions and be eligible for conversion as specified by the options SR_ISO7_HOST and SR_ISO7_language.
SR_DUMP_MSGS=BEFORE or AFTER	Puts a copy of each message processed by the Service Router on the Dump Server queue DUMP. If the value of SR_DUMP_MSGS is set to BEFORE, the message is copied before it is processed by the Service Router. If the value of SR_DUMP_MSGS is set to AFTER, the message is copied after it is processed by the Service Router.
SR_EXPAND_PDL=TRUE	Sets the Service Router to perform Public Distribution List (PDL) expansion. When this option is set, the active distribution list of any message that is addressed to a PDL is expanded. (Expanded means a PDL entry is replaced by the full list of the recipients that it represents.) When a PDL has been expanded, the message is re-submitted to the Service Router. The Service Router will expand PDLs that <i>cannot</i> be routed regardless of whether this option is set or not. SR_NO_ROUTE_PDL stops expansion when a message cannot be routed.
SR_FILTER_TYPES_OF_ATT=TRUE	Causes the Service Router to remove WINMAIL.DAT attachments, used by some clients.

Table 27: Service Router Options and their Descriptions

Option	Description
SR_ISO7_language= ISO7_characters	<p>Specifies how text using the ISO-7 extensions is converted to the ISO8859/1 character set by the Service Router. <i>ISO7_characters</i> is a list of ISO8859/1 characters to which the 12 special "ISO-7" characters are mapped. The IA5 characters used as special ISO-7 characters are: # \$ @ [] ^ { } ~</p> <p>The ISO8859/1 equivalents (<i>ISO7_characters</i>) must be specified in the same order. Ensure that the ISO8859/1 equivalents are entered into the file using the ISO8859/1 character set! The <i>language</i> must correspond to the language set in the SR_CONVERT_ISO7_LANG option and the SR_CONVERT_ISO7_LANG option must be present to activate this option.</p> <p>See also SR_CONVERT_ISO7_FROM_UNIX and SR_CONVERT_ISO7_FROM_X400. Also UAL_ISO7_language.</p>
SR_LD_BYPASS_LSERV=TRUE	<p>When this option is set to TRUE (the default setting) the Service Router can bypass the Local Delivery Service and route a local message directly to the queue of one of the following Scalix services:</p> <ul style="list-style-type: none"> Public Folder Server Directory Synchronization Server Error Manager Server Print Server Request Server <p>By minimizing traffic through the Local Delivery Service, this option can reduce the amount of time required for Directory synchronization, and increase the speed of other local traffic.</p> <p>If an ACL is associated with the Local Delivery Service, you must set this option to FALSE to prevent the ACL being bypassed when a message is being routed directly to one of the Scalix services listed above.</p>
SR_MAX_HOP_COUNT= <i>hop_count</i>	Specifies the number of hops that a message can make before it is assumed to be looping. The default is 100.
SR_MAX_NEST_LEVEL= <i>nest_level</i>	Specifies the maximum level of nesting allowed in a message before further nested parts are flattened by the Service Router. A value of zero means that all messages are flattened. See also LD_MAX_NEST_LEVEL.
SR_NO_ROUTE_PDL=TRUE	Stops Public Distribution List (PDL) expansion by the Service Router when a message cannot be routed. (Normally, if a message cannot be routed when there is a PDL within the message's distribution list, the Service Router will expand the PDL, or PDLs, and try to route the message again before returning a Non-Delivery Notification.)
SR_Q_TIME_OUT= <i>seconds</i>	Specifies the time, in seconds, between checking the Service Router queue for new messages and checking for deferred messages that are due for submittal to the Service Router. The default is 30 seconds.

Table 27: Service Router Options and their Descriptions

Option	Description
SR_RESOLVE_MASK= <i>number_of_ORname_fields</i>	<p>Specifies the directory attributes that are retained in the recipient address when the address is automatically resubmitted by the Service Router following a delivery failure. These attributes are specified as internal or language dependant attribute tags separated by forward slashes (/). If a message cannot be routed or delivered using the full recipient address, you can resubmit the recipient names with a less fully specified address by specify how many O/R Name fields are retained when the name is resubmitted.</p> <p>For example: S/G//Q: This will retain only the Personal Names attributes (Surname, GivenName, Initials and Generation) CN/OU1: This will retain only the Common Name and OrgUnit 1 attributes</p>
SR_ROUTE_X400_TO_OMX400_ <i>n</i> = <i>route_match</i>	<p>Allows messages to be rerouted from the x400 queue to the OMX400 queue for recipients whose address matches the specified values. The x400 queue is used for messages routed to non-Scalix X.400 systems; the OMX400 queue is used for messages routed to other Scalix systems. This option must be set on the Scalix system that contains the X.400 gateway where this rerouting is required.</p> <p><i>n</i> is a number between 1 and 8. This enables you to specify up to 8 unique instances of this option in the General Configuration File.</p> <p><i>route_match</i> specifies the route to be matched, using a series of O/R Address attributes and values, separated by forward slash characters (/). Attributes are specified as <i>TAG=value</i> pairings, where TAG is one of the following O/R Address attributes:</p> <p>TAG: O/R Address Attribute OU1: Organizational Unit Name 1 OU2: Organizational Unit Name 2 OU3: Organizational Unit Name 3 OU4: Organizational Unit Name 4 O: Organization Name P: Private Domain Name A: Administrative Domain Name C: Country OU1-TX: Teletex Organizational Unit Name 1 OU2-TX: Teletex Organizational Unit Name 2 OU3-TX: Teletex Organizational Unit Name 3 OU4-TX: Teletex Organizational Unit Name 4 O-TX: Teletex Organization Name</p> <p>The <i>value</i> specified is not case sensitive, and wildcard characters (*) can be used. If an attribute is not specified, it is treated as if it were fully wildcarded; that is, any value for that attribute is matched. No hierarchical rules are applied regarding which attributes can be specified and wildcarded.</p>
SR_SYNC_P2_WITH_P1=TRUE	Sets the Service Router to modify the original value of the O/R Address in the P2 distribution list as well as the P1

Table 27: Service Router Options and their Descriptions

Option	Description
SR_USEX500_DIR= TRUE or <i>X.500_Directory_Name</i>	Specifies that an X.500 Directory is used by the Service Router to resolve a DDN. If SR_USEX500_DIR is set to TRUE, the first X.500 Directory found is used. If the name of an X.500 Directory is specified, this Directory is used by the Service Router.
OMLIMIT_MIN_WARN_INTERVAL	<p>NOTIFY messages for the OMLIMIT-EXCEEDED sanction are only sent if the NOTIFY message has not been sent within the time specified by this setting.</p> <p>The default value for the OMLIMIT_MIN_WARN_INTERVAL option is one day (1d).</p> <p>Example settings for this option are:</p> <p>1h40m20s (1 hour 40 minutes and 20 seconds)</p> <p>2d40 (2 days and 40 seconds)</p> <p>6000 (6000 seconds/100 minutes)</p> <p>If the "omlimit -e u" sanction is enabled, the OMLIMIT_MIN_WARN_INTERVAL option also manages the interval during which omlimit-related messages are sent to a user.</p>

UAL Client Interface Options

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAKD_CONNRATE_LIMIT= <i>number_of_connections_per_second</i>	Specifies the maximum number of client connection processes that can be started per second. Specify a value here in order to limit the rate at which client connections are attempted. This can prevent delays caused by connection processes waiting for Server resources.
UAKD_LISTEN_Q_SIZE= <i>number_of_connections</i>	Specifies the number of TCP/IP Socket connections that can be queued to the UAL Server <code>listen.daemon</code> process during busy periods. This reduces the possibility of "UAL unable to connect" errors. <i>number_of_connections</i> can be between zero and the operating system limit. The default is 20 connections.
UAKD_NICE_VALUE=20<= <i>value</i> <=20	Increases or reduces the priority of TCP/IP Socket connections to the UAL Server <code>listen.daemon</code> process, over other activities performed by the Scalix Server. <i>value</i> is a number between -20 and 20, where negative values increase the priority of client signon. The default is -10.
UAKD_SERVER_PUSH_NOTIFS=TRUE	Determines whether the Server-push mechanism is enabled. The default is <code>TRUE</code> . The Server-push mechanism allows certain clients to receive notifications automatically, without having to poll for them. Set this option to <code>FALSE</code> to disable the Server-push mechanism, forcing clients to poll for notifications. Setting this option to <code>FALSE</code> can result in increased performance, but do not set the option to <code>FALSE</code> if: There are a significant number of Outlook clients in use. This will cause a large increase in network traffic. Any IMAP clients are in use. IMAP clients cannot receive notifications if this option is <code>FALSE</code> .
UAL_5_40_PERF_CHANGES=TRUE	Switches on or off the performance changes to the UAL Client Interface that were introduced in Scalix Release 5.40. The default is <code>TRUE</code> . Set this option to <code>FALSE</code> if you suspect that one or more of the performance enhancements is causing problems.
UAL_ALLOW_DISABLED_CLIENTS=FALSE	If this is set to <code>TRUE</code> , those clients specified in the <code>UAL_DISABLED_CLIENTS</code> option are permitted to sign on to the Server. Such sign-ons are logged with a Warning logging level. This can be used to find out which users are using a particular client so that they can be warned before the client is actually disabled.

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_FLDR_ACL_DEFAULT = <i>permissions</i>	<p>Specifies the permissions that are granted to the default user when a public folder is created. These permissions apply to each user unless the ACL has an entry that is more specific to that user.</p> <p>Set the value of <i>permissions</i> to a string of up to six characters, selected from the following: o (owner), c (contact), c (create), r (read), f (folder), E (edit all), e (edit own), D (delete all), d (delete own) v (visible).</p>
UAL_DEAD_TIMEOUT= <i>number_of_minutes</i>	<p>Specifies the amount of time that Scalix will wait for the next UAL command of any type (as opposed to UAL_IDLE_TIMEOUT, which is triggered by active commands only) from a UAL client before assuming a timeout. Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client.</p> <p>If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command.</p> <p>For serial connections, UAL_DEAD_TIMEOUT is overridden by UAL_SERIAL_DEAD_TIMEOUT. For local UAL clients, UAL_DEAD_TIMEOUT is overridden by UAL_LOCAL_DEAD_TIMEOUT.</p>
UAL_DEF_MIME_MN_OVERRIDE = <i>route</i>	<p>By default, a message generated using the Outlook client, with the "Send in RTF" flag unchecked, will be routed according to the default MIME mailnode entry in the Routing Table. Set this option if you want to use a different route for such messages.</p> <p>The route you specify must ultimately point to a MIME gateway. Specify a route as an OR address pattern, in the format specified for the -m option in the omaddr.t man page. For example, UAL_DEF_MIME_MN_OVERRIDE="internet,ux". If you specify the address pattern in this route as a Teletex value, you must specify an appropriate display character set in the UAL_DEF_MIME_MN_OVERRIDE_CS option.</p>
UAL_DEF_MIME_MN_OVERRIDE_CS= <i>character-set</i>	<p>This option specifies the display character set for the address pattern you specify in the UAL_DEF_MIME_MN_OVERRIDE option, if you entered it as a Teletex value. The default character set is ISO8859_1.</p> <p>This option has no effect if the UAL_DEF_MIME_MN_OVERRIDE option is not set.</p>

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_DEF_TNEF_MN_OVERRIDE= <i>route</i>	By default, a message generated using the Outlook client, with the "Send in RTF" flag checked, will be routed according to the default TNEF mailnode entry in the Routing Table. Set this option if you want to use a different route for such messages. The route you specify must ultimately point to a TNEF gateway. Specify a route as an OR address pattern, in the format specified for the <code>-m</code> option in the <code>omaddrtr</code> man page. For example, <code>UAL_DEF_TNEF_MN_OVERRIDE="internet,ux"</code> . If you specify the address pattern in this route as a Teletex value, you must specify an appropriate display character set in the <code>UAL_DEF_TNEF_MN_OVERRIDE_CS</code> option.
UAL_DEF_TNEF_MN_OVERRIDE_ CS= <i>character-set</i>	This option specifies the display character set for the address pattern you specify in the <code>UAL_DEF_TNEF_MN_OVERRIDE</code> option, if you entered it as a Teletex value. The default character set is <code>ISO8859_1</code> . This option has no effect if the <code>UAL_DEF_TNEF_MN_OVERRIDE</code> option is not set.
UAL_DIR_LIST_SORT_ORDER= <i>list_of_internal_attributes</i>	Specifies the order in which Directory attributes are sorted. The order is specified as a list of internal attribute names with each attribute separated by a <code>/</code> . The internal attribute names, which are numbers for the core Scalix attributes, are listed using the command <code>omshowatt -u</code> .
UAL_DIR_LIST_SORT_PROG= <i>absolute_program_name</i>	Specifies the program that sorts lists of Directory entries for UAL clients. The value <i>absolute_program_name</i> must specify the full path name of the sorting program together with any parameters that are necessary. The default Scalix sort program is <code>/bin/sort -f</code> . This is used if <code>UAL_DIR_LIST_SORT_PROG</code> is not set.
UAL_DIR_MOD_FULL_NAME=TRUE	Specifies that Full Name Checking is always done on the <code>UAL_CHKLIST</code> , <code>UAL_CHKNAM</code> , <code>UAL_DELENT</code> and <code>UAL_MODENT</code> commands.
UAL_DISABLE_BB=FALSE	Disables or enables public folder access. If this tag is set to <code>TRUE</code> , then if the user attempts to perform an action involving Public Folders the client displays an error message stating that the user has insufficient access capabilities to perform the action. The default is <code>FALSE</code> .
UAL_DISABLED_CLIENTS= <i>strings</i>	Specifies those UAL clients that are disabled from signing on to the Server. <i>strings</i> is a list of space-separated, quoted strings (use single quotes only). Each string is a client identity string as passed in the <code>UAL_INIT</code> command. Strings can contain wildcards. Sign-on attempts by identified clients are refused, and logged with an Error logging level. See also <code>UAL_ALLOW_DISABLED_CLIENTS</code> .
UAL_DISABLE_NESTED_BBS=TRUE	Stops the UAL Client Interface creating new nested Public Folders under top-level Public Folders.

