



Mirapoint XSAN Configuration Guide

Release 4.1
March 2009
Part number 010-00610f

This manual supports Messaging Operating System (MOS) releases 3.10.3 and 4.1.3, and later MOS releases until replaced by a newer edition.

The Mirapoint Software and Mirapoint documentation are Copyright © 1998-2009 Mirapoint Software, Inc. All Rights Reserved. You may not print, copy, reproduce, modify, distribute or display this work in hard copy, electronic, or any other form, in whole or in part, by any electronic, mechanical, or other means, without the prior written consent of Mirapoint Software, Inc., except that you are permitted to make one copy for archival purposes only in connection with the lawful use and operation of this software.

Mirapoint, RazorGate, and the Mirapoint logo are registered trademarks of Mirapoint Software, Inc. Mirapoint Message Server, Mirapoint Directory Server, Mirapoint Operations Console, RazorSafe, DirectPath, WebMail Direct, WebCal Direct, and GroupCal Direct are trademarks of Mirapoint Software, Inc.

Mirapoint integrates third party software programs within the Mirapoint Software, which are subject to their own license terms. If the user does not agree to abide by the applicable license terms for the integrated third party software programs as defined by the Mirapoint Software License Agreement, then you may not install or operate the Mirapoint Software. These software license agreements can be viewed at <http://www.mirapoint.com/licenses/thirdparty/eula.php>.

Portions of this product are Copyright © 1982, 1986, 1989, 1991, 1993 the Regents of the University of California. All Rights Reserved.

Portions of this product are Copyright © 2008 Red Hat, Inc. All Rights Reserved. The “Red Hat” trademark and the “Shadowman” logo are registered trademarks of Red Hat, Inc. in the U.S. and other countries.

Portions of this product are Copyright © 1997, 1998 FreeBSD, Inc. All Rights Reserved.

Portions of this product are Copyright © 1996-1998 Carnegie Mellon University. All Rights Reserved.

Portions of this product are Copyright © 1997-1998 the Apache Group. All Rights Reserved.

Portions of this product are Copyright © 1987-1997 Larry Wall. All Rights Reserved. See <http://www.perl.org>.

Portions of this product are Copyright © 1990, 1993-1997 Sleepycat Software. All Rights Reserved.

This software is derived in part from the SSLava™ Toolkit, which is Copyright © 1996-1998 by Phaos Technology Corporation. All Rights Reserved.

Portions of this product are Copyright © 1998, 1999, 2000 Bruce Verderaime. All Rights Reserved.

Portions of this product are Copyright © 1999-2003 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Macintosh is a trademark of Apple Computer, Inc.

Windows, Outlook, Exchange, and Active Directory are trademarks of Microsoft Corporation.

Java and Solaris are trademarks of Sun Microsystems, Inc.

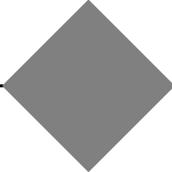
Linux is a registered trademark of Linus Torvalds.

All other trademarks are the property of their respective owners.

OTHER THAN ANY EXPRESS LIMITED WARRANTIES THAT MIRAPOINT PROVIDES TO YOU IN WRITING, MIRAPOINT AND MIRAPOINT'S LICENSORS PROVIDE THE SOFTWARE TO YOU "AS IS" AND EXPRESSLY DISCLAIM ALL WARRANTIES AND/OR CONDITIONS, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL MIRAPOINT'S LICENSORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY (INCLUDING NEGLIGENCE OR OTHER TORT), ARISING IN ANY WAY OUT OF YOUR USE OF THE SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF DAMAGES. MIRAPOINT'S LIABILITY SHALL BE AS LIMITED IN THE LICENSE AGREEMENT.

To view license agreements, including the Mirapoint Software License Agreement, Red Hat End User License Agreement, and Open LDAP Public License, see <http://www.mirapoint.com/licenses/thirdparty/eula.php>.

[2008.02.25.EULA]



Contents

- Preface 5
 - About Mirapoint Documentation5
 - Finding More Documentation.....5
 - Getting Customer Support5
 - Typographic Conventions6

- 1
- Mirapoint SAN Configurations 7
 - Recommended Configurations.....7
 - Terminology7

- 2
- Configuring Mirapoint SAN 9
 - Hardware Setup9
 - Site Requirements9
 - Filer Power and Cabling9
 - Power On10
 - XSAN Basic Setup.....11
 - Network Setup.....11
 - License.....13
 - XSAN OS Upgrade13
 - SAN Network Setup14
 - Fibre Channel Protocol14
 - FCP Configuration.....14
 - Cluster Failover (CF) Mode Setup.....16
 - Igroup Initiator Group.....17
 - XSAN Storage Setup18
 - Disk18
 - Storage Allocation19
 - Aggregate, Aggr20



Volume and LUN Mapping for Mirapoint	21
--	----

3

Administering Mirapoint SAN	25
Data Protection	25
Snapshot and Backup	25
Snaprestore and Selective Restore	25
Auto Support	27
Shelf Expansion	27
XSAN Filer Failover	27
Operational Procedures	27
Adding LUN to a Message Server	27
Adding New Message Server to Existing N+1 Cluster	30
Backup a Message Server	34
Restore a Message Server for Mailbox Restore	35
Troubleshooting	35
Unexpected Loss of LUN	35
Dual HBA Configuration	36
Multipath Configuration	36
Factory Reset	36
High Level Reset	36
Complete Factory Reset	36

A

Usable Space	37
--------------------	----



Preface

About Mirapoint Documentation

Documentation for all Mirapoint products is available through the Mirapoint Technical Library (MTL) on the Customer Support website:

<http://support.mirapoint.com/secure/MTL/MTL>

The MTL provides the Hardware and Software documentation for all supported Mirapoint releases and appliances, and the Support Knowledge Base.

The Support site is accessible to all customers with a valid Support Contract. If you have a valid Mirapoint Support Contract but need a Support login ID, email support-admin@mirapoint.com.

Finding More Documentation

For more documentation on RAID storage and the SAN controller, called the Filer in this manual, see this website:

<http://now.netapp.com/NOW/knowledge/docs>

Access requires a login and password. You can obtain an account automatically by supplying the serial number of your Filer.

Getting Customer Support

If you experience problems with your system, contact Customer Support by e-mail or by telephone:

E-mail: support@mirapoint.com

Telephone: 1-877-MIRAPOINT (1-877-647-2764)

When contacting Customer Support, please be prepared with the following information about your system: MOS version (**Version** command in the CLI), the host ID (**License Hostid** command), serial number (**Model Get Serial** command), and hardware model (**Model Get Chassis** command, and label on the bezel).

Typographic Conventions

Table 1 Typefaces Used in This Book

Typeface	Use	Example
Sabon Roman	Ordinary text	The e-mail server organizes mailboxes hierarchically.
Sabon Bold	Definitions; also screen elements such as page titles, option labels, and commands	A folder is a container that stores e-mail messages. Use the Ldap Set command to turn on auto-provisioning.
<i>Sabon Italic</i>	Emphasis and book titles	Specify <i>at least two</i> DNS servers. See the <i>Mirapoint Administrator's Protocol Reference</i> .
Typewriter	Screen display text and command names	Enter your IP address:
Typewriter Bold	Text that you type exactly as shown	user list
Typewriter Italic	Placeholders for variables you provide	<code>your_IP_address</code>

Mirapoint SAN Configurations

A Storage Area Network (SAN) is a subnetwork of shared storage devices. A SAN consists of switches, high-speed network connectors, and disk arrays for storage. One big advantage of SAN is that Mirapoint software can mediate storage so that any system outage in a failover cluster can be quickly replaced by a standby system using the same storage area as the failed system.

Recommended Configurations

Mirapoint provides multiple XSAN storage models in various configurations. For the list of available Mirapoint XSAN configurations, see the current *XSAN Configurations* article in the Knowledge Base on the Mirapoint Support website.

All of the XSAN configurations use RAID 6 DP (double parity) format, which enables the array to continue operating with two failed disks instead of one.

Available options include:

- ◆ An additional shelf is available to extend storage by either 300GB or 450GB per disk. The shelf options are listed in the Mirapoint sales catalog under the names XSAN300-c7 and XSAN300-c11, respectively.
- ◆ Cabinets are available, and strongly recommended.

Terminology

The following terms are used throughout this document:

Cluster More than one Filer is attached to the SAN fabric. For example, the 3020C is clustered with two Filers, whereas the 3020 is not.

DP Double parity, a form of RAID 6 that can handle two disk failures, instead of just one, without loss of data.

Fabric Refers to the switch and Fibre Channel connectors together.

Filer The SAN controller (or I/O server) that manages filesystems and volumes. The Filer mediates between RAID storage and the SAN fabric.

Head A Mirapoint Message Server with processor, memory, and system software supporting user mailboxes and inbound/outbound SMTP.

- N+1 Mirapoint Message Servers can be clustered into a failover configuration where multiple systems use a single standby head for take-over. N equals the number of active systems, while +1 represents the standby system. N can be any number up to nine (9) although four (4) is the generally recommended maximum.
- Switch A selector device with Fibre Channel connectors that supports multipoint to multipoint communication, device name lookup, security, and redundancy.

Configuring Mirapoint SAN

Mirapoint “S” Message Servers (M5000s and S6000) employ SAN (storage area networks) instead of direct-attached disk shelves. The SAN enables N+1 failover clustering, where a spare server takes over whenever any server in the cluster fails. N can be a number from 1 through 9, although 4 is the recommended maximum in most cases. After takeover, the replacement server uses the same SAN partition as the failed server it replaced. Once repaired, the failed server functions as the new spare.

As an example of SAN configuration, this chapter describes the procedure for setting up an XSAN340 for use with Mirapoint Message Servers to provide email services. Passwords and IP addresses are intentionally omitted.

Hardware Setup

This chapter describes how to configure XSAN storage for use with Mirapoint Message Servers. Operational procedures such as volume and LUN management, snapshots, and backup integration, are covered in the next chapter.

Site Requirements

Before you configure SAN storage, check for the following:

- ◆ Optical fiber cables to connect all FC ports (Mirapoint and Filer) to the XSAN.
- ◆ SFP (small form-factor pluggable) connectors, also called GBIC (gigabit interface connector). The SAN Switch and Filer require SFP connectors on all ports. These are normally shipped with the Filer and Switch.

Filer Power and Cabling

Refer to this *Site Requirements Guide* for power requirements:

<http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/site.pdf>

Refer to these quick-start *Installation and Setup Instructions* for help cabling XSAN300-Series Filers:

<http://now.netapp.com/NOW/knowledge/docs/ontap/re172/pdfs/filer/f3000isi.pdf>

Mirapoint recommends use of “single image mode” rather than “standby mode” on the XSAN appliance. This simplifies connections between switches and Filers, requiring only two instead of eight, saving six fiber-optic cables. See [Figure 1](#).

Connections from the first HBA (host bus adapter) to the first switch are blue. Connections from the second HBA to the second switch are shown in red. The XSAN to RAID connections are black, and XSAN Filer interconnects are green.

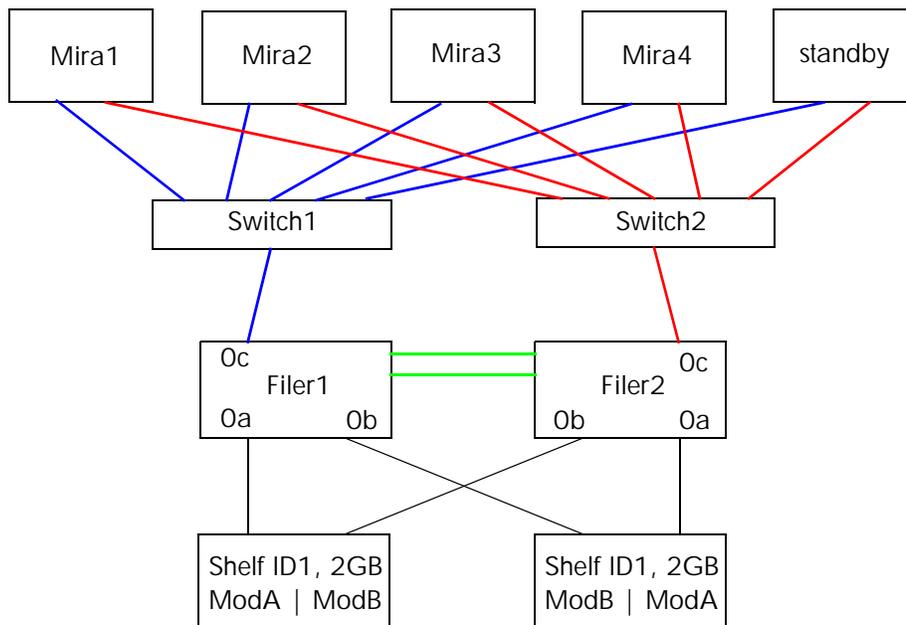


Figure 1 Switch and XSAN Cabling (MOS 3.x Systems)

To connect FC (Fibre Channel) ports on the XSAN Filers:

1. Connect FC port **0c** to the SAN switches.
2. **MOS 4.x systems only:** Connect FC port **0d** to the opposite SAN switches, as shown in [Figure 2](#) on page 11. Do not connect port **0d** on MOS 3.x systems, which cannot support SAN Multipath.
3. Connect FC port **0a** to Module A on the disk shelves.
4. Connect FC port **0b** to Module B on the disk shelves.
5. If this is a clustered Filer, attach the Filer interconnect cables.
6. For the serial console, attach a serial cable (null modem) to the console port of each Filer. XSAN should include a RJ45 to DB9 connector.

Power On

Power on all servers as close to the same time as possible.

If the two XSAN nodes are booted more than five minutes apart, failover might occur. If an XSAN comes up in “failover” mode, issue the following command on that node:

```
XSAN1 (takeover)> cf giveback
```

You can identify the **takeover** node by its prompt.

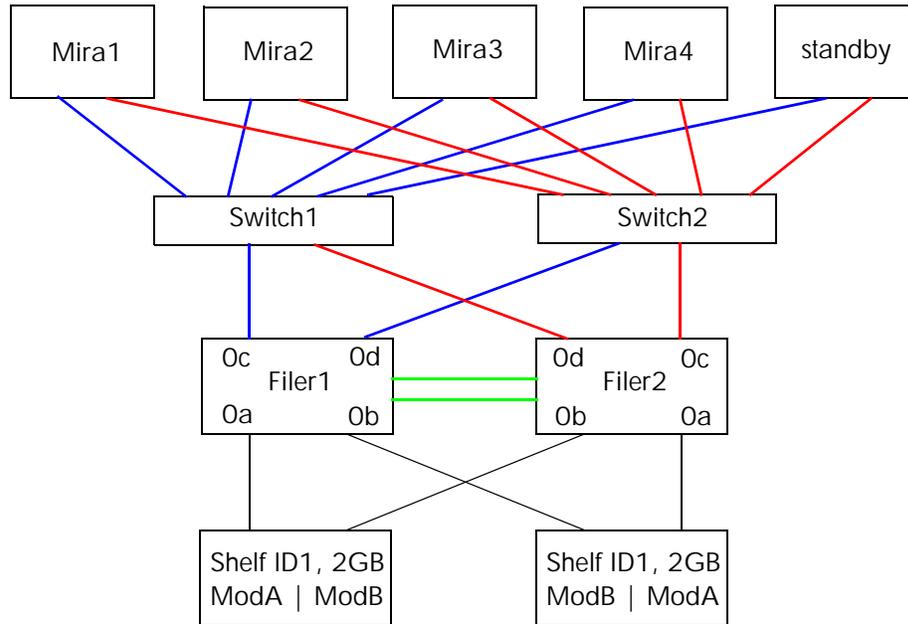


Figure 2 Switch and XSAN Cabling (MOS 4.x Systems)

XSAN Basic Setup

Network Setup

A serial console must be connected to the XSAN. If this is a clustered XSAN, you must connect and configure each Filer node separately.

Login Information

Login: **root**

Password: no password (if the password is set, it is usually "XSAN1")

Network Information

Be prepared with the following network information:

- ◆ Hostname (two, if clustered)
- ◆ IP address (two, if clustered)
- ◆ Netmask
- ◆ Router
- ◆ DNS server
- ◆ Timezone (can be configured later)

Running XSAN Setup

The setup command rewrites the `/etc/rc`, `/etc/exports`, `/etc/hosts`, `/etc/hosts.equiv`, `/etc/dgateways`, `/etc/nsswitch.conf`, and `/etc/resolv.conf` files, saving the original contents of these files in `.bak` files (for example, `/etc/exports.bak`).

1. Run the **setup** command, supplying network information as outlined above:

```

XSAN> setup
Are you sure you want to continue? [yes]
NetApp Release 7.2.1: Sun Dec 10 01:33:06 PST 2006
System ID: 0101194245 (XSANa); partner ID: 0101194307 (12b)
System Serial Number: 3078515 (XSAN a)
System Rev: C0
slot 0: System Board
  Processors:          4
  Memory Size:        3072 MB
slot 0: Dual 10/100/1000 Ethernet Controller VI
  e0a MAC Address:    00:a0:98:06:c4:26 (auto-100tx-fd-up)
  e0b MAC Address:    00:a0:98:06:c4:27 (auto-unknown-cfg_down)
  e0c MAC Address:    00:a0:98:06:c4:24 (auto-100tx-fd-up)
  e0d MAC Address:    00:a0:98:06:c4:25 (auto-unknown-cfg_down)
slot 0: Fibre Channel Target Host Adapter 0a
slot 0: Fibre Channel Target Host Adapter 0b
slot 0: Fibre Channel Target Host Adapter 0c
slot 0: Fibre Channel Target Host Adapter 0d
slot 0: SCSI Host Adapter 0e
slot 0: NetApp ATA/IDE Adapter 0f (0x000001f0)
  0f.0                977MB
slot 1: FC Host Adapter 1a
slot 1: FC Host Adapter 1b
slot 2: FC Host Adapter 2a
  28 Disks:           7616.0GB
  2 shelves with ESH2
slot 2: FC Host Adapter 2b
  42 Disks:           11424.0GB
  3 shelves with ESH2
slot 3: NVRAM
  Memory Size:        512 MB
Please enter the new hostname []: XSANa01
Do you want to configure virtual network interfaces? [n]:
Please enter the IP address for Network Interface e0a []: X.X.X.X
Please enter the netmask for Network Interface e0a []: X.X.X.X
Should interface e0a take over a partner IP address during failover? [n]:
Please enter media type for e0a {100tx-fd, tp-fd, 100tx, tp, auto (10/100/
1000)} [auto]:
Please enter flow control for e0a {none, receive, send, full} [full]:
Do you want e0a to support jumbo frames? [n]:
Please enter the IP address for Network Interface e0b []:
Should interface e0b take over a partner IP address during failover? [n]:
Please enter the IP address for Network Interface e0c []: X.X.X.X
Please enter the netmask for Network Interface e0c []: X.X.X.X
Should interface e0c take over a partner IP address during failover? [n]:
Please enter media type for e0c {100tx-fd, tp-fd, 100tx, tp, auto (10/100/
1000)} [auto]:
Please enter flow control for e0c {none, receive, send, full} [full]:
Do you want e0c to support jumbo frames? [n]:
Please enter the IP address for Network Interface e0d []:
Should interface e0d take over a partner IP address during failover? [n]:
Would you like to continue setup through the web interface? [n]:
Please enter the name or IP address of the default gateway []: X.X.X.X
  The administration host is given root access to the filer's
  /etc files for system administration. To allow /etc root access