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_DISALLOW_AUTO_PASSWORD=TRUE	<p>If set, a client cannot sign on to Scalix if the client has explicitly indicated that its password was obtained from a configuration file rather than having been entered interactively by a user. See also UAL_DISALLOW_NON_USER_PASSWORD.</p> <p>The user-specific setting of this option overrides the system-wide setting.</p> <p>Note that this mechanism is not intended to provide a secure indication.</p>
UAL_DISALLOW_NON_USER_PASSWORD=TRUE	<p>If set, a client cannot sign on to Scalix if the client has <i>not</i> explicitly indicated that its password was obtained interactively from a user. See also UAL_DISALLOW_AUTO_PASSWORD.</p> <p>Note that this option will only work with clients that supply the "password origination status". If a client does not support this element, then it will not be able to sign on even if the password is actually entered interactively by the user.</p> <p>The user-specific setting of this option overrides the system-wide setting.</p> <p>Note that this mechanism is not intended to provide a secure indication.</p>
UAL_DL_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the Distribution List area size limit. in kilobytes. A value of zero (0) means no size limit.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>
UAL_DO_LONG_INET_CHECK=FALSE	<p>Specifies whether the POP3 and IMAP4 Servers perform full checking of Internet addresses.</p> <p>The default value is FALSE, and this causes the POP3 and IMAP4 Servers to only look to see if a name has a DDA of type RFC-822 when checking if there is an Internet version of the name. This allows greater efficiency in cases where the names are either in a DDA or held in a Directory, since a check to see if the name is routable to a Unix queue is omitted.</p> <p>Set this option to TRUE to cause the full range of address conversions to be applied (according to the <code>unixmap.in</code> and <code>unixin.rules</code> steering files).</p>
UAL_FC_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the Filing Cabinet size limit. The value is set in kilobytes. A value of zero means no size limit.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_FORCE_IA_IN_ORN=FALSE	<p>Determines whether the UAL is forced to put the Internet address, as configured in the <code>SYSTEM</code> Directory, of all DL names in the message, doing additional directory lookups if necessary.</p> <p>Assuming that the <code>SYSTEM</code> Directory is populated with Internet addresses, when this option is set to <code>FALSE</code> (the default), the UAL will insert the Internet addresses of DL entries in the message, where it can do so without additional directory lookups. When a message reaches the Internet Mail Gateway, or is browsed by a POP3 or IMAP4 client, any ORNs without an Internet address are looked up in the Directory to retrieve the Internet address. So the Directory lookup overhead is at the Internet Mail Gateway, the POP3 or the IMAP4 interfaces.</p> <p>Set this option to <code>TRUE</code> to switch the overhead to the time when the message is sent.</p>
UAL_FORCE_TRACE_LEVEL= <i>trace_level</i>	<p>Sets the UAL trace level on a system-wide basis, overriding any trace value supplied by a client or set in the <code>user.cfg</code> file. <i>trace_level</i> can be any valid trace level, including 0 (zero), which switches off tracing.</p>
UAL_GIVE_GROUP5_INET_NAME_ STRICT=FALSE	<p>Determines whether the UAL Server gives a Group 5 Internet name in all cases.</p> <p>When set to <code>FALSE</code> (the default), the UAL Server gives a Group 5 Internet name in all cases where the address contains a DDA of type RFC-822. This occurs even for those addresses that are tunneled through Scalix, and causes replies to be incorrectly routed through the default Internet Mail Gateway.</p> <p>Set the option to <code>TRUE</code> to cause the UAL Server to only give a Group 5 Internet name when the address in the DDA is routable to the local Internet Mail Gateway.</p> <p>There are two circumstances under which setting this option to <code>TRUE</code> will have no effect:</p> <ul style="list-style-type: none"> If Scalix users and PDLs have Internet addresses configured. If Internet addresses of external users are put into the Group 5 Internet address at the incoming Internet Mail Gateway. You can prevent this by setting the option <code>INET_NO_IA_IN_ORN</code> to <code>TRUE</code>.
UAL_IDLE_SHUTDELAY= <i>number_</i> <i>of_minutes</i>	<p>Specifies the additional delay in shutting down a UAL client connection that has timed out.</p> <p><code>UAL_IDLE_SHUTDELAY</code> is used with <code>UAL_IDLE_TIMEOUT</code>. See also <code>UAL_IDLE_TIMEOUT</code>.</p> <p>For serial connections, <code>UAL_IDLE_SHUTDELAY</code> is overridden by <code>UAL_SERIAL_IDLE_SHUTDELAY</code>. For local UAL clients, <code>UAL_IDLE_SHUTDELAY</code> is overridden by <code>UAL_LOCAL_IDLE_SHUTDELAY</code>.</p>

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
<p>UAL_IDLE_TIMEOUT=<i>number_of_minutes</i></p>	<p>Specifies the amount of time that Scalix will wait for the next <i>active</i> UAL command from a UAL client before assuming a timeout (<i>PRE-PARE MESSAGE</i>, <i>ATTACH ITEM</i> are examples of active UAL commands, and <i>NEW MESSAGES</i> and <i>LIST ACK</i> are examples of passive UAL commands). Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client. (The sign off can be delayed using <i>UAL_IDLE_SHUTDELAY</i>.)</p> <p>If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command.</p> <p><i>UAL_IDLE_TIMEOUT</i> is used with <i>UAL_IDLE_SHUTDELAY</i>. For example:</p> <p>If <i>UAL_IDLE_TIMEOUT</i> is set to 30 minutes, and <i>UAL_IDLE_SHUTDELAY</i> is not set, the client is disconnected from the Server 30 minutes after the last active UAL command was issued.</p> <p>If <i>UAL_IDLE_TIMEOUT</i> is set to 30 minutes, and <i>UAL_IDLE_SHUTDELAY</i> is set to 10 minutes, 30 minutes after the last active UAL command is issued, the client displays a dialog box asking if the user wants to retain the connection. This dialog box is displayed for up to the 10 minutes specified by <i>UAL_IDLE_SHUTDELAY</i>.</p> <p>If the user responds within this time with a Yes, that is considered an active UAL command, and the <i>TIMEOUT</i> countdown restarts from the beginning.</p> <p>If the user responds with a No, the connection is closed.</p> <p>If the user does not respond within the 10 minutes, the connection is closed.</p> <p>For serial connections, <i>UAL_IDLE_TIMEOUT</i> is overridden by <i>UAL_SERIAL_IDLE_TIMEOUT</i>. For local UAL clients, <i>UAL_IDLE_TIMEOUT</i> is overridden by <i>UAL_LOCAL_IDLE_TIMEOUT</i>.</p>
<p>UAL_INTRAY_SIZE_LIMIT=<i>no_of_kilobytes</i></p>	<p>Sets the In Tray size limit in kilobytes. A value of zero (0) means no size limit. Set Message Store size limits using the <i>omlimit</i> command.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>
<p>UAL_ISO7_FROM_HOST=<i>language</i></p>	<p>This option is the same as <i>UAL_ISO7_HOST</i> except that the character set conversion only occurs when text is passed from Scalix to the client and not when it is passed back to the Server.</p>

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_ISO7_HOST= <i>language</i>	<p>This option allows clients to interoperate with a Scalix Message Store containing IA5 text that uses the ISO-7 extensions.</p> <p>If set, activates the option UAL_ISO7_<i>language</i> for any clients using <i>language</i>. IA5 text with the ISO-7 extensions are converted to the ISO8859/1 character set when downloaded to or displayed by a client and conversely, the ISO8859/1 characters are mapped back into IA5 with ISO-7 extensions when entering the Scalix system from a client.</p> <p>Only one instance of this option can be used and the <i>language</i> string must match a string in the <code>~/sys/LangMap</code> file.</p> <p>See also UAL_ISO7_FROM_HOST, UAL_ISO7_TO_HOST, and SR_ISO7_HOST.</p>
UAL_ISO7_ <i>language</i> = <i>ISO7_characters</i>	<p>Specifies how text using the ISO-7 extensions is converted to the ISO8859/1 character set. <i>ISO7_characters</i> is a list of ISO8859/1 characters to which the 12 special "ISO-7" characters are mapped. The IA5 characters used as special ISO-7 characters are:</p> <p><code>#\$@[\\]^`{}~</code></p> <p>The ISO8859/1 equivalents (<i>ISO7_characters</i>) must be specified in the same order. Ensure that the ISO8859/1 equivalents are entered into the file using the ISO8859/1 character set! The <i>language</i> must correspond to the language set in the UAL_ISO7_HOST option and the UAL_ISO7_HOST option must be present to activate this option.</p> <p>See also SR_ISO7_<i>language</i>.</p> <p>UAL_ISO7_TO_HOST=<i>language</i></p> <p>This option is the same as UAL_ISO7_HOST except that the character set conversion only occurs when text is passed from the client to Scalix and not when it is passed back to the client.</p>
UAL_KILL_REMOTE_SIGNON_2= TRUE	<p>Allows the Scalix Server to kill a current user session in order to allow the user to sign on again.</p> <p>If a user session is terminated abnormally (for example, if a user reboots their PC), the session can continue to exist on the Server. This could prevent the user from signing on again. Setting this option to TRUE allows the Scalix Server to kill the user's oldest session, so allowing the user to sign on.</p> <p>The Scalix Server permits 17 concurrent signons. If this option is set to TRUE, and the user tries to connect for the eighteenth time, the Scalix Server kills the user's oldest session, and then allows them to sign on again. If the option is set to FALSE, the Scalix Server will not permit the user to sign on.</p> <p>Note that some clients can set a lower limit for the number of concurrent signons.</p>
UAL_LIST_CACHE_SIZE= <i>number_of_message_parts</i>	<p>Specifies the number of message parts that can be held in memory by a UAL process. This entry reduces I/O by forcing Scalix to keep the message in memory instead of creating and then opening one file for each message part and the message header. The default is 4, which equates to a header record and three body parts.</p>

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_LOCAL_DEAD_TIMEOUT= <i>number_of_minutes</i>	<p>Specifies the amount of time that Scalix will wait for the next UAL command of any type (as opposed to UAL_LOCAL_IDLE_TIMEOUT, which is triggered by active commands only) from a <i>local</i> UAL client before assuming a timeout. Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client.</p> <p>If a timeout period is not specified, Scalix assumes the connection to the local UAL client is good regardless of how long it has been waiting for another command.</p> <p>UAL_LOCAL_DEAD_TIMEOUT overrides UAL_DEAD_TIMEOUT. To remove a timeout for local UAL clients that was set using UAL_DEAD_TIMEOUT, set UAL_LOCAL_DEAD_TIMEOUT to 0.</p>
UAL_LOCAL_IDLE_SHUTDELAY= <i>number_of_minutes</i>	<p>Specifies the additional delay in shutting down a local UAL client connection that has timed out.</p> <p>UAL_LOCAL_IDLE_SHUTDELAY is used with UAL_LOCAL_IDLE_TIMEOUT.</p>
UAL_LOCAL_IDLE_TIMEOUT= <i>number_of_minutes</i>	<p>Specifies the amount of time that Scalix will wait for the next <i>active</i> UAL command from a local UAL client before assuming a timeout (PREPARE MESSAGE, ATTACH ITEM are examples of active UAL commands, and NEW MESSAGES and LIST ACL are examples of passive UAL commands). Once the timeout period has been reached, Scalix assumes the serial connection to the client has been lost and signs off the user of the UAL client. (The sign off can be delayed using UAL_LOCAL_IDLE_SHUTDELAY.)</p> <p>If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command.</p> <p>UAL_LOCAL_IDLE_TIMEOUT is used with UAL_LOCAL_IDLE_SHUTDELAY.</p> <p>UAL_LOCAL_IDLE_TIMEOUT overrides UAL_IDLE_TIMEOUT. To remove a timeout for local UAL clients that was set using UAL_IDLE_TIMEOUT, set UAL_LOCAL_IDLE_TIMEOUT to 0.</p>
UAL_LOCAL_IGNORE_PASSWORD= TRUE	<p>Removes the password entry stage from the sign on process. The sign on will succeed only if the user has logged in using their Scalix mailbox Linux login, and if the user is using a local UAL Client.</p>

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_MAP_ALIAS_AT_MAIL=FALSE	By default, recipient O/R Addresses entered as aliases in a distribution list are displayed as those aliases to the recipients. Set this option to TRUE to cause aliases to be rewritten as their <i>real</i> O/R Addresses when the message is submitted to Scalix (unless both sender and recipient are using the Outlook mail client). This enables the user to use and see alias names when preparing a message, but the recipients of the message only see the <i>real</i> names in the distribution list, not the alias names. If both sender and recipient are using the Outlook mail client, then setting this option to TRUE has no effect, and the recipients continue to see the alias names in the distribution list.
MAX_SIGNON_PER_USER= <i>number</i>	Specifies the number of simultaneous signons that a user can have. The default is 17.
UAL_MOD_BB_ITEMS=TRUE	Determines whether items attached to Public Folders can be modified. Set this option to FALSE to prevent modification of Public Folder items. In this case, users can still be able to add top-level items to Public Folders, or delete top-level items from Public Folders, depending on the Public Folder's ACL. When set to TRUE (the default), master items can be modified, although slave items cannot. Public folders can be accessed only by Premium users. For more information, see "About Scalix Product Editions".
UAL_MSTORE_SIZE_LIMIT= <i>no_of_kilobytes</i>	Sets the overall message store size limit. The value is set in kilobytes. A value of zero means no size limit. Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.
UAL_NAMED_PIPE_BLOCK_SIZE= <i>block_size</i>	Sets the physical block size for TCP/IP Named Pipes client connections. The default is 1380 bytes.
UAL_NMP_DELAY= <i>number_of_milliseconds</i>	Specifies the time delay before returning a reply to a command from a UAL client using a TCP/IP Named Pipes connection. By default, there is no time delay, but this can mean the receiving client system can <i>miss</i> the reply when it is sent. To overcome this problem set the delay to between 10 and 100 milliseconds.
UAL_NO_AUTOGEN_IA=FALSE	Determines, for certain UAL clients, whether automatic Internet address mapping is enabled. This option only affects those UAL clients that use the UAL_ENTADD command to add Directory entries. Set this option to TRUE to override automatic Internet address mapping for these UAL clients. The default is FALSE . See the option INET_USE_AUTO_IAM .
UAL_NO_DESIGNATE_SIGNON=TRUE	Removes the designate sign on feature.