```

to all NFS clients enter RETURN below.
 Please enter the name or IP address of the administration host:
 Please enter timezone [GMT]: Continent/City
 Where is the filer located? []: Park Royal Site
 Do you want to run DNS resolver? [n]: y
 Please enter DNS domain name []: XXXXXXXX
 You may enter up to 3 nameservers
 Please enter the IP address for first nameserver []: X.X.X.X
 Do you want another nameserver? [n]: y
 Please enter the IP address for alternate nameserver []: X.X.X.X
 Do you want another nameserver? [n]:
 Do you want to run NIS client? [n]:
 The initial aggregate currently contains 3 disks; you may add more
 disks to it later using the "aggr add" command.

2. If you know an NTP server, make use of it for the Filer:

```
XSANa01> options timed.servers NTPServerIPaddr
```

3. Type **reboot** for all these changes to take effect.

```
XSAN> reboot
```

After reboot, log in over the network and continue setup using **telnet**.

License

All licenses should be installed on the XSAN system. FCP and Snapstore must be installed for Mirapoint interoperation. If this is a clustered XSAN, in addition to these licenses, you should see cluster licenses on both nodes.

To list and install the license keys:

```
XSANa01> license
XSANa01> license add KEY
```

If you do not see pre-installed licenses, login to <http://now.netapp.com> > My Now, and search for license keys by your system serial number or ID, which you can find in output from the **sysconfig -a** command.

XSAN OS Upgrade

Currently, XSAN Data Ontap 7.2.2 is the certified release for Message Servers. Mirapoint strongly recommends that you upgrade to this XSAN version and stay on it until advised to upgrade to a newer Mirapoint-certified release.

To see your current Data Ontap version:

```
XSANa01> version
```

To upgrade, first login to <http://now.netapp.com>, locate the 7.2.2 binary under Download Software > Data ONTAP, and read the 7.2.2 release notes:

<http://now.netapp.com/NOW/knowledge/docs/ontap/re1722/pdfs/ontap/rnote.pdf>

1. Unlike Mirapoint's **Update** command, you must manually download the binary from <http://now.netapp.com>, place the binary on a local HTTP server, then install it. For example, after obtaining a Data Ontap 7.2.2 binary named **722_setup_e.exe**, issue the following command to upgrade:

```
XSANa01> software get http://httpserver/722_setup_e.exe
```

2. That places the binary into an internal directory, /etc/software. The next three commands unzip and install the binaries, update on-board flash memory with new software, and reboot to complete the upgrade. For example:

```
XSANa01> software install 722_setup_e.exe
XSANa01> download
XSANa01> reboot
```



Check for software updates here:

http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

3. After reboot, run this command to verify the OS. XSAN OS can be revoked.

```
XSANa01> version -b
```

4. If this is a clustered XSAN, repeat the upgrade procedure for Node B.

SAN Network Setup

Connect all Fibre Channel cables from the Mirapoint Message Servers to the SAN switches, and from the SAN switches to the XSAN Filer(s).

Most XSAN300-Series Filers come with four onboard FC ports (0a, 0b, 0c 0d) and four additional FC ports on two extra cards for each node. Use ports 0a and 0b to connect the disk shelves to the XSAN Filer(s). Use port 0c (and 0d on release 4.x) to connect a SAN switch to an XSAN Filer. See the cabling diagrams, [Figure 1](#) on page 10 or [Figure 2](#) on page 11.

For SAN switch configuration, consult the switch vendor documentation.

Fibre Channel Protocol

This section describes setting up Fibre Channel (FC) protocol on an XSAN.

FCP Configuration

Typically, an XSAN Filer comes with all FC ports configured as “initiators.” “Initiator” is the FC setting for connecting the disk shelves. Because the Mirapoint server is the initiator while the XSAN Filer is the target, you need to convert some of the FC ports to “target” so that Mirapoint Message Server can attach to the XSAN Filers.



Releases prior to Mirapoint MOS 4 do not support multipathing, so one of the target FC ports must be manually disabled. If the port is not disabled, multiple SAN paths would appear, and Mirapoint servers would fail to boot.

Before starting, make sure all ports (0a, 0b, 0c and 0d) on the XSAN are connected properly to the SAN switch. If this is a clustered XSAN, make sure that you connect the FC ports from each Filer node to multiple switches as shown in [Figure 1](#) on page 10.

To show all the fiber interfaces and their current settings:

1. Type the following commands:

```
XSANa01> fcdadmin config
XSANa01> fcp config
```

2. If there is no “target” FC port, convert 0c (and optionally 0d) then reboot.

```
XSANa01> fcdadmin config -t target 0c
XSANa01> fcdadmin config -t target 0d
XSANa01> reboot
```

3. After reboot, verify that configuration changed, and check the target FC ports:

```
XSANa01> fcdadmin config
XSANa01> fcp config
```

If the FC ports are not up, check the fiber connection.

Table 2 shows differences in configuration of target FC adapters, before and after.

Table 2 Before and After FC Configuration

System and FCP Status																					
Before	<pre>XSANa01> fcp nodename Fibre Channel nodename: 00:00:00:00:00:00:00:00 (0000000000000000) XSANa01> fcp config XSANa01> fcp show cfmode fcp show cfmode: unknown XSANa01> fcdadmin config</pre> <table border="1"> <thead> <tr> <th>Adapter</th> <th>Type</th> <th>Local State</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td>0a</td> <td>initiator</td> <td>UNDEFINED</td> <td>online</td> </tr> <tr> <td>0b</td> <td>initiator</td> <td>UNDEFINED</td> <td>offline</td> </tr> <tr> <td>0c</td> <td>initiator</td> <td>UNDEFINED</td> <td>online</td> </tr> <tr> <td>0d</td> <td>initiator</td> <td>UNDEFINED</td> <td>offline</td> </tr> </tbody> </table> <pre>XSANa01> fcp show initiator XSANa01></pre>	Adapter	Type	Local State	Status	0a	initiator	UNDEFINED	online	0b	initiator	UNDEFINED	offline	0c	initiator	UNDEFINED	online	0d	initiator	UNDEFINED	offline
Adapter	Type	Local State	Status																		
0a	initiator	UNDEFINED	online																		
0b	initiator	UNDEFINED	offline																		
0c	initiator	UNDEFINED	online																		
0d	initiator	UNDEFINED	offline																		

Table 2 Before and After FC Configuration (Continued)

	System and FCP Status
After	<pre> XSANa01*> fcp nodename Fibre Channel nodename: 50:0a:09:80:86:28:1a:05 (500a098086281a05) XSANa01*> fcp config 0c: ONLINE <ADAPTER UP> PTP Fabric host address 021b00 portname 50:0a:09:81:86:28:1a:05 nodename 50:0a:09:80:86:28:1a:05 mediatype auto partner adapter None speed auto 0d: ONLINE <ADAPTER UP> PTP Fabric Standby host address 022200 portname 50:0a:09:81:00:01:6f:d9 nodename 50:0a:09:81:00:01:6f:d9 mediatype auto partner adapter 0c speed auto XSANa01pr*> fcp show cfmode fcp show cfmode: standby XSANa01*> fcdadmin config Adapter Type Local State Status ----- 0a initiator CONFIGURED online 0b initiator CONFIGURED online 0c target CONFIGURED online 0d target CONFIGURED online XSANa01*> fcp show initiator Initiators connected on adapter 0c: Portname Group 21:00:00:e0:8b:8a:85:b8 21:00:00:e0:8b:82:5c:c0 21:00:00:e0:8b:1b:de:ba 21:00:00:e0:8b:18:d6:a7 21:00:00:e0:8b:08:e7:73 21:00:00:e0:8b:17:04:bc 21:00:00:e0:8b:8a:87:b8 Initiators connected on adapter 0d: None connected. </pre>

Cluster Failover (CF) Mode Setup

This section applies only to clustered Filers. XSAN supports several cluster failover (CF) modes that have multipathing implications.

If the XSAN system is running Data Ontap 7.2.2, FAS3020C (and similar models) should be configured with Single System Image Mode.