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_NO_IA_IN_ORN=FALSE	Determines whether the UAL puts the Internet addresses of the sender, recipients or DL names, into the message. When set to <code>FALSE</code> (the default), Internet addresses are inserted into the message if they can be determined without additional directory lookups.
UAL_NO_REPLY_BLOCKING=TRUE	If set, multiple UAL Client Interface replies are not blocked up before being sent to a UAL remote client. This is used to overcome data-comm problems that result from large blocks being sent from the Server.
UAL_NO_WB_EMPTY=TRUE	Stops a user's Waste Basket being emptied when the user has finished using a UAL client and signs off. If this option is set, use the command <code>omtidy</code> or <code>omtidyall</code> to ensure Waste Baskets continue to be emptied regularly.
UAL_OUTTRAY_SIZE_LIMIT= <i>no_of_kilobytes</i>	Sets the Out Tray size limit in kilobytes. A value of zero (0) means no size limit. Set Message Store size limits using the <code>omlimit</code> command. Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.
UAL_PASSWORD_AGED=IGNORE WARN or ERROR	This option determines the effect of an expired password on a user signing on to Scalix through a client. The default value is <code>ERROR</code> . If the user's password has expired, an error is generated when the user attempts to signon and the signon fails. The signon can only succeed when a valid new password is supplied. If the value is set to <code>WARN</code> and the user's password has expired, the user can sign on using the expired password but a warning message is placed in their In Tray stating that their password has expired and should be changed immediately. (This message appears in the In Tray for the first signon of the day.) If the value is set to <code>IGNORE</code> any user password expiration condition is ignored (a Scalix user will be allowed to signon even if their password has expired.)
UAL_PEND_SIZE_LIMIT= <i>no_of_kilobytes</i>	Sets the Pending Tray size limit in kilobytes. A value of zero (0) means no size limit. Set Message Store size limits using the <code>omlimit</code> command. Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.
UAL_POP3_HOSTNAME= <i>hostname</i>	Set the POP3 Server or the IMAP4 Server hostname so that addresses display as <code>name@hostname</code> . By default, the name of the Scalix host is used.
UAL_POP3_LANG= <i>language</i>	Specifies the Scalix language to use for error messages returned by the POP3 or IMAP4 client. The default is <code>C</code> .

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_POP3_TIMEOUT= <i>no_of_seconds</i>	Specifies the number of seconds of inactivity allowed before connection to the POP3 Server will timeout. The default is 600 seconds.
UAL_POP3_TRACE=TRUE	If set, information from the <code>in.pop3d</code> process is traced and placed in the <code>~Scalix/tmp</code> directory. You can set this option to <code>DETAIL</code> to generate more detailed logging. Set the option to <code>FALSE</code> to prevent logging.
UAL_PRINT_SERVER_ONLY=TRUE	If set, all printing by UAL clients goes through the Print Server. See also <code>UAL_PRINT_SPECIFICATION</code> .
UAL_PRINT_SPECIFICATION= <i>print_command</i>	If set, <i>print_command</i> overrides any printer specification supplied by a UAL client. The <i>print_command</i> can either be a Linux printer command line or a Print Server printer specification.
UAL_PWD_WARNING_DAYS= <i>days</i>	Activates the mechanism to generate advisory messages to users whose mailbox passwords are due to expire within the period specified by <i>days</i> . The warning message appears as a new message in the user's In Tray for the first signon of the day. Use this option if clients are being used that do not recognize the <i>password expired</i> signon error. These clients cannot signon successfully once the user's password has expired.
UAL_SCK_DELAY= <i>number_of_milliseconds</i>	Specifies the time delay before returning a reply to a command from a UAL client using a TCP/IP Sockets connection. By default, there is no time delay, but this can mean the receiving client system can <i>miss</i> the reply when it is sent. To overcome this problem set the delay to between 10 and 100 milliseconds.
UAL_SEND_OBJECT_FILES=TRUE	Determines whether an object file (created with the UAL <code>addobj</code> call) is mailed with the message to which it is attached. When set to <code>TRUE</code> , the UAL submits both the message and any attached object files (assuming that the object files do not have the <code>UAL_ADDOBJ_NOT_MAIL</code> flag set).
UAL_SERIAL_DEAD_TIMEOUT= <i>number_of_minutes</i>	Specifies the amount of time that Scalix will wait for the next UAL command of any type (as opposed to <code>UAL_SERIAL_IDLE_TIMEOUT</code> , which is triggered by active commands only) from a UAL client <i>using a serial connection</i> before assuming a timeout. Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client. If a timeout period is not specified, Scalix assumes the <i>serial</i> connection to the UAL client is good regardless of how long it has been waiting for another command. <code>UAL_SERIAL_DEAD_TIMEOUT</code> overrides <code>UAL_DEAD_TIMEOUT</code> . To remove a timeout for UAL clients that was set using <code>UAL_DEAD_TIMEOUT</code> , set <code>UAL_SERIAL_DEAD_TIMEOUT</code> to 0. This removes the timeout for all UAL clients using a serial connection.

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_SERIAL_IDLE_SHUTDELAY= <i>number_of_minutes</i>	Specifies the additional delay in shutting down a UAL client <i>serial</i> connection that has timed out. UAL_SERIAL_IDLE_SHUTDELAY is used with UAL_SERIAL_IDLE_TIMEOUT.
UAL_SERIAL_IDLE_TIMEOUT= <i>number_of_minutes</i>	Specifies the amount of time that Scalix will wait for the next <i>active</i> UAL command from a UAL client <i>using a serial connection</i> before assuming a timeout (PREPARE MESSAGE, ATTACH ITEM are examples of active UAL commands, and NEW MESSAGES and LIST ACL are examples of passive UAL commands). Once the timeout period has been reached, Scalix assumes the serial connection to the client has been lost and signs off the user of the UAL client. (The sign off can be delayed using UAL_SERIAL_IDLE_SHUTDELAY.) If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command. UAL_SERIAL_IDLE_TIMEOUT is used with UAL_SERIAL_IDLE_SHUTDELAY. UAL_SERIAL_IDLE_TIMEOUT overrides UAL_IDLE_TIMEOUT. To remove a timeout for UAL clients that was set using UAL_IDLE_TIMEOUT, set UAL_SERIAL_IDLE_TIMEOUT to 0. This removes the timeout for all UAL clients using a serial connection.
UAL_SET_ROC_ON_ND= TRUE or FALSE	Overrides client setting of "Return of contents on non-delivery". Set to TRUE to request a return of contents for all non-delivered messages. Set to FALSE to prevent return of contents for all non-delivered messages.
UAL_SHOW_8BIT_T61_AS_1167= TRUE	If set, teletex body parts (file code 1736) are presented to clients as normal 8-bit text (file code 1167). This enables existing Western European clients to work with no loss of features when handling teletex body parts that contain 8-bit character sets.
UAL_SIGNON_ALIAS=YES or ONLY	Specifies whether aliases are used for sign on. Any UAL_SIGNON_ALIAS entries in <code>user.cfg</code> take precedence over the UAL_SIGNON_ALIAS entry in <code>general.cfg</code> . (This enables you to set a default use of aliases in <code>general.cfg</code> and then set overrides for specific users in <code>user.cfg</code> .) The YES value means aliases can be used to sign on with, users can also continue to use their Personal Name if they want to. The ONLY value means the aliases only can be used to sign on with, the Personal Name cannot be used any more. UAL_SIGNON_ALIAS is used with UAL_SIGNON_ALIAS_CONFIG and UAL_USE_SIGNON_ALIAS.

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_SIGNON_ALIAS_CONFIG=SYS or USER	Allows users to sign on using an alias. The <code>SYS</code> value means that everyone can sign on using an alias. The <code>USER</code> value means that alias sign on entries in <code>user.cfg</code> will be used when they exist, and will take precedence over any alias sign on entries in <code>general.cfg</code> . UAL_SIGNON_ALIAS_CONFIG is used with UAL_SIGNON_ALIAS and UAL_USE_SIGNON_ALIAS.
UAL_SINGLE_TEMP_DIR= <i>temp-directory</i>	Specifies the location of user temporary files. Normally (when this option is not set), the user temporary files are stored in the directory files for each user. Specify a value for this option to cause all user temporary files to be written to a single directory. This allows you to use a high-speed (and possibly low-recovery) file system (for example, a RAM disk) to store these temporary files. The directory you specify must have: permissions 771 an owner of <code>scalix</code> a group of <code>scalix</code> a path length of 225 characters or less For example, to create a directory called <code>usr_tmp</code> , enter the following commands: <code>mkdir \$(omrealpath '~/usr_tmp')</code> <code>chown scalix:scalix \$(omrealpath '~/usr_tmp')</code> <code>chmod 771 \$(omrealpath '~/usr_tmp')</code>
UAL_SIZE_ERR_TO_USER=TRUE	Specifies that a UAL error message is generated when a user tries to create an item in their Filing Cabinet or Distribution List area once it has exceeded the limit set by <code>omlimit</code> .
UAL_SIZE_MSG_TO_ENU=TRUE	Specifies that a message is sent to the Error Notification user when a user's In Tray, Pending Tray or Waste Basket exceeds the set warning limit, boundary limit, or maximum limit. See UAL_SIZE_WARNING_BOUNDS and UAL_SIZE_WARNING_LIMIT.
UAL_SIZE_MSG_TO_USER=TRUE	Specifies that a message is sent to the user when their In Tray, Pending Tray or Waste Basket exceeds the set warning limit, boundary limit, or maximum limit. See UAL_SIZE_WARNING_BOUNDS and UAL_SIZE_WARNING_LIMIT.
UAL_SIZE_ON_RECEIPT=FALSE	Specifies whether a user whose message store components exceed their configured limits is prevented from receiving messages. When this option is set to <code>FALSE</code> (the default), users are not prevented from receiving messages even if the size of their message store component is greater than its configured limit.

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_SIZE_ON_SEND=FALSE	Specifies whether a user whose message store components exceed their configured limits is prevented from sending messages. When this option is set to TRUE , then message delivery rules can be implemented that limit a user's ability to send messages. These rules utilize the OMLIMIT-EXCEEDED message attribute filter. When this option is set to FALSE (the default), then rules based on the OMLIMIT-EXCEEDED filter have no effect.
UAL_SIZE_WARNING_BOUNDS= <i>percent_increase</i>	Specifies the boundary levels for warnings between the warning limit and the maximum limit. For example, if set to 5, warnings will be sent when the size of an In Tray, Pending Tray or Waste Basket increases by 5% or more since the last warning. To enable warning messages to be sent, you must have set at least one of these options: UAL_SIZE_MSG_TO_ENU and UAL_SIZE_MSG_TO_USER .
UAL_SIZE_WARNING_LIMIT= <i>percentage_of_max_limit</i>	Specifies the percentage of the maximum limit, set on the size of the In Tray, Pending Tray or Waste Basket, that should be reached before a warning messages is generated. For example, if set to 80, warnings will be sent when the In Tray, Pending Tray or Waste Basket area reaches 80 percent of its maximum limit. To enable warning messages to be sent, you must have set at least one of these options: UAL_SIZE_MSG_TO_ENU and UAL_SIZE_MSG_TO_USER .
UAL_SOCKET_BLOCK_SIZE= <i>block_size</i>	Sets the physical block size for Sockets client connections. The default is 1380 bytes.
UAL_TTX_NAME_FORMAT_LANG= <i>attribute order</i>	Specifies the order in which Personal Name attributes are displayed for clients using the UAL Client Interface display program (item.browse). The four Personal Name Attributes are represented with the following letters: S: Surname F: Given Name I: initials G: Generation Qualifier Enter the letters in the order you want the attributes to be displayed. The LANG part of the option specifies the language the format is used for. For example, to display native Japanese names in their natural form, specify the option like this: UAL_TTX_NAME_FORMAT_NIPPON=SF
UAL_TTX_NAME_SHOW_ALL=TRUE	Sets the UAL Client Interface display program (item.browse) to display all teletex O/R Address attributes regardless of whether the correct client character set is configured. By default, these address attributes are not displayed unless a suitable client character set is configured.

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
UAL_TTX_NAME_SUBST=TRUE	Substitutes the teletex O/R Address attributes with the corresponding printable string attributes before displaying the message. This option is for clients using the UAL Client Interface display program (<code>item.browse</code>).
UAL_USE_SIGNON_ALIAS=FALSE or TRUE	Specifies whether the alias is used after sign on. If you set <code>UAL_USE_SIGNON_ALIAS</code> to <code>FALSE</code> , the UAL client reverts to using the user's Personal Name for the remaining time the user is signed on (the alias or Personal Name is used on the "Creator" part of a message). If you set <code>UAL_USE_SIGNON_ALIAS</code> to <code>TRUE</code> , the alias is used for the remaining time the user is signed on. <code>UAL_USE_SIGNON_ALIAS</code> is used with <code>UAL_SIGNON_ALIAS</code> and <code>UAL_SIGNON_ALIAS_CONFIG</code> .
UAL_WB_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the Waste Basket size limit in kilobytes. A value of zero (0) means no size limit. Set Message Store size limits using the <code>omlimit</code> command.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p> <p><code>UXO_CHECK_TYPES_OF_DDA=DDA_type</code></p> <p>Specifies the DDA types that are acceptable as valid internet addresses, to enable the UAL Client Interface and the Internet Mail Gateway to route the message to the correct destination for the recipient. <i>DDA_type</i> is one of the following:</p> <p>One or more valid DDA types, for example:</p> <p>RFC-822 HPMEXT1 HPMEXT2 HPMEXT3 HPMEXT4</p> <p>You can specify up to 10 DDA types; if you do specify more than one type, separate them with commas. For example:</p> <p>RFC-822,HPMEXT1,HPMEXT2,HPMEXT3,HPMEXT4</p> <p>FALSE</p> <p>No DDA type checking is performed. This behavior is as for Scalix systems before B.05.10.</p> <p>If your Scalix Directory contains DDAs with no type specified and they are not valid internet addresses, you are recommended to set this option to:</p> <p><code>UXO_CHECK_TYPES_OF_DDA=RFC-822,HPMEXT1,HPMEXT2,HPMEXT3,HPMEXT4</code></p> <p>If this option is not present, the default setting is:</p> <p><code>UXO_CHECK_TYPES_OF_DDA=,RFC-822,HPMEXT1,HPMEXT2,HPMEXT3,HPMEXT4</code></p> <p>The leading comma means that DDAs with no type specified are also acceptable as valid internet addresses.</p>

Table 28: UAL Client Interface Options and their Descriptions

Option	Description
OMLIMIT_MIN_WARN_INTERVAL	<p>If the "omlimit -e u" sanction is enabled, the OMLIMIT_MIN_WARN_INTERVAL option also manages the interval during which omlimit-related messages are sent to a user. The default value for the OMLIMIT_MIN_WARN_INTERVAL option is one day (1d).</p> <p>Example settings for this option are:</p> <p>1h40m20s (1 hour 40 minutes and 20 seconds)</p> <p>2d40 (2 days and 40 seconds)</p> <p>6000 (6000 seconds/100 minutes)</p> <p>The OMLIMIT_MIN_WARN_INTERVAL option also manages the interval between OMLIMIT-EXCEEDED notifications when the service router processes message delivery rulesets.</p>

Virus Protection Options

Table 29: Virus Protection Options and their Descriptions

Option	Description
SR_VS_DO_VIRUS_SCAN=FALSE	<p>In the absence of the <code>ALL-ROUTES.VIR</code> ruleset file, this option determines whether virus scanning is active. If the <code>ALL-ROUTES.VIR</code> file exists, then the rules within that file determine the virus scanning/cleaning action that will be taken.</p> <p>Set the option to <code>TRUE</code> to cause the Service Router to check all message attachments for viruses. If a virus is found, the message is not routed, and Scalix generates a non-delivery notification.</p> <p>Set the option to <code>FALSE</code> to disable virus checking.</p> <p>Note that the performance of your Scalix system can be degraded if you enable virus checking and a large number of viruses are detected, since each virus detected will cause Scalix to generate a non-delivery report.</p>
SR_VS_IGNORE_ITEM_TYPES= <i>filetype-no</i>	<p>Specifies the filetypes of items that will not be scanned for viruses. By default, when virus scanning is enabled, either by setting the <code>SR_VS_DO_VIRUS_SCAN</code> option to <code>TRUE</code> or by creating the <code>ALL-ROUTES.VIR</code> ruleset file, all filetypes are scanned for viruses. Use this option to prevent certain filetypes from being scanned.</p> <p><i>filetype-no ...</i> is a colon-separated list of numerical file codes, as specified in <code>~/nls/language/filetype</code>.</p> <p>For example, set <code>SR_VS_IGNORE_ITEM_TYPES</code> to <code>1167</code> to prevent text files from being scanned.</p>
SR_VS_TEST_SCAN_SL= <i>string</i>	<p>Specifies the location of the test virus shared library. If this file is in its default location, you must set this option to <code>/opt/scalix/version/lib/libom_testvs.sl</code> if you want to test your virus scanning configuration.</p>
SR_VIRUS_SCAN_TYPE= <i>string</i>	<p>Specifies whether virus checking is operating in test mode. Set this option to <code>"Test Scan"</code> to cause the Service Router to check messages and generate a non-delivery notification if the first five characters of any attachment is <code>VIRUS</code>.</p> <p>Note: if you set this option to <code>"Test Scan"</code>, you must also set the <code>SR_VS_TEST_SCAN_SL</code> option to the location of the test virus shared library.</p> <p>If you set this option to <code>"Generic"</code>, <code>~/sys/omvscan.cfg</code> determines the virus scanning engine to use. Also, you must copy <code>omvscan.map</code> to the <code>~/rules/</code> directory to enable virus scanning.</p>

Export Process Options

Table 30: Export Process Options and their Descriptions

Option	Description
XP_START_IMPORT_DELAY= <i>seconds</i>	If set, specifies the number of seconds that the <code>xport.in</code> process will wait when invoked by the Service Router in recovery mode before starting to process messages and put them on the Service Router queue. The delay is only observed to a resolution of 5 seconds. A value of zero will cause <code>xport.in</code> to start processing the messages immediately. The default is 60 seconds.

Miscellaneous Options

Table 31: Miscellaneous Options and their Descriptions

Option	Description
AK_ACK_MSG_PRI=2	Determines the priority to use for acknowledgments if <code>AK_ACK_SAME_PRI</code> is not set or the priority of the message being acknowledged is not known. Valid values are 0 (normal), 1 (non-urgent), and 2 (urgent). The default is 2 (urgent).
AK_ACK_SAME_PRI=TRUE	When set to <code>TRUE</code> , this option causes the priority of an acknowledgment to be the same as that of the message being acknowledged where that can be determined. Otherwise, the value of <code>AK_ACK_MSG_PRI</code> is used. The default is <code>FALSE</code> .
CT_OLD_PENDING_SIZE_MODE= FALSE	Determines whether items in the message store that are pending deletion are included in the reported size of the message store. When this option is set to <code>TRUE</code> , then items that are pending deletion are included in the reported size. The size is not reduced until the items are actually deleted. Leave this option <code>FALSE</code> (the default) if you want the reported size of the message store to be reduced as soon as items are marked for deletion. This is useful, for example, when users have limits configured on their message store sizes: if set to <code>FALSE</code> , users will see an immediate effect on their reported message store sizes when they delete items. If set to <code>TRUE</code> , the message store size will not change until the user logs out of Scalix (for multiple signons, the last user must log out). If you change the value of this option, you should run the <code>omscan -A -S</code> command to regularize the reported message store sizes.
CT_OLD_SIZE_METHOD=FALSE	Specifies how the size of containers in the message store is updated. When set to <code>FALSE</code> (the default), container sizes are written to the Container Access Monitor, so that all processes have access to them. This ensures that all processes use the same reported container sizes, and is particularly useful when message store size limits are configured. If set to <code>TRUE</code> , container sizes are stored within an individual process before being written to disk.

Table 31: Miscellaneous Options and their Descriptions

Option	Description
DIA_NO_T61_SUBJECT=FALSE	Determines whether the <code>omcontain</code> command attempts to display T61 subjects. When set to <code>FALSE</code> (the default), <code>omcontain</code> does not display T61 subjects.
IM_MAKE_MSG_ID_GLOBAL_UNIQUE=TRUE	Specifies that Scalix message IDs should use the long format. Set the value to <code>FALSE</code> if you want message IDs to use the short format. Note, however, that this can cause message IDs not to be globally unique. If you change the value of this option, you must restart Scalix for the change to take effect.
SMTPD_PWD_TRANSITION=FALSE	Determines whether the SMTP relay generates SASL passwords and stores them in the user list. Set this option to <code>TRUE</code> to cause the SMTP Relay to generate alternative SASL passwords when a PLAIN password authentication succeeds. These SASL passwords are then stored in the user list. Reset this option to <code>FALSE</code> (the default) when the user list contains all the SASL passwords.
USRL_AUTO_GEN_AUTHID=G_I_S	Specifies the method used to generate the name part of the Authentication ID. If you want to use the OMID method instead, enter <code>DEFAULT</code> for this value.

Client-specific Configuration Options

A subset of the options used in the `general.cfg` file can be specified for individual Scalix clients. The options for each client are held in a file with the name of the client host's Fully Qualified Domain Name (FQDN), in the directory `~/sys/client.cfg`. For example, for a client on the host `north.sales.alpha.com`, the client configuration options file would be named `~/sys/client.cfg/north.sales.alpha.com`.

Note	Scalix Connect for Microsoft Outlook can be used only by Premium users. For more information, see "About Scalix Product Editions".
-------------	--

The `client.cfg` directory does not exist by default and must be created. It should be owned by the user `scalix` with permissions of `555 (dr-xr-xr-x)`. Client files within this directory should be owned by the user `scalix` with permissions of `444 (-r--r--r--)`.

See the descriptions in the <Xref_Color>“System-Wide Configuration Options” section for the following options:

- IMAP_AUTOMATIC_MDN=FALSE
- IMAP_BB_FOLDER_PREFIX=#bb
- IMAP_BB_FOLDER_SEPARATOR=/
- IMAP_DELETE_SUBFOLDERS=FALSE
- IMAP_FOLDER_PREFIX=
- IMAP_FOLDER_SEPARATOR=/
- IMAP_IDLE_TIMEOUT=30
- IMAP_LOGFILE=~/.tmp/imap.%h
- IMAP_LOGLEVEL=0
- IMAP_MDSENT_FLAG=\$MdnSent
- IMAP_MIN_SIZE_ESTIMATE=0
- IMAP_REMOTE_UAL_ENABLED=TRUE
- IMAP_SEARCH_TIMEOUT=0
- IMAP_UAL_TRACE_LEVEL=0
- IMAP_X_NETSCAPE_URL=

The following table lists the IMAP4 options you can specify for individual clients.

Table 32: IMAP4 Options and their Descriptions

Option	Description
IMAP_CAPABILITIES= <i>capabilities-list</i>	See the description in the <Xref_Color>“System-Wide Configuration Options” section. Note that capabilities you specify here are <i>added</i> to those specified on a system-wide and per-user basis.
IMAP_MAILSTORE_HOST= <i>hostname</i>	Specifies the fully qualified domain name of the Scalix host to which the IMAP Server connects. Use this option when the IMAP4 Server does not reside on the same machine as the relevant Scalix message store.

Outlook Configuration Options

The `mapi.cfg` file sets parameters for all Scalix MAPI Client users and provides a way to customize some Outlook client functionality. This file is in the `~/nls/C/` directory on the Scalix Server if you are using Auto-upgrades (see the Scalix installation Guide for more information).

After a user logs into Outlook for the first time, the `mapi.cfg` file is automatically downloaded to the local system from the Scalix Server.

The local `mapi.cfg` file is downloaded to the `C:\Documents and Settings\user\Local Settings\Application Data\Scalix\Scalix\MAPI\Profiles\profile_name\Scalix` directory.

The tables below list and describe the parameters you can configure in the `mapi.cfg` file.

Caution

If you modify any of these parameters then install (manually or automatically) an updated version of the MAPI service provide, the `mapi.cfg` file is replaced (overwrites the existing file) and the changes will be lost.

[AutoUpgrade] Parameters

Use this section to set Auto-upgrade options.

Table 33: MAPI Configuration Options and their Descriptions

Parameter	Description
n	<p>The mapi.cfg version number that is used to determine whether auto-upgrades occur.</p> <p>This number is also used to determine whether mapi.cfg is downloaded to update other administrative settings.</p> <p>NOTE: If the version number of the mapi.cfg file on the user system is <i>greater</i> than the version number of the mapi.cfg file on the server, Scalix does not upgrade Scalix Connect on the user system with the latest version of the MAPI service provider and/or update the mapi.cfg file.</p>
SetupPath (8.2 to 9.1.1) HTTPSetupPath (9.2 and up)	<p>The path to the shared drive/directory that contains the source Scalix Connect installation files. The SetupPath value must be a valid UNC path. The HTTPSetupPath must be a valid http:// address.</p>
HTTPUpdateInstallMgr	This value is set to 1. Do not modify this value.
HTTPUpgradeExemptList	Allows you to specify users that you do not want to upgrade to the latest version of Scalix Connect.
MinimumScalixVersion	The version number of the Scalix Connect dynamic link libraries.
ForwardInstallLogsTo	<p>The administrator mailbox to which auto-upgrade results are sent.</p> <p>Enter 0 to disable error logging.</p>
ForwardInstallLogsFrom	The text that displays in the From field of the Auto-install log message.
ForwardInstallLogsSubject	The Subject line of the e-mail that includes the auto-upgrade results.
UseLocalTimeVSGMT	Specify whether you want to use local time or Greenwich Mean Time (GMT) to auto-upgrade users. Enter 1 to use local time, or 0 to use GMT.
UpgradeIntervalTimeCheck	<p>The (metric) time at which Scalix polls client systems to verify whether they are using the latest version of Scalix Connect. For example, enter 8 to poll for Auto-upgrade status information at 8 am. Enter 22 to poll for information at 10 pm.</p> <p>Entering value of 24 or greater causes Scalix to poll for Auto-upgrade information in intervals (by seconds). For example, if you want to poll client systems every hour, enter 3600.</p>

[Startup] Parameters

Use this section to set startup options.