1. Check the setting of cluster failover (CF) mode:

```
XSANa01> fcp show cfmode
```

2. If the XSAN is not properly configured, change the CF mode as follows:

```
XSANa01> fcp status
```

```
XSANa01> fcp stop
```

```
XSANa01> priv set advanced
```

```
XSANa01*> fcp set cfmode single_image
```

```
Are you sure you want to change the cfmode?
```

```
The cfmode setting must be the same on both filers in a cluster.
```

```
Improper setting of this option can cause host outages, continue? yes
```

```
Mon Nov 13 20:58 GMT [XSANa01: scsitarget.ispfct.cfmodeChanged:warning]:
```

```
FCP cfmode changed from partner to single_image.
```

```
XSANa01*> priv set admin
```

```
XSANa01> reboot
```

3. After reboot, restart FCP and check status:

```
XSANa01> fcp show cfmode
XSANa01> fcp start
XSANa01> fcp status
```

4. If necessary, repeat all the above steps on the second node.

Igroup Initiator Group

A world wide port name (WWPN) is a unique identifier in the SAN. An Igroup is a group of WWPNs, used to map into LUNs created on the XSAN Filer. The XSAN Filer offers its LUNs to this group of initiators.

In standard (non-N+1) configurations, LUNs need to be mapped only for the Mirapoint sever that binds to it. In N+1 failover configurations, LUNs must be accessible to all Mirapoint's WWPNs so that any Mirapoint head can change itself into standby node. To make a cluster ready for N+1 configuration, Mirapoint recommends creating an Igroup with the same name as the N+1 cluster, "c101" for instance, containing all the WWPNs of all Mirapoint servers in the cluster.

1. Make sure the FC ports (0a, 0b, 0c) on the XSAN Filer are connected properly to the RAID and SAN switch. If Mirapoint servers are configured with dual HBA, each HBA should be connected to a different SAN switch.
2. Get the WWPNs of the HBA on all Mirapoint SAN systems

```
ms14.domain> storage getport port0 portwwn
0x210000e08b82f9bb
OK Completed
ms14.domain> storage getport port1 portwwn
0x210000e08b821fbf
OK Completed

ms15.domain>
...
```

3. If cables are connected properly, these WWPNs should also display on the XSAN Filer. Run the following command to make sure they do. Note that the fourth and ninth lines, after removing colons, correspond to the WWPNs.

```
XSANa01> fcp show initiator
Initiators connected on adapter 0c:
Portname          Group
21:00:00:e0:8b:82:28:be
21:00:00:e0:8b:82:bb:3b
21:00:00:e0:8b:17:6e:dd
21:00:00:e0:8b:82:f9:bb
21:00:00:e0:8b:18:a0:a1
21:00:00:e0:8b:0d:b4:c1
21:00:00:e0:8b:17:26:bc
21:00:00:e0:8b:82:a5:be
21:00:00:e0:8b:82:1f:bf
21:00:00:e0:8b:18:ea:2d
Initiators connected on adapter 0d:
None connected.
```

4. Create an Igroup and add the WWPNs to it. For **Igroup Add**, take the WWPN of Mirapoint ports from the **Storage Getport** commands in step 2.

```
XSANa01> igroup create -f -t linux c101
XSANa01> igroup add c101 0x210000e08b82f9bb
```

```

XSANa01> igroup add c101 0x210000e08b821fbf
...
XSANa01*> igroup show
  c101 (FCP) (ostype: linux):
    21:00:00:e0:8b:82:f9:bb (logged in on: 0c)
    21:00:00:e0:8b:82:28:be (logged in on: 0c)
    21:00:00:e0:8b:82:bb:3b (logged in on: 0c)
Now you have an Igroup named "c101" (same as the Mirapoint cluster name),
which you can map to LUNs you'll create in LUN Configuration on page 23.

```

XSAN Storage Setup

This section tells how to configure the following RAID storage items on the XSAN:

- ◆ [Disk](#) on page 18
- ◆ [Storage Allocation](#) on page 19
- ◆ [Aggregate, Aggr](#) on page 20
- ◆ [Volume and LUN Mapping for Mirapoint](#) on page 21

Disk

To check status of storage on the disk shelves, use the **Aggr Status** command. It shows which disks are in use, which disks are free to use for expansion, and which disks are available to another Filer. The first portion of the output shows which disks are currently in use; there are two of them for parity. The second portion shows the spare disks that can be used for expansion. The final portion, Partner disks, shows which disks are available to the other XSAN Filer node in the cluster.

```

XSANa01> aggr status -r
Aggregate aggr0 (online, raid_dp) (block checksums)
Plex /aggr0/plex0 (online, normal, active)
RAID group /aggr0/plex0/rg0 (normal)
RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/b1ks) Phys (MB/b1ks)
-----
dparity 2a.16 2a 1 0 FC:A - FCAL 10000 272000/557056000 274845/562884296
parity 2a.32 2a 2 0 FC:A - FCAL 10000 272000/557056000 274845/562884296
data 2a.17 2a 1 1 FC:A - FCAL 10000 272000/557056000 274845/562884296

Spare disks
RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/b1ks) Phys (MB/b1ks)
-----
spare 2a.18 2a 1 2 FC:A - FCAL 10000 272000/557056000 274845/562884296
spare 2a.19 2a 1 3 FC:A - FCAL 10000 272000/557056000 274845/562884296
spare 2a.20 2a 1 4 FC:A - FCAL 10000 272000/557056000 274845/562884296
spare 2a.21 2a 1 5 FC:A - FCAL 10000 272000/557056000 274845/562884296
spare 2a.22 2a 1 6 FC:A - FCAL 10000 272000/557056000 274845/562884296

Partner disks
RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/b1ks) Phys (MB/b1ks)
-----
partner 2b.32 2b 2 0 FC:B - FCAL 10000 0/0 274845/562884296
partner 2b.48 2b 3 0 FC:B - FCAL 10000 0/0 274845/562884296
partner 2b.52 2b 3 4 FC:B - FCAL 10000 0/0 274845/562884296
partner 2b.55 2b 3 7 FC:B - FCAL 10000 0/0 274845/562884296

```

Storage Allocation

Before building aggregates and allocating disks, plan the storage layout. You can take the amount of storage available and calculate how much usable space is available. Mirapoint recommends the opposite approach: start from the amount of usable space that is required and calculate how much storage to allocate.

Message Servers and Aggregates

Mirapoint used to suggest allocating 40% or less of each volume for snapshots. This is because email folders change only when new messages arrive and when messages are deleted. User mailboxes often store a large volume of messages that remain in place month after month, not generating any snapshot.

You might decide to retain snapshots for an extended period. For longer-lived snapshots the standard recommendation for sizing is $2x + \text{Delta}$. This implies an allocation of 60% on each volume for snapshots. If you choose this configuration, you must ensure that there is room to grow the aggregate, if necessary.

To allow a maximum Message Server size of 2.8TB (initially 2TB with 40% extra for growth), and allocating 60% of a volume for snapshots, calculate as follows:

$$7 \text{ TB} * 60\% = 4.2 \text{ TB (snapshot space)}$$

$$7 \text{ TB} * 40\% = 2.8 \text{ TB (LUN space)}$$

A single aggregate is limited to 16 TB, which is enough for two Message Servers.

It is possible to move a volume (and therefore a Message Server) from one aggregate to another. This involves a service outage for data transfer.

Volume and Aggregate Space Allocation

For the initial N+1 clustered Message Server, the following volume allocation is required. This is based on a 40% snapshot allocation (that is, $Z = 0.4$).

Table 3 Volumes Required

	Calculate	Answer
100GB allocated for special volumes: /vol/vol0 /vol/mira0 /vol/chYY	These are constant: /vol/vol0 : 40GB /vol/mira0 : 20MB /vol/chYY : 50GB	100GB
Volume size allocated for the 2 x 1TB LUNs in the volume	$S = 2 / (1 - Z)$	$2 / 0.6 = 3400\text{GB}$
Total Aggregate size required		$3400\text{GB} + 100\text{GB} = 3500\text{GB}$

You can later add an additional Message Server to the aggregate. The following volume allocation would then be required, based on a 40% snapshot allocation (that is, $Z = 0.4$).

Table 4 Volumes Required After Clustering

Space Requirements	Calculate	Answer
100GB allocated for special volumes: /vol/vol0 /vol/mira0 /vol/chYY	These are constant: /vol/vol0 : 40GB /vol/mira0 : 20MB /vol/chYY : 50GB	100GB
Volume size allocated for the 2 x 1TB LUNs in the volume	$S = 2 / (1 - Z)$	$2 / 0.6 = 3400\text{GB}$
Second Message Server: Volume size allocated for the 2 x 1TB LUNs in the volume	$S = 2 / (1 - Z)$	$2 / 0.6 = 3400\text{GB}$
Total Aggregate size required		$3400\text{GB} + 3400\text{GB} + 100\text{GB} = 6900\text{GB}$

Aggregate, Aggr

Aggregate is the virtual layer for disks that form the RAID-DP underneath. The XSAN Filer has a default bootable aggregate **aggr0**, minimally a three-drive aggregate with one data disk and two parity disks to form the RAID-DP array.

Instead of creating a new aggregate, it is best to expand the existing aggr0 for the first Message Server. This maximizes storage space and uses as many disk spindles as possible for better throughput. However, save some disks as hot spares. Use a RAID size of 26, so that for an aggregate containing two shelves with 28 drives, you have two parity disks and two hot spares.