Table 34: Startup Parameters and their Descriptions

Parameter	Description
AddressBookDownloadReminderInterval	<p>This option displays the number of days since you last downloaded a copy of the Address Book from the Scalix Server. Scalix Connect also reads the value of <code>PreviousABDownloadDate</code> in the registry key of <code>HKEY_LOCAL_MACHINE\SOFTWARE\SCALIX\MAPI</code>.</p> <p>Scalix Connect calculates the difference between the two dates. If the difference is greater than the value displayed in the <code>AddressBookDownloadReminderInterval</code> option, Scalix Connect displays a reminder to users to download a copy of the Address Book from the Scalix Server. To remind users to download Address Books monthly, set the value in the <code>AddressBookDownloadReminderInterval</code> option to 30.</p>
AlwaysShowLogon	<p>The <code>AlwaysShowLogon=1</code> option specifies that the user is always prompted for a password at startup. If you enabled password storing at logon, you are not be prompted for a password.</p>

[Addressing] Parameters

Addressing parameters affect the interpretation of Scalix addresses on messages.

Table 35: Addressing Parameters and their Descriptions

Parameter	Description
InternetToOM	<p>On an incoming message, Scalix Connect converts a Scalix address that includes a DDA (Domain Defined Attribute) type of RFC-822 to an address type of SMTP. For more information on DDAs, see the <i>Scalix Administration Guide</i>.</p> <p>Scalix Connect then uses the DDA for the revised address. For example, Scalix Connect replaces the Scalix address of a message such as: <code>chris/linux/dd.RFC-822=cwolfe@pwd.scalix.com</code> with an SMTP address such as <code>cwolfe@pwd.scalix.com</code>.</p> <p>You can override this behavior and keep the address as an Scalix type by including the setting <code>InternetToOM=1</code> in this section.</p>
HPMEXTToSMTP	<p>You can extend the conversion of Scalix addresses to SMTP addresses that include a DDA type of the form <code>HPMEXTn</code> by including the setting <code>HPMEXTToSMTP=1</code> in this section. The <code>InternetToOM=1</code> option takes precedence over the <code>HPMEXTToSMTP=1</code> option.</p>

[Display] Parameters

The options in the `Display` section specify the following:

- the maximum number of items within a container
- which attributes are displayed
- the maximum line length in plain text messages

- how Internet Addresses are formatted

Settings in the [Display] section affect the presentation of Scalix addresses, for example, the displayed part of an address but not the underlying message address or type.

Table 36: Display Parameters and their Descriptions

Parameter	Description
MaxContainerSize	The maximum number of messages that are listed on opening a folder can be configured using the <code>MaxContainerSize</code> setting. This setting can take an integer value between 20 and 32767. The default value is 32767. If the configured limit is exceeded then a warning message is displayed. Archiving (or auto-archiving) a folder that contains more than <code>MaxContainerSize</code> items causes the container-limit warning message to be displayed. Scalix Connect archive those items within the limit of 32767.
ShowMailnodes	Either the mailnode attributes or custom attributes can be displayed, but not both. You can set only one of the following options for attribute display. The format of each option is described in the following sections. The <code>ShowMailnodes=1</code> option specifies that the mailnode attributes are displayed in message headers along with the name (<code>Personal Name / OU1 , OU2</code>). This is useful when selecting similar entries from the directory. Without this setting you must scroll across the window to see the mailnode. The setting also applies to the display of addresses when either composing or reading a message. A way to resolve an unresolved address is to right-click on the address. This displays possible alternatives, which include the mailnode.
ShowCustomAttributes	The <code>ShowCustomAttributes=1</code> option specifies that custom attributes, other than the mailnode details (<code>Personal Name / OU1 , OU2</code>), are displayed in message headers. If you set this option, you must also set the <code>UserDefinedAttributes</code> option to specify the attributes to be displayed.
UserDefinedAttributes	<code>UserDefinedAttributes=%(attr_tag)%(attr_tag)</code> where <code>attr_tag</code> is the internal attribute tag for an Scalix Directory attribute type defined in the <code>~/sys/dir.attrs</code> file. This tag can be either a predefined Scalix system attribute type (a numerical value) or a custom attribute type you have defined. Use the <code>omshowatt -u</code> command on the Scalix server to list the internal attribute tags. For example, specifying the line: <code>UserDefinedAttributes=%(1)%(8)%(9)</code> displays the Surname, Organization, and Country Code as the internal attribute tags for these Scalix attribute. The types are 1, 8, and 9 , respectively, in the <code>dir.attrs</code> file. If you have defined a custom attribute type of <code>JobTitle</code> in the <code>dir.attrs</code> file, specifying the line: <code>UserDefinedAttributes=%(1)%(JobTitle)</code> displays the Surname and Job Title.

Table 36: Display Parameters and their Descriptions

Parameter	Description
MaxContainerSize	<p>The maximum number of messages that are listed on opening a folder can be configured using the <code>MaxContainerSize</code> setting. This setting can take an integer value between 20 and 32767. The default value is 32767. If the limit is exceeded, a warning message displays.</p> <p>Archiving (or auto-archiving) a folder that contains more than the <code>MaxContainerSize</code> value causes the container-limit warning message to display. Scalix Connect archives those items within the limit of 32767.</p>
LineLength	<p>The maximum length of a line in a plain text message can be specified in the <code>Display</code> section. The format of the option is as follows: <code>LineLength=n</code> Specifies the maximum number of characters in each line of a plain text message, where <i>n</i> can be a value up to 80. For example, <code>LineLength=60</code> ensures that the message text has no more than 60 characters per line. If the <code>LineLength=n</code> entry is missing or invalid then lines default to a maximum of 72 characters. To disable wrapping, add the entry <code>LineLength=0</code>.</p>
TabStops	<p>The <code>TabStops=n</code> specifies the number of spaces for tab stops used by plain text messages, where <i>n</i> must not exceed the <code>LineLength</code> value or 20 (whichever is smaller). The default for <i>n</i> is 4 if the <code>TabStops</code> setting is missing.</p>
ShowCompleteInternetAddress	<p>The <code>ShowCompleteInternetAddress=1</code> option specifies that the entire Internet address is displayed.</p> <p>The <code>ShowCompleteInternetAddress=1</code> causes a Scalix address that contains a DDA type of RFC-822 to be displayed as an Internet address. For example, the address of an incoming message such as <code>chris/linux/dd.RFC-822=cwolfe@pwd.scalix.com</code> displays as <code>chris</code> (by default).</p> <p>If you set the value <code>ShowCompleteInternetAddress=1</code>, Scalix Connect displays the address as <code>cwolfe@pwd.scalix.com</code>.</p> <p><code>ShowCompleteInternetAddress=1</code> takes precedence over <code>ShowMailnodes=1</code>.</p> <p>The behavior of <code>ShowCompleteInternetAddress</code> is also affected by settings in the <code>[Addressing]</code> section. A setting of <code>InternetToOM=1</code> prevents interpretation of the RFC-822 DDA as an Internet address and causes <code>ShowCompleteInternetAddress</code> to be ignored.</p> <p><code>HPMEXTToSMTP=1</code> extends the behavior of <code>ShowCompleteInternetAddress</code> to <code>HPMEXT</code> DDAs.</p>

Recovery Folder Options

Table 37: Recovery Folder Options and their Descriptions

Option	Description
RECOVERY_FOLDER_EXPIRY_TIME=<time_period>	Where <time_period> is the amount of time that deleted items remain in the "Scalix Recovered Items" folder before being removed from the system. The default is seven days (7d). Example settings for this option are: 4d12h (4 days and 12 hours) or 240h (240 hours).

Guidelines for the User Defined Attributes

For the UserDefinedAttributes option to be applied correctly, note the following guidelines:

The [Display] section must also include the following setting:

```
ShowCustomAttributes=1
```

The [Display] section must not include the setting ShowMailnodes=1 since this setting takes precedence over UserDefinedAttributes and results in an address of the following form:

```
Eric Smith / ou1, ou2
```

The UserDefinedAttributes setting can include up to ten entries, including a maximum of six custom attributes.

Any attribute specified in UserDefinedAttributes must also appear in the properties option of the Address Book. If the attribute is not one of the standard X.400 addressing fields or teletext equivalents (Scalix internal format Group 1 and 3), as displayed in the Name/Address Fields and the DDA pages, then you must add it to the custom attribute page through the mapi.cfg [Name Attributes] section.

In other words, any tags in UserDefinedAttributes outside the value ranges 1-23 and 51-68 must also appear in the [Name Attributes] section. For example:

```
[Name Attributes]
Heading=Custom
1=Phone: , 116
2=My Own Data: , myown
[Display]
ShowMailnodes=0
ShowCustomAttributes=1
UserDefinedAttributes=%(2)%(1)%(116)%(myown)
```


[Directories] Parameters

The option in the [Directories] section enables additional directories from the server to be included. The server default Directory is always opened by the MAPI Address Book provider.

Table 38: Directories Parameters and their Descriptions

Parameter	Description
n=directory name	<p>Specifies an additional Directory, where:</p> <ul style="list-style-type: none"> <i>n</i> is a consecutive sequence number (1-20). <i>directory name</i> is the name of the additional directory to be included. Directory names are case-sensitive. <p>For example, to include Directories for your Sales and Overseas departments, specify the following lines:</p> <pre>[Directories] 1=SALES 2=OVERSEAS</pre> <p>Make sure these entries are case sensitive, for example if you have the directory MYOWNONE Shared LOCAL DB config update read modify-self (as shown by <code>omlistdirs</code>), the entry in this section must be:</p> <pre>1=MYOWNONE</pre> <p>You also need to add these directories to the CDA server (<code>omaddcda</code>) and then run <code>omexecdda</code>.</p> <p>Add the <code>-d <directoryname></code> option to enable type-down searching on this directory.</p>

[Name Attributes] Parameters

The option in the Name [Attributes] section enables additional attributes from the Scalix Directory to be included as a name/address properties page and as the Search page of the MAPI client.

Table 39: Name Attribute Parameters and their Descriptions

Parameter	Description
heading= <i>text</i> <i>n</i> = <i>label</i> , <i>tag</i>	<p>This specifies the additional attributes for the Properties and Search pages, where:</p> <ul style="list-style-type: none"> <i>text</i> is the text used as the page Tab heading. You can specify up to 16 characters. <i>n</i> is the sequence number for the custom attribute. You can specify up to six attributes (1-6). <i>label</i> is the text for the label displayed for the attribute on the page. You can specify up to 24 characters. <i>tag</i> is the numeric value of the internal attribute tag for the corresponding Scalix Directory attribute. To see the list of available tags, use the <code>omshowatt -u</code> command. <p>For example, to include three additional custom attributes, specify the following lines:</p> <pre>[Name Attributes] heading = Custom 1=Job Title:, 111 2=Department:, 115 3=Phone:, 116</pre>

[PAW] Parameters

The [PAW] (Personal Assistant Wizard) section in `mapi.cfg` has the following options:

Table 40: PAW Parameters and their Descriptions

Parameter	Description
URL	The URL pointing to the web server for the user, for example: URL= <code>http://zaphod.pwd.scalix.com</code> This option is required for PAW to be available to the user.
AutoLogon	The automatic logon option uses either 1 or 0 as a value. 1 bypasses the web-based logon page and logs on automatically to the PAW home page. 0 does not bypass the web-based logon page and you must enter your username and password. This option is not required. PAW is still available to a user even if this option is not set.
Profile	The profile option determines the language for the PAW application at start-up. For example: Prof= <code>PAW-ENGLISH</code> This option is not required. PAW is still available to a user even if this option is not set.

Restrictions

Before Scalix Connect includes PAW as a menu option (Tools > Personal Administration Wizard), Scalix Connect checks the following requirements:

- A valid Scalix Server profile.
- A valid `mapi.cfg` file with the [PAW] section included in the configuration file.
- If the [PAW] section exists within the `mapi.cfg` file, Scalix Connect checks that the URL is defined for the PAW server.

The PAW option in the Outlook Tools menu is unavailable if any of these requirements are missing or invalid.

User-Specific Configuration Options

A subset of the options used in the file `general.cfg` can be specified for individual users. The options for each user are held in files with the name of that user's Scalix ID number in the directory `~/sys/user.cfg`.

For example, if the Scalix user `Chris Wolf/ny,hq,mis` has a Scalix ID of 103, then options specific to him will be in the file `~/sys/user.cfg/103`.

To get the user's ID number:

- 1 Use the `omshowu` command with the `-G` option, then search for the Internal User ID. For example:

```
omshowu -n "Jane Smith" -G
```

The `user.cfg` directory does not exist by default and must be created. It must be owned by the user Scalix with permissions of 555 (`dr-xr-xr-x`). User files within this directory must be owned by the user Scalix with permissions of 444 (`-r--r--r--`).

See the descriptions in the <Xref_Color>“System-Wide Configuration Options” section for the following options:

- IMAP_AUTOMATIC_MDN=FALSE
- IMAP_BB_FOLDER_PREFIX=#bb
- IMAP_BB_FOLDER_SEPARATOR=/
- IMAP_DELETE_SUBFOLDERS=FALSE
- IMAP_FOLDER_PREFIX=
- IMAP_FOLDER_SEPARATOR=/
- IMAP_IDLE_TIMEOUT=30
- IMAP_LOGLEVEL=0
- IMAP_MDSENT_FLAG=\$MdnSent
- IMAP_MIN_SIZE_ESTIMATE=0
- IMAP_SEARCH_TIMEOUT=0
- IMAP_X_NETSCAPE_URL=

IMAP Client User-specific Options

Table 41: IMAP Client User-Specific Options and their Descriptions

Option	Description
IMAP_CAPABILITIES= <i>capabilities-list</i>	See the description in the <Xref_Color>"System-Wide Configuration Options" section. Note that capabilities you specify here are <i>added</i> to those specified on a system-wide and client-wide basis.
IMAP_LOGFILE=~/.tmp/imap.%h	See the description in the <Xref_Color>"System-Wide Configuration Options" section. Note that, if logging has been enabled in the system-wide or client-specific configuration file, this option has no effect.

UAL Client Interface User-specific Options

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_DEAD_TIMEOUT= <i>number_of_minutes</i>	Specifies the amount of time that Scalix will wait for the next UAL command of any type (as opposed to UAL_IDLE_TIMEOUT, which is triggered by active commands only) from a UAL client before assuming a timeout. Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client. If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command. For serial connections, UAL_DEAD_TIMEOUT is overridden by UAL_SERIAL_DEAD_TIMEOUT. For local UAL clients, UAL_DEAD_TIMEOUT is overridden by UAL_LOCAL_DEAD_TIMEOUT.
UAL_DIR_LIST_SORT_ORDER= <i>list_of_internal_attributes</i>	Specifies the order in which Directory attributes are sorted. The order is specified as a list of internal attribute names with each attribute separated by a /. The internal attribute names, which are numbers for the core Scalix attributes, are listed using the command omshowatt -u.
UAL_DIR_MOD_FULL_NAME=TRUE	Specifies that Full Name Checking is always done on the UAL_CHKLIST, UAL_CHKNAME, UAL_DELENT and UAL_MODENT commands.
UAL_DISALLOW_AUTO_PASSWORD=TRUE	If set, a client cannot sign on to Scalix if the client has explicitly indicated that its password was obtained from a configuration file rather than having been entered interactively by a user. See also UAL_DISALLOW_NON_USER_PASSWORD. The user-specific setting of this option overrides the system-wide setting. Note that this mechanism is not intended to provide a secure indication.

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_DISALLOW_NON_USER_PASSWORD=TRUE	<p>If set, a client cannot sign on to Scalix if the client has <i>not</i> explicitly indicated that its password was obtained interactively from a user. See also UAL_DISALLOW_AUTO_PASSWORD.</p> <p>Note that this option will only work with clients that supply the "password origination status". If a client does not support this element, then it will not be able to sign on even if the password is actually entered interactively by the user.</p> <p>The user-specific setting of this option overrides the system-wide setting.</p> <p>Note that this mechanism is not intended to provide a secure indication.</p>
UAL_DL_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the Distribution List area size limit. in kilobytes. A value of zero (0) means no size limit.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>
UAL_FC_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the Filing Cabinet size limit. The value is set in kilobytes. A value of zero means no size limit.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>
UAL_IDLE_SHUTDELAY= <i>umber_of_minutes</i>	<p>Specifies the additional delay in shutting down a UAL client connection that has timed out.</p> <p>UAL_IDLE_SHUTDELAY is used with UAL_IDLE_TIMEOUT.</p> <p>See also UAL_IDLE_TIMEOUT</p> <p>For serial connections, UAL_IDLE_SHUTDELAY is overridden by UAL_SERIAL_IDLE_SHUTDELAY. For local UAL clients, UAL_IDLE_SHUTDELAY is overridden by UAL_LOCAL_IDLE_SHUTDELAY.</p>

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
<p>UAL_IDLE_TIMEOUT= <i>number_of_minutes</i></p>	<p>Specifies the amount of time that Scalix will wait for the next "active" UAL command from a UAL client before assuming a timeout (PRE-PARE MESSAGE, ATTACH ITEM are examples of active UAL commands, and NEW MESSAGES and LIST ACK are examples of passive UAL commands). Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client. (The sign off can be delayed using UAL_IDLE_SHUTDELAY.)</p> <p>If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command.</p> <p>UAL_IDLE_TIMEOUT is used with UAL_IDLE_SHUTDELAY.</p> <p>For example:</p> <p>If UAL_IDLE_TIMEOUT is set to 30 minutes, and UAL_IDLE_SHUTDELAY is not set, the client is disconnected from the Server 30 minutes after the last active UAL command was issued.</p> <p>If UAL_IDLE_TIMEOUT is set to 30 minutes, and UAL_IDLE_SHUTDELAY is set to 10 minutes, 30 minutes after the last active UAL command is issued, the client displays a dialog box asking if the user wants to retain the connection. This dialog box is displayed for up to the 10 minutes specified by UAL_IDLE_SHUTDELAY.</p> <p>If the user responds within this time with a Yes, that is considered an active UAL command, and the TIMEOUT countdown restarts from the beginning.</p> <p>If the user responds with a No, the connection is closed.</p> <p>If the user does not respond within the 10 minutes, the connection is closed.</p> <p>For serial connections, UAL_IDLE_TIMEOUT is overridden by UAL_SERIAL_IDLE_TIMEOUT. For local UAL clients, UAL_IDLE_TIMEOUT is overridden by UAL_LOCAL_IDLE_TIMEOUT.</p>
<p>SpUAL_INTRAY_SIZE_LIMIT= <i>no_of_kilobytes</i></p>	<p>Sets the In Tray size limit in kilobytes. A value of zero (0) means no size limit. Set Message Store size limits using the <code>omlimit</code> command.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>
<p>UAL_ISO7_FROM_HOST=<i>language</i></p>	<p>This option is the same as UAL_ISO7_HOST except that the character set conversion only occurs when text is passed from Scalix to the client and not when it is passed back to the Server.</p>

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_ISO7_HOST= <i>language</i>	<p>This option allows clients to interoperate with a Scalix Message Store containing IA5 text that uses the ISO-7 extensions.</p> <p>If set, activates the option UAL_ISO7_<i>language</i> for a client using language. IA5 text with the ISO-7 extensions are converted to the ISO8859/1 character set when downloaded to or displayed by the client and conversely, the ISO8859/1 characters are mapped back into IA5 with ISO-7 extensions when entering the Scalix system from the client.</p> <p>Only one instance of this option can be used and the language string must match a string in the <code>~/sys/LangMap</code> file.</p> <p>See also UAL_ISO7_FROM_HOST, UAL_ISO7_TO_HOST, and SR_ISO7_HOST.</p>
UAL_ISO7_TO_HOST= <i>language</i>	<p>This option is the same as UAL_ISO7_HOST except that the character set conversion only occurs when text is passed from the client to Scalix and not when it is passed back to the client.</p>
UAL_LOCAL_DEAD_TIMEOUT= <i>number_of_minutes</i>	<p>Specifies the amount of time that Scalix will wait for the next UAL command of any type (as opposed to UAL_LOCAL_IDLE_TIMEOUT, which is triggered by active commands only) from a local UAL client before assuming a timeout. Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client.</p> <p>If a timeout period is not specified, Scalix assumes the connection to the local UAL client is good regardless of how long it has been waiting for another command.</p> <p>UAL_LOCAL_DEAD_TIMEOUT overrides UAL_DEAD_TIMEOUT. To remove a timeout for local UAL clients that was set using UAL_DEAD_TIMEOUT, set UAL_LOCAL_DEAD_TIMEOUT to 0.</p>
UAL_LOCAL_IDLE_SHUTDELAY= <i>number_of_minutes</i>	<p>Specifies the additional delay in shutting down a local UAL client connection that has timed out.</p> <p>UAL_LOCAL_IDLE_SHUTDELAY is used with UAL_LOCAL_IDLE_TIMEOUT.</p>

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_LOCAL_IDLE_TIMEOUT= <i>number_of_minutes</i>	<p>Specifies the amount of time that Scalix will wait for the next "active" UAL command from a local UAL client before assuming a timeout (PREPARE MESSAGE, ATTACH ITEM are examples of active UAL commands, and NEW MESSAGES and LIST ACL are examples of passive UAL commands). Once the timeout period has been reached, Scalix assumes the serial connection to the client has been lost and signs off the user of the UAL client. (The sign off can be delayed using UAL_LOCAL_IDLE_SHUTDELAY.)</p> <p>If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command.</p> <p>UAL_LOCAL_IDLE_TIMEOUT is used with UAL_LOCAL_IDLE_SHUTDELAY.</p> <p>UAL_LOCAL_IDLE_TIMEOUT overrides UAL_IDLE_TIMEOUT. To remove a timeout for local UAL clients that was set using UAL_IDLE_TIMEOUT, set UAL_LOCAL_IDLE_TIMEOUT to 0.</p>
UAL_LOCAL_IGNORE_PASSWORD= TRUE or FALSE	<p>Specifies whether a password check is made during sign on. Set the option to TRUE to remove the password entry stage from the sign on process. Set the option to FALSE to add the stage back into the sign on process if it has been removed by setting UAL_LOCAL_IGNORE_PASSWORD in the <code>general.cfg</code> file.</p> <p>The sign on will succeed only if the user has logged in using their Scalix mailbox Linux login, and if the user is using a local UAL Client.</p>
UAL_MSTORE_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the overall message store size limit. The value is set in kilobytes. A value of zero means no size limit.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>
UAL_NMP_DELAY= <i>number_of_milliseconds</i>	<p>Specifies the time delay before returning a reply to a command from a UAL client using a TCP/IP Named Pipes connection. By default, there is no time delay, but this can mean the receiving client system can "miss" the reply when it is sent. To overcome this problem set the delay to between 10 and 100 milliseconds.</p>
UAL_NO_DESIGNATE_SIGNON= TRUE or FALSE	<p>Specifies whether the designate sign on feature is used. Set the value to TRUE to remove the designate sign on feature. Set the value to FALSE to add the designate sign on feature if it has been removed by setting UAL_NO_DESIGNATE_SIGNON in the <code>general.cfg</code> file.</p>
UAL_NO_WB_EMPTY=TRUE	<p>Stops a user's Waste Basket being emptied when the user has finished using a UAL client and signs off. If this option is set, use the command <code>omtidyu</code> or <code>omtidyallu</code> to ensure Waste Baskets continue to be emptied regularly.</p>

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_PASSWORD_AGED= <i>IGNORE WARN or ERROR</i>	<p>This option determines the effect of an expired password on a user signing on to Scalix through a client.</p> <p>The default value is <code>ERROR</code>. If the user's password has expired, an error is generated when the user attempts to signon and the signon fails. The signon can only succeed when a valid new password is supplied.</p> <p>If the value is set to <code>WARN</code> and the user's password has expired, the user can sign on using the expired password but a warning message is placed in their In Tray stating that their password has expired and should be changed immediately. (This message appears in the In Tray for the first signon of the day.)</p> <p>If the value is set to <code>IGNORE</code> any user password expiry condition is ignored (a Scalix user will be allowed to signon even if their password has expired.)</p>
UAL_PEND_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the Pending Tray size limit in kilobytes. A value of zero (0) means no size limit. Set Message Store size limits using the <code>omlimit</code> command.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>
UAL_PWD_WARNING_DAYS= <i>days</i>	<p>Activates the mechanism to generate advisory messages to users whose mailbox passwords are due to expire within the period specified by days. The warning message appears as a new message in the user's In Tray for the first signon of the day. Use this option if clients are being used that do not recognize the password expired signon error. These clients cannot signon successfully once the user's password has expired.</p>
UAL_SCK_DELAY= <i>number_of_milliseconds</i>	<p>Specifies the time delay before returning a reply to a command from a UAL client using a TCP/IP Sockets connection. By default, there is no time delay, but this can mean the receiving client system can "miss" the reply when it is sent. To overcome this problem set the delay to between 10 and 100 milliseconds.</p>
UAL_SERIAL_DEAD_TIMEOUT= <i>number_of_minutes</i>	<p>Specifies the amount of time that Scalix will wait for the next UAL command of any type (as opposed to <code>UAL_SERIAL_IDLE_TIMEOUT</code>, which is triggered by active commands only) from a UAL client using a serial connection before assuming a timeout. Once the timeout period has been reached, Scalix assumes the connection to the client has been lost and signs off the user of the UAL client.</p> <p>If a timeout period is not specified, Scalix assumes the serial connection to the UAL client is good regardless of how long it has been waiting for another command.</p> <p><code>UAL_SERIAL_DEAD_TIMEOUT</code> overrides <code>UAL_DEAD_TIMEOUT</code>. To remove a timeout for UAL clients that was set using <code>UAL_DEAD_TIMEOUT</code>, set <code>UAL_SERIAL_DEAD_TIMEOUT</code> to 0. This removes the timeout for all UAL clients using a serial connection.</p>

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_SERIAL_IDLE_SHUTDELAY= <i>number_of_minutes</i>	Specifies the additional delay in shutting down a UAL client serial connection that has timed out. UAL_SERIAL_IDLE_SHUTDELAY is used with UAL_SERIAL_IDLE_TIMEOUT.
UAL_SERIAL_IDLE_TIMEOUT= <i>number_of_minutes</i>	Specifies the amount of time that Scalix will wait for the next "active" UAL command from a UAL client using a serial connection before assuming a timeout (PREPARE MESSAGE, ATTACH ITEM are examples of active UAL commands, and NEW MESSAGES and LIST ACL are examples of passive UAL commands). Once the timeout period has been reached, Scalix assumes the serial connection to the client has been lost and signs off the user of the UAL client. (The sign off can be delayed using UAL_SERIAL_IDLE_SHUTDELAY.) If a timeout period is not specified, Scalix assumes the connection to the UAL client is good regardless of how long it has been waiting for another command. UAL_SERIAL_IDLE_TIMEOUT is used with UAL_SERIAL_IDLE_SHUTDELAY. UAL_SERIAL_IDLE_TIMEOUT overrides UAL_IDLE_TIMEOUT. To remove a timeout for UAL clients that was set using UAL_IDLE_TIMEOUT, set UAL_SERIAL_IDLE_TIMEOUT to 0. This removes the timeout for all UAL clients using a serial connection.
UAL_SIGNON_ALIAS=YES, ONLY or NONE	Specifies whether aliases are used for sign on. Any UAL_SIGNON_ALIAS entries in user.cfg take precedence over the UAL_SIGNON_ALIAS entry in general.cfg. (This enables you to set a default use of aliases in general.cfg and then set overrides for specific users in user.cfg.) The YES option means aliases can be used to sign on with--users can also continue to use their Personal Name if they want to. The ONLY option means the aliases only can be used to sign on with--the Personal Name cannot be used any more. The NONE option means aliases are not used. UAL_SIGNON_ALIAS is used with UAL_SIGNON_ALIAS_CONFIG and UAL_USE_SIGNON_ALIAS.
UAL_SIZE_ON_RECEIPT=FALSE	Specifies whether a user whose message store components exceed their configured limits is prevented from receiving messages. When this option is set to FALSE (the default), users are not prevented from receiving messages even if the size of their message store component is greater than its configured limit.