In the previous section, aggregate0 was sized for two Message Servers, requiring 6900 GB. Three disk shelves provide 7900 GB (after reserving three hot spares), so allocate up to three disk shelves to aggregate0 to provide the required capacity for two Message Servers plus a small amount for additional use if needed.

For a table estimating the number of disks required in an aggregate to achieve a desired size, See [Appendix A, Usable Space](#).

For example, on XSANa01, there are initially 28 disks in total. Three are already allocated to aggregate 0, so add 23 more disks to aggregate0, leaving two spare (unallocated). This leads to a usable space of 5400 GB.

On XSANb01, there are initially 42 disks in total. Three are already allocated to aggregate0, so add 36 more disks to aggregate0, leaving three spare (unallocated). This leads to a usable space of 7900 GB.

In a clustered XSAN, each Filer node owns its disk shelf, essentially, so the number of available disks on each node is cut in half, and you need to do the configuration procedure separately on two nodes.

Clustered Aggregate

Follow this procedure very carefully. If you make a mistake in aggregate configuration, you could configure unnecessary redundant parity or leave no spare disk. Also, once the disks are added into aggregate, they cannot be removed without performing a factory reset, a time consuming process as it involves disk zeroing.

1. Change the RaidSet to 26. This allows 24 data disks, 2 parity disks and 2 hot spares per 28 disks (2 shelves). The Filer starts new RaidSets automatically as required. Check the setting, then change it:

```
XSANa01> aggr status -v
XSANa01> aggr options aggr0 raidsize 26
```

2. Verify the setting, and check the status of disks and storage size:

```
XSANa01> aggr status -v
XSANa01> aggr status -r
XSANa01> df -Ag
```

3. Add the disks to **aggr0** until it reaches the desired aggregate size.

There are two ways of adding disk to aggr0. Use the safer approach, which is adding a few disks at a time, rather than adding them all together.

```
XSANa01> aggr add aggr0 -d diskId1 diskId2 ...
```

For example, to add three disks at a time:

```
XSANa01> aggr add aggr0 -d 2a.19 2a.19 2a.20
```

Or, to add only one disk at a time:

```
XSANa01> aggr add aggr0 -d 2a.19
```

4. Verify the status of the disks and storage size:

```
XSANa01> aggr status -r
XSANa01> df -Ag
```

5. If this is a clustered XSAN, repeat the same configuration on Filer Node B.

Volume and LUN Mapping for Mirapoint

There are several rules for LUN configuration:

- ◆ For Mirapoint servers running MOS 3, LUN Zero must exist as a pilot to communicate with the XSAN. The LUN ID for this special LUN must be Zero. Clustered XSAN configurations also need only one LUN Zero for all nodes.



Mirapoint servers running MOS 4 automatically map LUN Zero.

- ◆ Each Mirapoint head must have a dedicated volume. This ensures that the snap policy (such as snapmirror or snapvault) can be applied independently for each Mirapoint Message Server.
- ◆ All LUNs belonging to a Mirapoint server must be provisioned on the same volume. Without this, it is impossible to ensure mail store consistency, which

requires all the LUNs in the message store to be saved by snapshot at the same time. XSAN calls this an “atomic” snapshot.

Initial Layout

[Table 5](#) shows all the volumes allocated from aggregate0. The snap reserve is 0% for all volumes, but 40% space is allocated in each volume for snapshots.

Table 5 Layout for Snapshot

XSAN	Volume	LUN ID	Igroup	Notes
3020a01	vol0 size=40GB	N/A	N/A	For the XSAN Filer's own use
	mira0 size=20MB	lun0 size=4MB ostype: linux	cl01: with id=0 ostype: linux	LUN 0, required by Mirapoint appliances
	ms14 size=3400GB flexible volume	lun1 size=1023GB	cl01: with id=11 ostype: linux	Message Store 14 LUN 1 (first cluster node)
		lun2 size=1023GB	cl01: with id=12 ostype: linux	Message Store 14 LUN 2 (optional first cluster node)
3020b01	vol0 size=40GB	N/A	N/A	For the XSAN's own use
	ch01 size=50GB	lun1 size=35GB	cl01: with id=1 ostype: linux	Standby node
	ms15 size=3400GB flexible volume	lun1 size=1023GB	cl01: with id=21 ostype: linux	Message Store 15 LUN 1 (second cluster node)
		lun2 size=1023GB	cl01: with id=22 ostype: linux	Message Store 15 LUN 2 (optional second cluster node)

Volume Configuration

After you come up with your own volume and LUN mapping table, as in [Table 5](#), proceed with configuration commands. Some commands are for node A and some for node B because this is a clustered XSAN. For single system image (SSI) mode, there is no need to create **/vol/mira0** on the second node.

1. To perform basic volume configuration:

```
XSANa01> vol size /vol/vo10 40g
XSANa01> vol create /vol/mira0 aggr0 20m
XSANb01> vol size /vol/vo10 40g
XSANb01> vol create /vol/chXX aggr0 50g
```

Replace *chXX* with the name of the N+1 standby head, for example *ch01*.

```
XSANb01> vol create /vol/ch01 aggr0 50g
```

2. To perform active mail store volume configuration:

```
XSANa01> vol create /vol/msXX aggr0 S
...
```

Replace *S* with the volume storage for active mail store, for example 3400GB. Replace the volume name, *msXX*, with the name of the message store that uses it, for example *ms14* and *ms15*.

```
XSANa01> vol create /vol/ms14 aggr0 3400GB
XSANb01> vol create /vol/ms15 aggr0 3400GB
```

For information about XSAN maintenance commands to delete a volume, see **Help Vol Offline** and **Help Vol Destroy**.

LUN Configuration

To configure LUNs, you can use either the command line or XSAN's interactive "LUN Setup" wizard. This procedure uses the command line.

Replace the volume name *msXX* with the name of the message store that uses it, for example *ms14*. Replace *chXX* with the name of the N+1 standby head, for example *ch01*. Replace *c1XX* with the name of the N+1 cluster, for example *c101*.

1. Configure LUN zero for the required *mira0* volume:

```
XSANa01> lun create -s 4m -t linux /vol/mira0/lun0
XSANa01> lun map /vol/mira0/lun0 c101 0
```

The *mira0* LUN must be assigned with ID 0. If this command returns an error, find out which LUN was mapped to ID 0, unmap the LUN, and try again.

2. Configure a LUN for the standby node, as in [Table 5](#) on page 22. Map it to the *Igroup*.

```
XSANa01> lun create -s 35g -t linux /vol/ch01/lun1
XSANa01> lun map /vol/ch01/lun1 c101 1
```

3. Configure two 1 TB LUNs for *ms14*, as in [Table 5](#) on page 22. Map them to the *Igroup*.

```
XSANa01> lun create -s 1023g -t linux /vol/ms14/lun1
XSANa01> lun create -s 1023g -t linux /vol/ms14/lun2
XSANa01> lun map /vol/ms14/lun1 c101 1
XSANa01> lun map /vol/ms14/lun2 c101 2
```

4. If this is a clustered XSAN, repeat the LUN configuration commands on Filer Node B.

For information about the XSAN maintenance commands to delete a LUN, see **Help Lun Unmap** and **Help Lun Destroy**.

Snapshot Configuration

Mirapoint recommends that you turn off snapshot reservation. This example shows how to disable snapshot reservation for the **ms14** volume:

```
XSANa01> df -g ms14
Filesystem                total      used   avail capacity  Mounted on
/vol/ms14/                2720GB    2050GB   669GB    75%   /vol/ms14/
/vol/ms14/.snapshot      680GB         0GB   680GB     0%   /vol/ms14/.snapshot
XSANa01*> snap reserve ms14
Volume ms14: current snapshot reserve is 20% or 713031680 k-bytes.
XSANa01*> snap reserve ms14 0
XSANa01*> df -g ms14
Filesystem                total      used   avail capacity  Mounted on
/vol/ms14/                3400GB    2050GB  1349GB    60%   /vol/ms14/
/vol/ms14/.snapshot       0GB         0GB     0GB     ---%   /vol/ms14/.snapshot
```

Repeat the same configuration for all Mirapoint volumes:

```
XSANa01> snap reserve mira0 0
XSANa01> snap reserve ch01 0
XSANa01> snap reserve ms15 0
```

If this is a clustered XSAN, repeat commands on the second node:

```
XSANb01> snap reserve mira0 0
XSANb01> snap reserve ch01 0
XSANb01> snap reserve ms14 0
XSANb01> snap reserve ms15 0
```

Change of Snap Schedule

Instead of using the default snap schedule, Mirapoint recommends setting a new snap schedule. Snapshots can be created daily, saved for a week, then weekly snapshots are saved for two weeks. For example:

To set a snapshot schedule for the **ms14** volume:

1. Check the existing snap schedule:

```
XSANa01> snap sched ms14
```

2. Change the schedule, and check the change:

```
XSANa01> snap sched ms14 2 6
XSANa01> snap sched ms14
```

3. Repeat the schedule change for all Mirapoint volumes, and on both nodes:

```
XSANb01> snap sched ch01 2 6
XSANb01> snap sched ms15 2 6
...
```

You cannot auto-delete snapshots that are in use by a clone-created LUN. Ensure that your backup scripts always release the clone-created LUNs and delete the snapshot, otherwise that volume might fill up and the LUNs will return an ENOSPC (out of space) error when a write is attempted, thus causing a service outage on the Message Servers.