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_SIZE_ON_SEND=FALSE	<p>Specifies whether a user whose message store components exceed their configured limits is prevented from sending messages. When this option is set to <code>TRUE</code>, then message delivery rules can be implemented that limit a user's ability to send messages. These rules utilize the <code>OMLIMIT-EXCEEDED</code> message attribute filter. When this option is set to <code>FALSE</code> (the default), then rules based on the <code>OMLIMIT-EXCEEDED</code> filter have no effect.</p>
UAL_TRACE_FILE= <i>file_specification</i>	<p>The name stub of the file to which UAL trace information is logged (UAL logging must be enabled). <code>%p</code> in the <i>file_specification</i> is replaced with the PID of the UAL process, <code>%s</code> by the notification session-ID, and <code>%u</code> by the Scalix UID. A leading <code>~</code> represents the Scalix home directory. Note that the existing log file is overwritten. Without a leading <code>~</code> or <code>/</code> character, the file(s) are created in the <code>~scalix/tmp</code> directory. UAL trace output is enabled by using the <code>UAL_TRACE_LEVEL</code> option. The default value is <code>OM%u</code>.</p> <p>The substitutions in the log file name allow log files to be created on a per-ual-process, per-notif-session, or per-user basis. This allows MAPI and Scalix Web Access sessions that use concurrent UAL sessions to be traced without any loss of data. UAL client session must be restarted to enable the changes to this option.</p> <p>Example: <code>UAL_TRACE_FILE=ual.%u.%p</code> creates log files in the <code>~scalix/tmp</code> directory with a stub of <code>ual.user_id.pid</code>. For example: <code>ual.102.1773U.log</code> <code>ual.102.1773U.f0001</code></p> <p>Example: <code>UAL_TRACE_FILE=/tmp/ual-logs/%u.%p</code> creates log files in the <code>/tmp/ual-logs</code> directory with a stub of <code>user_id.pid</code>. For example: <code>/tmp/ual-logs/102.1773U.log</code> <code>/tmp/ual-logs/102.1773U.f0001</code></p> <p>In this example, the <code>/tmp/ual-logs</code> directory must be created before any trace files can be written.</p> <p>See <code>UAL_TRACE_LEVEL</code> for more information.</p>

Table 42: UAL Client Interface User-Specific Options and their Descriptions

Option	Description
UAL_TRACE_LEVEL= <i>trace_level</i>	<p>Activates UAL Client Interface tracing. The trace files are placed in the <code>~/tmp</code> directory. If this directory cannot be found, they are placed in the <code>/tmp</code> directory. File names begin with <code>OMuser-no</code>, where <i>user-no</i> is the Scalix user number, and end according to the trace level set. If you require several different kinds of trace information, add the numbers for the levels you require and set the entry to the total.</p> <p>0: No tracing. The default.</p> <p>1: Raw (unformatted) command/reply tracing (file name: <i>nameN.trc</i>).</p> <p>2: Command statistics.</p> <p>4: Message Store file name mapping. No trace file. The subject of an item listed or displayed in the client is replaced by its corresponding Message Store file name.</p> <p>8: Full tracing of command/reply and file transfer data. This can be used to rerun a session (file name: <i>nameU.log</i> and <i>nameU.fnnnn</i>).</p> <p>16: Raw (unformatted) command/reply tracing and file transfer data (file name: <i>nameN.trc</i>).</p> <p>Also use this entry to set Event Log logging on the Server for the client. Set the entry to the required Event Log logging level multiplied by 100.</p>
UAL_LINUX_PASSWORD=TRUE or FALSE	<p>Specifies whether the user uses their Linux password instead of their Scalix password when signing on. <code>TRUE</code> sets Scalix to use the Linux password, and <code>FALSE</code> (the default) sets Scalix to use the Scalix password.</p>
UAL_USE_SIGNON_ALIAS=FALSE or TRUE	<p>Specifies whether the alias is used after sign on. If you set <code>UAL_USE_SIGNON_ALIAS</code> to <code>FALSE</code>, the UAL client reverts to using the user's Personal Name for the remaining time the user is signed on (the alias or Personal Name is used on the "Creator" part of a message). If you set <code>UAL_USE_SIGNON_ALIAS</code> to <code>TRUE</code>, the alias is used for the remaining time the user is signed on.</p> <p><code>UAL_USE_SIGNON_ALIAS</code> is used with <code>UAL_SIGNON_ALIAS</code> and <code>UAL_SIGNON_ALIAS_CONFIG</code>.</p>
UAL_WB_SIZE_LIMIT= <i>no_of_kilobytes</i>	<p>Sets the Waste Basket size limit in kilobytes. A value of zero (0) means no size limit. Set Message Store size limits using the <code>omlimit</code> command.</p> <p>Once you configure a limit, you must then configure the sanctions to be applied to those users who exceed it. If you do not configure sanctions, users can exceed their limits without effect.</p>

Offline Folder Synchronization Options (Outlook Clients)

Table 43: Offline Folder Synchronization Options and their Descriptions

Option	Description
OFS_LOG_SIZE_LIMIT= <i>kilobytes</i>	<p>Specifies, in kilobytes, the maximum size of the folder synchronization change log. Set a value between 20 and 10,000 KB. The default is 100 KB.</p> <p>When the size of a change log exceeds this value, the older entries can be deleted when the change log file is compacted.</p> <p>Use this option to prevent change logs from growing too large when logging is activated but synchronization is not taking place for any reason. Take care when you set this value, since removal of any valid entries will cause the entire folder to be resynchronized.</p> <p>A value you set in <code>general.cfg</code> can be overridden on a per-user basis by setting a lower value for this option in the relevant user-specific configuration file.</p>
OFS_LOG_AGE_LIMIT= <i>days</i>	<p>Specifies, in days, the maximum age of entries in the folder synchronization change log. Set a value between 1 and 18,000 days. The default is 90 days.</p> <p>When the age of a change log entry exceeds this value, it can be deleted when the change log file is compacted.</p> <p>Use this option to prevent change logs from growing too large when logging is activated but synchronization is not taking place for any reason. Take care when you set this value, since removal of any valid entries will cause the entire folder to be resynchronized.</p> <p>A value you set in <code>general.cfg</code> can be overridden on a per-user basis by setting a lower value for this option in the relevant user-specific configuration file.</p>

Scalix Command Line Reference Guide

This chapter summarizes the following collections of Scalix CLI commands.

- Access Control List Commands232
- Audit Log Commands232
- Client Directory Access (CDA) Commands232
- Configuration and Installation Commands233
- Directory Commands233
- Directory Relay Server Commands233
- Directory Synchronization Commands234
- Error Manager Commands235
- Event Log Commands235
- Internet Address Commands235
- Internet Mail Gateway Commands235
- LDAP Commands236
- Mailbox Access Commands236
- Mailnode Commands236
- Message Store Commands237
- Miscellaneous Commands237
- Public Distribution List (Group) Commands238
- Public Folder Commands238
- Routing Table Commands240
- Service, Queue and Daemon Commands240
- System Configuration and Maintenance Commands240
- User Entry and Management Commands241

Introduction

After you log into your Scalix server, you can enter Scalix commands directly in a Linux terminal window, using the full Scalix command line interface. Scalix commands are constructed with a prefix of “om” followed by extensions that define the action and the object. For example, the command to add a user is **omaddu** (followed by the needed extensions).

Because interpretations of command characters vary per shell program, you might need to use escape sequences for some command elements. For example, if you enter parentheses while using the command **omsearch**, some shell programs might require that the parentheses be escaped with backslashes (\) before the shell program can interpret the command correctly.

Note

Some Scalix commands can be used only by a Scalix Administrator with system administration privileges (for example, a root user).

The following table lists the location of additional information (manual pages) about command formats and matching rules not covered in this chapter.

Table 1: Man Page Locations

Command	Use
omaddress	Enter O/R Address, mailnodes, and pattern matching rules
omattribs	Attribute input and output formats
omdiratt	Formatting of attribute definition files
scalix	Services, queues, daemons, and commands

Starting with the following page, this publication provides a complete catalog of all Scalix CLI commands, listing them in common categories, and providing a simple description of how each command is used. For more information about syntax and extensions, see the man page for each command.

Access Control List Commands

Table 2: ACL Commands and their Uses

Command	Uses
omaddacl	Add an Access Control List
omaddacln	Add capabilities for users to an Access Control List
omchkacln	Check capabilities of a user in an Access Control List
omdelacl	Delete an Access Control List
omdelacln	Delete capabilities for users from an Access Control List
ommodacln	Modify capabilities for users in an Access Control List
omshowacl	Show the contents of an Access Control List

Audit Log Commands

Table 3: Audit Log Commands and their Uses

Command	Uses
omconfaud	Configure Audit Log logging levels
omshowaud	Show Audit Log logging levels

Client Directory Access (CDA) Commands

Table 4: Client Directory Access Commands and their Uses

Command	Uses
omaddcda	Add a Directory to the CDA Server configuration
omdelcda	Delete a Directory from the CDA Server configuration
omexeccda	Force the CDA Server to process a Directory immediately
ommodcda	Modify the CDA Server configuration for a Directory
omshowcda	Show the CDA Server configuration for a Directory

Configuration and Installation Commands

Table 5: Installation and Configuration Commands and their Uses

Command	Uses
omcptree	Copy or refresh a Directory hierarchy
omdelom	Delete a Scalix instance
ommakeom	Make a Scalix instance
ompatchom	Update a Scalix instance
omredirtcp	Redirect socket connections to the correct Scalix system in a multi-system environment

Directory Commands

Table 6: Directory Commands and their Uses

Command	Uses
omaddent	Add one or more entries to a Directory
omdelent	Delete one or more entries from a Directory
omdiropt	Optimize a Directory
omdoptall	Optimize all Directories
omfmtent	Format Directory and address attributes
omlistdirs	List Directories
ommoddir	Modify a Directory
ommodent	Modify a Directory entry
omremdir	Delete a Directory
omsearch	Search a Directory
omshowatt	Show available attribute types

Directory Relay Server Commands

Table 7: Directory Relay Server Commands and their Uses

Command	Uses
omresetmn	Reset a mailnode mapping file
omaddmnmp	Add entry to a mailnode mapping file
ommodmnmp	Modify entry in a mailnode mapping file

Table 7: Directory Relay Server Commands and their Uses

Command	Uses
omdelmnp	Delete entry in a mailnode mapping file
omshowmnp	List entries in a mailnode mapping file

Directory Synchronization Commands

Support for multiple Scalix servers and directory synchronization is available only in Scalix Enterprise Edition. For more information, see "About Scalix Product Editions".