Administering Mirapoint SAN

This chapter covers day-to-day administration topics, including:

- ◆ [Data Protection](#) on page 25
- ◆ [Auto Support](#) on page 27
- ◆ [Operational Procedures](#) on page 27
- ◆ [Troubleshooting](#) on page 35
- ◆ [Factory Reset](#) on page 36

Data Protection

This section describes various techniques for protecting Message Server data.

Snapshot and Backup

With snapshots scheduled, a consistent image of all the volumes gets saved. The snapshot alone is not a complete backup. Mirapoint recommends that you save snapshot images to alternate media. There are many options for data backups and data restores, including SnapVault or NDMP backup, and Snapmirror for site replication. Refer to the third-party documentation as needed for specific backup and restore procedures.

Snaprestore and Selective Restore

The LUN-based snapshot allows only system-wide restore, leaving two restore scenarios to consider.

Restore to Previous Snapped Image on Same Mirapoint Head

This is a quick restore of the Mirapoint mail store to a previous stage. All changes since the previous snapshot are lost. Follow these steps:

1. Shutdown the Mirapoint Message Server.
2. On the XSAN, run a command such as this:

```
XSANA01> snap restore -t vol -s hourly.2 /vol/ch01
```

Wait until snap restore is complete (normally less than a minute).

3. Reboot the Mirapoint Message Server.

Rebuild a System Image on Different Mirapoint Head

The following procedure is available only if you retain snapshots for a while. If you discard snapshots soon after (NDMP) tape backup, it is often impossible to go back far enough in time to rebuild the system image. You might need to review snapshot policies and the amount of disk space you are allocating for snapshots.

Mirapoint does not support selective mailbox restore for SAN-based mail stores. However, with retained snapshots, you could attach a new Mirapoint head to old snapshot LUNs. The new Mirapoint head would come up as another server but with an old mail store. You could then restore user mailboxes to a previous state.

Snapshot LUNs cannot be attached directly. You must clone them before attaching. To do this, prepare a fresh Mirapoint system. Attach its FC port to the SAN switch. Note the WWPN; you need it to build a new Igroup.

To clone a LUN image based on a previous snapshot, follow these steps:

1. To clone the LUNs under volume ms14 for Mirapoint 14, for example:

```
XSANa01*> snap list ms14
Volume ms14
working...
%/used      %/total    date          name
-----
0% ( 0%)    0% ( 0%)    Mar 19 12:15  hourly.2
XSANa01*> lun show
/vol/mira0/lun0          4m (4194304)      (r/w, online, mapped)
/vol/ms14/lun1          1023.1g (1098572103680) (r/w, online, mapped)
/vol/ms14/lun2          1023.1g (1098572103680) (r/w, online, mapped)
```

2. Take the hourly.2 snapshot and clone the LUNs under ms14:

```
XSANa01> lun clone create /vol/ms14/SelectiveRestoreLUN1 -b /vol/ms14/lun1 hourly.2 -o noreserve
XSANa01> lun clone create /vol/ms14/SelectiveRestoreLUN2 -b /vol/ms14/lun2 hourly.2 -o noreserve
```

The **snap** command now shows as busy:

```
XSANa01*> snap list ms14
Volume ms14
working...
%/used      %/total    date          name
-----
0% ( 0%)    0% ( 0%)    Mar 19 12:15  hourly.2          (busy,LUNs)
```

No actual disk space is consumed because they are on the same volume, however the parent snapshot is marked busy so it is undeletable. The option **noreserve** is required, otherwise Data Ontap attempts to reserve an additional 1TB of space per LUN, and that space is not available.

3. Create the Igroup and map its LUNs:

```
XSANa01> igroup create -f -t linux mira-new WWPN
XSANa01> lun map /vol/ms14/SelectiveRestoreLUN1 mira-new 101
XSANa01> lun map /vol/ms14/SelectiveRestoreLUN2 mira-new 102
XSANa01*> lun show -m
LUN path                Mapped to          LUN ID  Protocol
-----
/vol/mira0/lun0         mira-all           0       FCP
/vol/ms14/lun1         c101                11      FCP
/vol/ms14/lun2         c101                12      FCP
/vol/ms14/SelectiveRestoreLUN1  mira-new           101     FCP
/vol/ms14/SelectiveRestoreLUN2  mira-new           102     FCP
```

4. (MOS 4) On the newly prepared Mirapoint system, perform a **reboot**. This LUN should appear on the list.
(MOS 3) On the newly prepared Mirapoint system, perform a **Storage Scan** command and then a **Storage Listsan** command. This LUN should appear on the list.
5. You can now “reattach” the Mirapoint SAN head onto these LUNs.

Auto Support

XSAN Filers offer autosupport, similar to Mirapoint CustomerCare, enabled by default. Add `dl-netapp-install@mirapoint.com` to the autosupport list, plus any other email addresses you want to notify:

```
XSANa01> options autosupport.to dl-netapp-install@mirapoint.com your@address
```

You can also add autosupport addresses using XSAN’s web-based administration interface at `http://filename/na_admin` — after login, click FilerViewer and navigate to Filer > Configure Autosupport.

Shelf Expansion

After any shelf expansion, you get additional disks available for use in new LUNs by the Message Servers. See [Adding LUN to a Message Server](#) on page 27 for steps to take after adding new disks.

XSAN Filer Failover

In the event that one XSAN Filer fails and the other takes over, Mirapoint Message Servers running MOS 3 reboot due to loss of the LUNs, then try to access the LUNs through the standby Filer. Mirapoint Message Servers running MOS 4 with SAN MultiPath continue to use the LUN through other available paths, without the need to reboot.

There is some performance degradation in Filer failover situations. The proxy path through the active node introduces latency, and the active node must take over all the load from the failed node.

Operational Procedures

This section describes adding a LUN, adding a Message Server, backup, and restore.

Adding LUN to a Message Server

Adding a LUN involves nine steps, each in a subsection below.

Calculate Amount of Additional Space Required

In the equations below, *X* is the new required total size of the LUNs in the volume, in TB. The maximum LUN size is 8 TB. The LUN limit per Message Server is five. *Z* is the percentage of the volume allocated for snapshots, such as 40.

For example:

$$\text{Required volume size} = 100 * X / (100 - Z)$$

For instance, to get 3 TB of storage with 40% snapshot allocation:

$$\text{Required volume size} = 100 * 3 / (100 - 40) = 3 / 0.6 = 5 \text{ TB}$$

Add Disks to Aggregate Making it the Required Size

In this subsection, *S* is the required volume size you just calculated, and *msXX* is the name of the message store to be grown, for example *ms14*. *V* is the volume size, which you obtain with the command **df -Vg volume**, in this case 3400 GB.

```
XSANa01*> df -Vg ms14
Filesystem                total      used   avail capacity  Mounted on
/vol/ms14/                 3400GB    2059GB   1340GB     61%   /vol/ms14/
/vol/ms14/.snapshot        0GB       0GB      0GB      ---%   /vol/ms14/.snapshot
```

Using **vol status**, find the aggregate where the volume lives, in this case **aggr0**.

```
XSANa01*> vol status -v ms14
Volume State   Status           Options
ms14 online    raid_dp, flex   nosnap=on, nosnapdir=off,
                                minra=off, no_atime_update=off,
                                nvfail=off, ignore_inconsistent=off,
                                snapmirrored=off, create_ucode=on,
                                convert_ucode=off, maxdirsize=31457,
                                schedsnapname=ordinal,
                                fs_size_fixed=off,
                                guarantee=volume, svo_enable=off,
                                svo_checksum=off, svo_allow_rman=off,
                                svo_reject_errors=off, no_i2p=off,
                                fractional_reserve=0, extent=off,
                                try_first=volume_grow

Containing aggregate: 'aggr0'
                                Plex /aggr0/plex0: online, normal, active
                                RAID group /aggr0/plex0/rg0: normal
```

You can find the amount of free space in the aggregate, in this case *aggr0*, with the command **df -Ag aggrX**. In this case the answer is 110 GB.

```
XSANa01*> df -Ag aggr0
Aggregate                total      used   avail capacity
aggr0                    3550GB    3440GB   110GB     97%
aggr0/.snapshot          286GB     1GB     285GB     0%
```

Add disks to the aggregate, in this case *aggr0*, until there is (*S* - *V*) space available. See [Appendix A, Usable Space](#) for a table to help you estimate the number of disks required in an aggregate to achieve a desired size.

```
XSANa01> aggr add aggr0 -d diskId1 ...
XSANa01> df -Ag aggr0
...
```

Ensure There Are Sufficient Hot Spares Left

Check status of the disks, and that there are sufficient unallocated disks left, to be used as hot spares. There should be one unallocated disk per shelf of 14 disks.

```
XSANa01> aggr status -r
```

Grow the Volume

To grow the volume, use the **vol size** command. Specify the volume name, in this case /vol/ms14, and the required volume size *S* from above, in this case 5000 GB.

```
XSANA01*> vol size /vol/ms14 5000G
vol size: Flexible volume 'ms14' size set to 5000g.
```

Create the LUN

L is the number of the LUN to be created, and *msXX* is the volume to grow.

```
lun create -s 1023g -t linux /vol/msXX/lunL
```