Table 8: Directory Synchronization Commands and their Uses

Command	Uses
omaddds	Add a Directory Synchronization agreement
omdelds	Delete a Directory Synchronization agreement
omlistds	List Directory Synchronization agreements
ommodds	Modify a Directory Synchronization agreement
omresyncds	Resynchronize a Directory
omshowds	Show details of a Directory Synchronization agreement

Error Manager Commands

Table 9: Error Commands and their Uses

Command	Uses
omconfenu	Configure an Error Manager
omshowenu	Show the address of the Error Manager

Event Log Commands

Table 10: Event Log Commands and their Uses

Command	Uses
omconflvl	Configure Event Log logging levels
omshowlog	Show the Event Log
omshowlvl	Show Event Log logging levels

Internet Address Commands

Table 11: Internet Address Commands and their Uses

Command	Uses
omaddiam	Add entry to the Internet address mapping file
omdeliam	Delete an entry from the Internet address mapping file
omgeniamods	Generate a script to modify Internet addresses.
ommodiam	Modify a mapping between OR address and Internet address
ompreviewia	Preview the automatically generated Internet address
omshowiam	List mappings between OR addresses and Internet addresses

Internet Mail Gateway Commands

Table 12: Internet Mail Gateway Commands and their Uses

Command	Uses
omconfux	Configure the Internet Mail Gateway
omshowux	Show the configuration of the Internet Mail Gateway

LDAP Commands

Table 13: LDAP Commands and their Uses

Command	Uses
omldapadd	Add one or more entries to an LDAP Directory
omldapdelete	Delete one or more entries from an LDAP Directory
omldapmodify	Modify an LDAP Directory entry
omldapmoddn	Modify the DN of an LDAP entry
omldapsearch	Search an LDAP Directory

Mailbox Access Commands

Table 14: Mailbox Access Commands and their Uses

Command	Uses
omdelete	Delete a message
omlist	List messages
omlogoff	Terminate an omLogon connection to Scalix
omlogon	Obtain a connection to Scalix
omnew	List newly arrived messages
omread	Read a message
omsend	Send a message

Mailnode Commands

Table 15: Mailnode Commands and their Uses

Command	Uses
omaddmn	Add a mailnode
omdelmn	Delete one or more mailnodes
ommodmn	Modify a mailnode
omshowmn	List local mailnodes

Message Store Commands

Table 16: Message Store Commands and their Uses

Command	Uses
omcontain	Manipulate containers in the Message Store
omcpinu	Copy a user's Message Store data from a file
omcpoutu	Copy a user's Message Store data to a file
omdosur	Create a data file for restoring a single user
omdref	Convert a Scalix DirectRef into a readable description of the item represented, including Message Store item hierarchy
omdumpis	Write Item Structure database to standard output
omgetsur	Get files from an archive
omlimit	Set Message Store size limits globally or for a user
omnewis	Create an empty database
omprepsur	List files required for single user restore
omscan	Scan, report, and repair Scalix data inconsistencies
omshowis	Display the date omupdtis was last run
omsnoop	Report on potential Message Store conversion problems
omsuspend	Halt all client activity temporarily
omtidyallu	Delete items from the Message Store
omtidyu	Delete items from the Message Store for an individual and search nested folders for an item to delete
omupdtis	Read Item Structure log entries and update the database
tfbrowse	Convert between Scalix transaction file format and textual format

Miscellaneous Commands

Table 17: Miscellaneous Commands and their Uses

Command	Uses
ombconv	Convert a numeric value into a variety of numeric bases
ombprint	Print messages in batch mode
omenquire	Enquire about Scalix system status and report the results
omsolve	Display solutions to an error message

Public Distribution List (Group) Commands

Table 18: PDL Commands and their Uses

Command	Uses
omaddpdl	Add a Public Distribution List
omaddpdln	Add an entry to a Public Distribution List
omdelpdl	Delete one or more Public Distribution Lists
omdelpdln	Delete one or more entries from a Public Distribution List
ommodpdl	Modify a Public Distribution List
ommodpdln	Modify Public Distribution List entries
omshowpdl	List Public Distribution Lists
omshowpdln	List entries in a Public Distribution List
omaddaci	Add an Access Control Information member
omchkaci	Check Access Control Information capabilities for a user
omdelaci	Delete an Access Control Information member
ommodaci	Modify an Access Control Information member
omshowaci	Show the contents of Access Control Information

Public Folder Commands

Public folders can be accessed only by Premium users, using an IMAP client or SWA. For more information, see "About Scalix Product Editions".

Table 19: Public Folder Commands and their Uses

Command	Uses
omaddbb	Add a top-level Public Folder
omdelbb	Delete a top-level Public Folder
omlistbbs	List top-level Public Folders
ommaintbb	Maintain top-level Public Folders
ommodbb	Modify the subject of a top-level Public Folder
omshowbb	Show details of a top-level Public Folder
omaddbbsa	Add a Public Folder Synchronization agreement
omdelbbsa	Delete a Public Folder Synchronization agreement
omlistbbsa	List Public Folder Synchronization agreements

Table 19: Public Folder Commands and their Uses

Command	Uses
<code>ommodbbsa</code>	Modify a Public Folder Synchronization agreement

Routing Table Commands

Table 20: Routing Table Commands and their Uses

Command	Uses
omaddrt	Add a route
omdelrt	Delete a route
ommodrt	Modify a route
omshowrt	List routes and show how an address is routed

Service, Queue and Daemon Commands

Table 21: Service, Queue and Daemon Commands and their Uses

Command	Uses
omisoff	Check Scalix services are off
omoff	Stop one or more services
omon	Start one or more services
omrc	Start Scalix
omreset	Reset status of services or remove Scalix
omresub	Resubmit messages
omresubdmp	Resubmit messages processed by the Archive Server
omsetsvc	Display the status of a service in detail; configure auxiliary processes
omshut	Stop Scalix
omstat	List Scalix daemons

System Configuration and Maintenance Commands

Table 22: System Configuration and Maintenance Commands and their Uses

Command	Uses
omcheck	Check Scalix file permissions and ownership
omcnvinst	Configure converter, language, and character set files
ommon	Monitor the operation of Scalix
omstat	Show the status of the system
omvers	List version numbers of all binaries and scripts

User Entry and Management Commands

Table 23: User Entry and Management Commands and their Uses

Command	Uses
omaddu	Add a user
omadmidp	Configure system IDs for use by Scalix users
omconfpwd	Configure password controls
omdelu	Delete one or more users
ommoddl	Modify distribution list entries and auto-action addresses
ommodu	Modify a user
omshowpwd	Show password controls
omshowu	List users or display details about a specific user

Glossary

Some terms and acronyms in this manual may be unfamiliar to users. Here are some terms and definitions that are specific to the Scalix product and the Linux platform.

Table 1: Some Common Terms and their Definitions

Address Directories	In Scalix terminology, the address directories are databases that clients use to look up names and addresses. Scalix directories can hold addresses of both Scalix and non-Scalix users, and other information that an administrator can configure such as job titles and phone numbers. Directories can be searched by any number of attributes.
Management Console or SAC	The Scalix Management Console (SAC) is a browser-based application that enables most day-to-day system administration tasks on a Scalix messaging system through an easy-to-use GUI. It is a separate component of Scalix that users can access with any approved browser on either Microsoft Windows or Linux workstations. SAC provides efficient access to a wide range of Scalix server options, including user account management, starting and stopping server services, administering queues, public distribution list or group management, and changing low-level server configuration settings. It also provides system monitoring to assess the status of processes and resources.
ADUC	(Active Directory Users and Computers) ...
Authentication Identifier	The Scalix system has several ways of identifying users for different purposes: Display names, personal names, authentication IDs and Internet addresses. The display name (also known as a "common name") is used in Outlook and other clients as the "displayed" address. It can serve authentication purposes and determines the sort order in the Outlook address book. Authentication IDs support the concept of a separate login name and allow for integration with external authentication systems that may have their own naming rules. Internet addresses are SMTP addresses of the form name@domain. Personal names are used for internal addressing of email and are sometimes referred to as "X.400 addresses," "OpenMail addresses" or "ORN (originator Recipient name)."
Bulletin Board	In Scalix terminology, a bulletin board is a set of public folders where members can share files, ideas, documents and more. They are a shared area in the Scalix message store.
Clam AV	An open source freeware program that protects against viruses.
Community Edition	The free, single-server, unlimited-use version of the Scalix product. Does not include advanced groupware and collaboration functionality.

Table 1: Some Common Terms and their Definitions

DDR	
Display Names vs User Names vs Personal Names vs authentication ID vs Internet address	The Scalix system has several ways of identifying users for different purposes: Display names, personal names, authentication IDs and Internet addresses. The display name (also known as a "common name") is used in Outlook and other clients as the "displayed" address. It can be used for authentication purposes and determines the sort order in the Outlook address book. Authentication IDs support the concept of a separate login name and allow for integration with external authentication systems that may have their own naming rules. Internet addresses are SMTP addresses of the form name@domain. Personal names are used for internal addressing of email and are sometimes referred to as "X.400 addresses," "OpenMail addresses" or "ORN (originator Recipient name)."
Enterprise Edition	The company's flagship product, which includes multi-server support, unlimited number of Standard users, any number of Premium users, the full complement of Scalix advanced capabilities, and a wide variety of technical support options.
Gateway	Gateways are a way of passing messages out of the Scalix network to different mail environments. The gateway converts outgoing messages from a Scalix format to a format that external services can use to do send processes, and later to a format that target environments can receive such as an SMTP address. Scalix comes with a standard SMTP gateway that converts Scalix-formatted messages to SMTP and vice-versa. This SMTP gateway is called the Unix Mail Gateway or Internet Mail Gateway.
Groups and PDLs	In Scalix terminology, the terms "group" and "PDL" are used interchangeably to mean a group of people organized into a mailing list. PDLs can contain both local and remote users, and can contain nested PDLs.
IMAP	(Internet Message Access Protocol) A standard interface between an e-mail client program and the mail server. In Scalix, the iMAP4 server enables a client to: Access, list, read, and delete items from inboxes, filing cabinets and public folders; read parts of a message without downloading the entire thing, keep a record of which messages have been read, and update messages on the server from a client. IMAP extensions also provide for calendaring and contact management.
Internet Domains vs mailnodes	Mailnodes have no direct relationship to Internet domains. However, you can set up rules so that when a user is created on a mailnode, Internet address generation kicks in and creates an Internet address for the user. You can map multiple mailnodes to the same Internet domain name.
LDAP	(Lightweight Directory Access Protocol) A protocol used to access a directory listing. In Scalix, the LDAP server is a daemon process based on a client/server model that provides an interface to enable LDAP clients to store and retrieve data from a Scalix directory without any information about the operation of Scalix. It provides LDAP clients access to shared Scalix directories that do not have an associated password.
LVM	(Logical Volume Manager) Used for backing up Scalix directories.

Table 1: Some Common Terms and their Definitions

mailnodes	A logical structure used to organize users into administrative groupings. For example, some companies organize their email users by work group whereas others break their users down by employment status. Each Scalix server is associated with a single mailnode created during installation. After installation, you can use the Management Console to create additional mailnodes on a server, including customizing any new mailnodes with a specific Internet address or domain name.
MAPI	(Mail API) A programming interface from Microsoft that enables a client application to send to and receive mail from Exchange Server or a Microsoft Mail (MS Mail) messaging system. Microsoft applications such as Outlook, the Exchange client and Microsoft Schedule use MAPI.
Message Store	The message store is a collection of flat Linux files held in file system directories on the Scalix server. It holds new messages received as well as messages in transit. For clients that use the message store (server-based clients), it also holds old messages that are files for reference in folders, copies of outgoing messages, draft messages, private distribution lists, personal information such as calendaring, tasks, bulletin boards, public folders and more.
Mx Records	Mail exchanger records inside DNS servers. These decide which server is responsible for dealing with mail or domain DNS actions.
OpenMail	The original technology, licensed from Hewlett Packard, upon which the Scalix system is based.
O/R or Originator/Recipient Address	An attribute list that distinguishes one user, or distribution list, from another and defines the user's point of access to the message handling system or the distribution list's location.
PAM	(Pluggable Authentication Modules). A standard library in Linux that connects applications that require authentication with shared library modules interfacing with authentication mechanisms.
PDL	In Scalix terminology, the terms "group" and "PDL" are used interchangeably to mean a group of people organized into a mailing list. PDLs can contain both local and remove users, and can contain nested PDLs.
Personal Name	The Scalix system has several ways of identifying users for different purposes: Display names, personal names, authentication IDs and Internet addresses. The display name (also known as a "common name") is used in Outlook and other clients as the "displayed" address. It can be used for authentication purposes and determines the sort order in the Outlook address book. Authentication IDs support the concept of a separate login name and allow for integration with external authentication systems that may have their own naming rules. Internet addresses are SMTP addresses of the form name@domain. Personal names are used for internal addressing of email and are sometimes referred to as "X.400 addresses," "OpenMail addresses" or "ORN (originator Recipient name.)"
POP	(Post Office Protocol) A standard interface between an e-mail client program and the mail server. The Scalix POP3 server enables clients to list, read and delete items from the inbox area of the Scalix message store. The Scalix POP3 server does not provide access to any other areas of the message store such as public folders.

Table 1: Some Common Terms and their Definitions

Premium Users	Scalix has two levels of access and usage: Premium and Standard. Premium users have access to the full benefits and functionality of the Scalix email and calendaring system. Standard users gain access to a subset of Scalix functionality including email, personal calendar and contacts through Scalix Web Access and Novell Evolution as well as email access using POP/IMAP clients.
Realm	
SAC	The Scalix Management Console (SAC) is a browser-based application that enables most day-to-day system administration tasks on a Scalix messaging system through an easy-to-use GUI. It is a separate component of Scalix that users can access with any approved browser on either Microsoft Windows or Linux workstations. SAC provides efficient access to a wide range of Scalix server options, including user account management, starting and stopping server services, administering queues, public distribution list or group management, and changing low-level server configuration settings. It also provides system monitoring to assess the status of processes and resources.
Scalix Connect	A MAPI application that enables the use of the Outlook client interface and all of its functionality.
Sendmail	An SMTP-based message transfer agent (MTA) that runs under Unix and Linux. It is the mail transfer process used inside the Scalix system.
SSL	
Small Business Edition	A version of the Scalix system that targets organizations getting started with a commercial version of Scalix that do not have the higher end requirements of Enterprise Edition. It is functionally equivalent to Enterprise Edition except that it allows only single-server installations
SmartHost	
Spam Assassin	An open source freeware program that filters spam.
Standard Users	Scalix has two levels of access and usage: Premium and Standard. Premium users have access to the full benefits and functionality of the Scalix email and calendaring system. Standard users gain access to a subset of Scalix functionality including email, personal calendar and contacts through Scalix Web Access and Novell Evolution as well as email access using POP/IMAP clients.
SWA	Scalix Web Access, the browser-based email, calendar, contacts and public folders client that comes with any Scalix installation.
Transports	Transports are services that Scalix uses to pass Scalix format messages to other Scalix services. Scalix uses Sendmail and SMTP formatted messages to send messages between servers in the Scalix network, but other connections can be written. The transport service on the Scalix server is called the Sendmail Interface.
UAL	(User Access Layer) A proprietary Scalix protocol that enables communication between clients and the Scalix server.
WAP	(Wireless Application Protocol) A standard for providing cellular phones, pagers and other handheld devices with secure access to e-mail and text-based Web pages.