In this case, the third LUN for the Message Server ms14 is numbered LUN 3:

```
XSANA01*> lun create -s 1023g -t linux /vol/ms14/lun3
lun create: created a LUN of size: 1023.1g (1098572103680)
```

Present the LUN

L is the number of the LUN to be created, in this case LUN 3, *msXX* is the volume to grow, *cLXX* is the name of the **igroup** representing the N+1 cluster, *I* is the LUN ID for the new LUN; it should be one greater than the LUN ID for the previous Message Server LUN. Find the LUN ID using the first command, then present the LUN with the second command:

```
lun show -m
lun map /vol/msXX/lunL cLXX I
```

In this case:

```
XSANA01*> lun show -m
```

LUN path	Mapped to	LUN ID	Protocol
/vol/mira0/lun0	mira-all	0	FCP
/vol/ms14/lun1	c101	11	FCP
/vol/ms14/lun2	c101	12	FCP

```
XSANA01*> lun map /vol/ms14/lun3 c101 13
XSANA01*> lun show -m
```

LUN path	Mapped to	LUN ID	Protocol
/vol/mira0/lun0	mira-all	0	FCP
/vol/ms14/lun1	c101	11	FCP
/vol/ms14/lun2	c101	12	FCP
/vol/ms14/lun3	c101	13	FCP

Check that LUN is Visible to Message Server

To check LUN visibility on a Mirapoint Message Server running MOS 3, use the **Storage Scan** and **Storage Listsan** commands:

```
ms14.stg> Storage Scan
OK Completed
ms14.stg> Storage Listsan
0.0.0.0 NETAPP 3MB masked clear "" ""
0.0.0.1 NETAPP 35840MB masked standby ch01.stg.. c101
0.0.0.21 NETAPP 1047676MB masked system ms15.stg.. c101
0.0.0.22 NETAPP 1047676MB masked inuse ms15.stg.. c101
0.0.2.0 NETAPP 3MB masked clear "" ""
```

```

0.0.2.11 NETAPP 1047676MB bound system ms14.stg... c101
0.0.2.12 NETAPP 1047676MB bound inuse ms14.stg... c101
0.0.2.13 NETAPP 1047676MB masked clear "" ""
2.0.0.0 NETAPP 3MB masked clear "" ""
2.0.0.1 NETAPP 35840MB masked standby ch01.stg... c101
2.0.0.21 NETAPP 1047676MB masked system ms15.stg... c101
2.0.0.22 NETAPP 1047676MB masked inuse ms15.stg... c101
2.0.1.0 NETAPP 3MB masked clear "" ""
2.0.1.11 NETAPP 1047676MB bound system ms14.stg... c101
2.0.1.12 NETAPP 1047676MB bound inuse ms14.stg... c101
2.0.1.13 NETAPP 1047676MB masked clear "" ""
OK Completed

```

Mirapoint Message Server running MOS 4.x with SAN MultiPath will find that the above example devices 2.0.0.x and 2.0.1.x will not appear in the output of **Storage Listsan**. To see the LUNs, **reboot** the MOS 4.x appliance.

The new LUN should be visible. It can be identified by its Mirapoint format LUN ID, *M* in the first column. The LUN is visible twice because there are two HBAs. For the next step, pick the lower number because it represents the primary HBA.

Attach Extra LUN to the Message Server

M is the Mirapoint format LUN ID from above. The **Storage Configsan** command reboots the Message Server, taking it offline for approximately ten minutes while it adjusts partition tables.

```
Storage Configsan M Grow
```

In this case:

```

ms14.stg> Storage Configsan 0.0.2.13 Grow
Warning: non-'grow*' operations will reboot the system. Continue? (y/n) y
Checking prerequisites
Growing store
OK Started

```

Verify Backup Policies

Update all your backup policies to reflect the new LUN for this Message Server.

Adding New Message Server to Existing N+1 Cluster

Adding a new Message Server involves 12 steps, each in a subsection below.

Is a New Aggregate Required?

There should be a maximum of two active Message Servers per aggregate. If an aggregate already exists as *aggrX* supporting only one active Message Server, and *aggrX* uses no more than 7 TB of disk space, the new Message Server should be added to this aggregate.

Otherwise, create a new aggregate.

Calculate Amount of Additional Space Required

In the equations below, X is the new required total size of LUNs for the new Message Server, in TB. The maximum LUN size is 1 TB. The LUN limit per Message Server is four. Z is the percentage of the volume to be allocated for snapshots, such as 40.

$$\text{Required volume size} = 100 * X / (100 - Z)$$

In this case for 2TB of storage and 40% snapshot allocation, calculate S :

$$\text{Required volume size} = 100 * 3 / (100 - 40) = 3 / 0.6 = 5 \text{ TB}$$

Create a New Aggregate if Required

Perform this step only if a new aggregate is required.

1. Find existing aggregates on the system

```
XSANa01*> df -Ag
Aggregate          total      used      avail capacity
aggr0              3550GB    3440GB    110GB     97%
aggr0/.snapshot   286GB     1GB      285GB     0%
```

2. Create a new aggregate, in this case aggr1, with 26 disks RAID size:

```
XSANa01*> aggr create -t raid_dp -r 26 aggr1
```

Allocate New Disks to the Aggregate

S is the required volume size you just calculated. Add disks to the aggregate, until there is at least that much space available. For a table to help you estimate the number of disks required in an aggregate to achieve the desired amount of usable space, See [Appendix A, Usable Space.](#)

Up to three shelves, each of 14 x 300GB disks, should be allocated to an aggregate. This provides space for two Message Servers, each with 2 x 1TB LUNs, using 40% allocation for snapshots.

```
XSANa01> aggr add aggr1 -d diskId1 ...
XSANa01> df -Ag aggr1
...
```

Ensure There Are Sufficient Hot Spares Remaining

Check status of the disks, and that there are sufficient unallocated disks left, which are used as hot spares. There should be one unallocated disk per shelf of 14 disks.

```
XSANa01> aggr status -r
```

Create a New Volume

S is the required volume size from above, and $msXX$ is the new Message Server:

```
vol create /vol/msXX aggrX S
```

In this case, for Message Server ms16:

```
XSANa01> vol create /vol/ms16 aggr1 5000g
```



```
XSANA01> snap reserve ms16 0
```

Do not create a fractional reserve for snapshots. For example:

```
XSANA01> vol options ms16 fractional_reserve 0
XSANA01> snap sched ms16 2 6
```

Setting a fractional reserve means that if long-lived snapshots are created, it is possible to fill the LUN, which would lead to an ENOSPC (out of space) error on the LUN when a write is attempted, thus causing a service outage on the Message Server. Ensure that processes always delete snapshots once they are backed up.

Create LUNs and Present Them

Create 1 TB LUNs as required for the new mail store to achieve the desired size:

```
XSANA01*> lun create -s 1023g -t linux /vol/msXX/lun1
XSANA01*> lun create -s 1023g -t linux /vol/msXX/lun2
...
```

To present the LUNs, you must know which LUN ID to use. In this case, use ID 1 for the standby head of the cluster, 11-14 for the first active head in the cluster, 21-24 for the second active head, etc. In the following, A1 denotes the LUN ID for the first LUN, A2 the LUN ID for the second LUN, and *c/XX* is the name of Igroup corresponding to the N+1 cluster, in this case *c/01*.

```
XSANA01*> lun map /vol/ms16/lun1 c/01 A1
XSANA01*> lun map /vol/ms16/lun2 c/01 A2
...
```

Update the Igroups

There is an Igroup associated with the cluster, *c/XX*, which is also the name of the N+1 cluster, in this case *c/01*. You must add WWPNs of both HBAs on the new Message Server to this Igroup.

1. Get the WWPNs from the Mirapoint Message Server:

```
ms16.stg> Storage Getport port0 Portwwn
0x210000e08b82f9bb
OK Completed
ms16.stg> Storage Getport port1 Portwwn
0x210000e08b821fbf
OK Completed
```

2. Add these WWPNs to the Igroup on the XSAN Filer:

```
XSANA01> igroup add c/01 0x210000e08b82f9bb
XSANA01> igroup add c/01 0x210000e08b821fbf
```

3. With a clustered XSAN, repeat step 2 on the other Filer node.

Attach the New Message Server to the LUNs

The Message Server should be freshly installed, with a MOS CD in the drive and the SAN and either Dual HBA or Multipath licenses installed.

MOS 3.x appliances use the **Storage Scan** and **Storage Listsan** commands to check that the LUNs are visible.

The new LUN should be visible. It can be identified by the Mirapoint format LUN ID, *M* in the first column. Note that the LUN should be visible twice, since there are two HBAs. The lower number is the primary HBA.

```
ms16.stg> Storage Scan
OK Completed
ms16.stg> Storage Listsan
0.0.0.0 NETAPP 3MB masked clear "" ""
0.0.0.1 NETAPP 35840MB masked standby ch01.stg.swi.contact cl01
0.0.0.21 NETAPP 1047676MB masked system ms15.stg.swi.contact cl01
0.0.0.22 NETAPP 1047676MB masked inuse ms15.stg.swi.contact cl01
0.0.2.0 NETAPP 3MB masked clear "" ""
0.0.2.11 NETAPP 1047676MB bound system ms14.stg.swi.contact cl01
0.0.2.12 NETAPP 1047676MB bound inuse ms14.stg.swi.contact cl01
0.0.2.31 NETAPP 1047676MB masked clear "" ""
0.0.2.32 NETAPP 1047676MB masked clear "" ""
2.0.0.0 NETAPP 3MB masked clear "" ""
2.0.0.1 NETAPP 35840MB masked standby ch01.stg.swi.contact cl01
2.0.0.21 NETAPP 1047676MB masked system ms15.stg.swi.contact cl01
2.0.0.22 NETAPP 1047676MB masked inuse ms15.stg.swi.contact cl01
2.0.1.0 NETAPP 3MB masked clear "" ""
2.0.1.11 NETAPP 1047676MB bound system ms14.stg.swi.contact cl01
2.0.1.12 NETAPP 1047676MB bound inuse ms14.stg.swi.contact cl01
2.0.1.31 NETAPP 1047676MB masked clear "" ""
2.0.1.32 NETAPP 1047676MB masked clear "" ""
OK Completed
```

The LUN IDs to use here are 0.0.2.31 and 0.0.2.32. LUN IDs 2.0.2.31 and 2.0.2.32 are from secondary HBA. Attach the Message Server to the LUNs as follows:

```
ms16.stg> Storage Configsan 0.0.2.31 Init
```

The system goes through several steps including a reboot. When it finishes, attach the second LUN. This takes the Message Store offline approximately ten minutes for partition table adjustment.

```
ms16.stg> Storage Configsan 0.0.2.32 Grow
```

Repeat for additional LUNs, if any, in order. Each takes the Message Store offline approximately ten minutes, as before. In this case there are

```
ms16.stg> Storage Configsan M Grow
```

In this case, there are no additional LUNs, only alternate routes to the same LUNs.

Configure the New Message Server and Test

Do this according to the process for adding new production Message Servers at your site. Consider licensing, configuration, patches, and enabled services.

Verify Backup Procedures

Update all your backup policies to reflect the new LUN for this Message Server.

Add the Message Server to the N+1 Cluster

On the new active Message Server, join the cluster. In this case, cl01 the name of the N+1 cluster, and the name of the Igroup.

```
ms16.stg> Cluster Attach cl01
```

On the standby head, scan the cluster to recognize the new active head:

```
ms00.stg> Cluster Scan
```

The new Message Server is now part of the N+1 failover cluster.

Backup a Message Server

Suppose *msXX* is the Message Server to be backed up. The following must be carried out on the node that owns this volume. By convention, the Igroup named “media-servers” is the Igroup containing all the NetBackup media servers.

Take a Snapshot of the Volume

```
XSANA01*> snap create -V msXX nightly1
```

```
XSANA01*> snap list -V msXX
```

```
Volume msXX
```

```
working...
```

%/used	%/total	date	name
0% (0%)	0% (0%)	Mar 19 16:43	nightly1

List the LUNs in the Volume

List all the LUNs and pick those that match the Message Server to be backed up.

```
XSANA01*> lun show
```

/vol/mira0/lun0	4m (4194304)	(r/w, online, mapped)
/vol/ms14/lun1	1023.1g (1098572103680)	(r/w, online, mapped)
/vol/ms14/lun2	1023.1g (1098572103680)	(r/w, online, mapped)

Backup Each LUN in the Volume

In this procedure, you clone, present, backup, and discard the LUN.

1. Clone the LUN:

```
XSANA01*> lun clone create /vol/msXX/backup1unY -b /vol/msXX/lunY backup1 -o noreserve
```

2. Present the LUN to the media servers with LUN ID, *L*:

```
XSANA01*> lun map /vol/msXX/backup1unY media-servers L
```

3. From one of the media servers, it should now be possible to scan and see the LUN, then back it up to tape.
4. Once backup is complete, discard the LUN:

```
XSANA01*> lun destroy /vol/msXX/backup1unY
```

Repeat these steps for each LUN in volume Y. In this case Y = 1, 2.

Delete Clone-Created Snapshots

You cannot auto-delete snapshots that are in use by a clone-created LUN.

To delete a clone-created snapshot:

1. Ensure that your backup scripts always release the clone-created LUNs.
2. Delete the clone-created snapshot to avoid the danger when a write attempts of the volume filling and the LUNs returning an ENOSPC (out of space) error, which causes a service outage on the Message Server.

```
XSANa01*> snap delete -V msXX backup1
```

Restore a Message Server for Mailbox Restore

Dedicated space is allocated on the Filers for a mailbox restore.

To restore a Message Server for mailbox restore:

1. Identify sufficient 1 TB LUNs that are not currently in use.
2. Restore the LUNs from tape using the NetBackup media servers.
The LUNs should already be presented to the Mirapoint “restore head(s)”.
3. Ensure that the Mirapoint is freshly installed with the correct version of MOS, a System CD in the drive, plus SAN and dual HBA licenses.
4. Re-attach it to the LUNs using the existing restore process.

Troubleshooting

Unexpected Loss of LUN

Note: for 4.x systems, skip this section. Due to Multipath this does not occur.

With clustered XSAN Filers, during an XSAN failover on the XSAN, all Mirapoint systems that have LUNs on the failed XSAN will lose the path to their LUNs and reboot. This is expected behavior, though it creates an issue.

Some Mirapoint systems that have LUNs on the other XSAN Filer, which did not fail, also lose the paths to their LUNs and reboot. This occurs because the target IDs of the paths that are presented by XSAN change after the XSAN Filer failover. Mirapoint systems must reboot to start using the new paths. The following examples show path details.

Before XSAN Filer failover:

```
mp3.mirapoint.com> storage listsan
2.0.0.0 NETAPP 4MB masked clear "" ""
2.0.0.1 NETAPP 736308MB masked system mp1.mirapoint.com ""
2.0.0.2 NETAPP 730128MB masked system mp2.mirapoint.com ""
2.0.0.3 NETAPP 736256MB bound system mp3.mirapoint.com ""
```

After XSAN Filer failover:

```
mp3.mirapoint.com> storage listsan
```

```
2.0.1.0 NETAPP 4MB masked clear "" ""
2.0.1.1 NETAPP 736308MB masked system mp1.mirapoint.com ""
2.0.1.2 NETAPP 730128MB masked system mp2.mirapoint.com ""
2.0.1.3 NETAPP 736256MB bound system mp3.mirapoint.com ""
```

Dual HBA Configuration

The double-switch configuration shown in [Figure 1](#) on page 10 requires Dual HBA. For 3.x systems, Mirapoint recommends this Dual HBA configuration.

Multipath Configuration

Multipath permits all Message Servers to see any available LUN via any path through the SAN fabric. Like Dual HBA on 3.x, Multipath on 4.x requires a license. For 4.x systems, Mirapoint recommends the Multipath configuration.

Factory Reset

In rare circumstances, you could run into a situation where you cannot rectify the XSAN Filer configuration, so you must perform a factory reset.

There are two types of resets:

- ◆ Resetting the configurations except aggregate and volume. This is useful to start from scratch for the high level settings.
- ◆ A complete reset to factory form. This is required when you assign the disks wrongly to aggregate and leave no spare behind, or when you configure a small Raidset that ends up consuming more disks for parity.

High Level Reset

To perform the first type of reset, run these commands:

```
XSANa01> priv set advanced
XSANa01> halt -c factory
```

Complete Factory Reset

This is a time consuming process. It can take hours to erase all the disk drives.

1. From the serial port CLI, halt the XSAN and reboot.
2. Select the special boot menu with Ctrl+C when prompted.
3. Select option 4a to erase all disk drives and re-install a fresh operating system with aggregate.
4. Reboot from flash memory and step through setup.
5. Download the Data Ontap image to the fresh vol0 with HTTP, NFS or CIFS and run a normal install.
6. Reboot.

Usable Space

The following table provides an estimate of the amount of usable disk space for a FAS3020c with 300GB FCAL disks. The left hand column is the number of disks allocated to an aggregate. This excludes hot spares and is based on RAID-DP with a RAID size of 26.

The following figures do not include any allocation for snapshots. Any such allocation must be deducted from these figures. For instance, with 18 disks and a 40% allocation for snapshots, you have $3633 * (1 - 0.4) = 2179\text{GB}$ available after accounting for snapshots.

To calculate space for other types of disks, or Filer configurations, contact Mirapoint for more information.

Table 6 Usable Space Per Disk

No of Disks in Aggregate	Usable Space (300GB FCAL disks)
1	required for parity
2	required for parity
3	227
4	454
5	681
6	908
7	1135
8	1362
9	1589
10	1816
11	2043
12	2270
13	2497
14	2725
15	2952
16	3179

Table 6 Usable Space Per Disk

No of Disks in Aggregate	Usable Space (300GB FCAL disks)
17	3406
18	3633
19	3860
20	4087
21	4314
22	4541
23	4768
24	4995
25	5222
26	5450
27	required for parity
28	required for parity
29	5677
30	5904
31	6131
32	6358
33	6585
34	6812
35	7039
36	7266
37	7493
38	7720
39	7947
40	8175
41	8402
42	8629
43	8856
44	9083
45	9310
46	9537

Table 6 Usable Space Per Disk

No of Disks in Aggregate	Usable Space (300GB FCAL disks)
47	9764
48	9991
49	10218
50	10445
51	10672
52	10900
53	required for parity
54	required for parity
55	11127
56	11354
57	11581
58	11808
59	12035
60	12262
61	12489
62	12716

