



Unified Storage Server

v3.0.0



User's Guide

Installation & Basic Configuration

References

1. Chelsio iSCSI Target User's Guide.
2. IETF iSCSI standard RFC 3720.

Copyright and Trademarks

Copyright © 2013 by Chelsio Communications, Inc.

All rights reserved.

All terms in this document that describe a third party brand or trademark are the properties of their respective owners, (for example, 'Linux' is owned by Linus Torvalds).

The terms 'Chelsio', 'Chelsio Communications' and the Chelsio logo are registered trademarks of Chelsio Communications, Inc.

Address for communication:

Chelsio Communications, Inc., 370 San Aleso Avenue, Sunnyvale, CA 94085, U.S.A.

TEL: +1 (408) 962-3600, FAX: +1 (408) 962-3661, support@chelsio.com, www.chelsio.com

Copyright ©2013.Chelsio Communications. All Rights Reserved.

CONTENTS

1. INTRODUCTION	8
2. HARDWARE	11
2.1 RECOMMENDED HARDWARE	11
2.2 RESERVED MEMORY MATRIX	13
2.3 CONFIGURING THE HARDWARE	14
2.4 INSTALLING THE PRODUCT	14
3. FIRST BOOT	16
4. WEB-BASED MANAGEMENT	17
4.1 ACCESSING THE MANAGEMENT INTERFACE	17
4.2 LAYOUT AND NAVIGATION	22
4.3 LICENSING	24
5. HARDWARE & SOFTWARE FEATURES MATRIX	33
5.1 INSTALLATION AND BOOT	33
5.2 NETWORK CONTROLLER SUPPORT	34

5.3	STORAGE CONTROLLER SUPPORT	36
5.4	STORAGE MANAGEMENT/VIRTUALIZATION SUPPORT	39
5.5	HIGH AVAILABILITY / CLUSTERING	42
5.6	SYSTEM SUPPORT	43
5.7	CLIENTS SUPPORTED	44
5.8	CLIENT PROTOCOLS SUPPORT	46
6.	NETWORKING	48
6.1	NETWORKING OVERVIEW	48
6.2	CONFIGURE AND CONTROL NETWORK DEVICES AND SETTINGS	52
6.3	NETWORK TEAM CONFIGURATION	65
6.4	ACTIVE CONNECTIONS MONITORING	75
6.5	NETWORK TROUBLESHOOTING TOOLS.....	79
6.6	DHCP SERVER	89
7.	STORAGE	93
7.1	STORAGE OVERVIEW	93
7.2	STORAGE MANAGEMENT	96

7.3	STORAGE CONFIGURATION	112
7.4	OS DISK DEVICES CONFIGURATION	114
7.5	SOFTWARE RAID ARRAY CONFIGURATION	133
7.6	ISCSI INITIATOR	144
7.7	HARDWARE RAID ARRAY CONFIGURATION	157
7.8	FIBRE CHANNEL.....	175
7.9	VOLUME MANAGEMENT	201
7.10	REPLICATION.....	242
7.11	FILE SYSTEMS	256
7.12	LUSTRE	270
7.13	BACKUP (NON-HA)	275
8.	ISCSI SAN	300
8.1	ISCSI OVERVIEW	300
8.2	ISCSI DEPLOYMENT	301
8.3	CONFIGURING THE CHELSIO ISCSI TARGET	302
8.4	ISCSI SUMMARY	309
8.5	CREATING A NEW TARGET	313

8.6	ISCSI TARGET SUMMARY INTERFACE	317
8.7	ISCSI TARGET LUN CONFIGURATION	328
8.8	ISCSI TARGET NETWORK PORTALS.....	337
8.9	ISCSI TARGET CLIENTS / INITIATORS	342
9.	USS HIGH AVAILABILITY	352
9.1	CREATE CLUSTER	355
9.2	CLUSTER RESOURCES	363
9.3	CLUSTER SERVICES	365
10.	FILE SHARING	366
10.1	FILE SHARING OVERVIEW.....	366
10.2	GLOBAL SETTINGS.....	375
10.3	FILE SHARING CONFIGURATION	379
10.4	USER MANAGEMENT.....	409
11.	SYSTEM TOOLS	420
11.1	SYSTEM TOOLS USAGE	420

12. APPENDIX **462**

12.1 TROUBLESHOOTING **462**

12.2 CHELSIO END USER LICENSE AGREEMENT **475**

12.3 GNU GENERAL PUBLIC LICENSE..... **480**

12.4 OTHER LICENSING INFORMATION **490**

1. Introduction



Chelsio's Unified Storage Server is a powerful easy-to-use turnkey solution for creating high performance file and block storage solutions. USS is designed to meet the storage performance and ROI challenges faced by data center cluster environments, including high-performance computing environments. It is a middleware that is best-of-breed in the market, and provides an easy integration path for VARs and OEMs, and ease of use for end users.

Key Features and Benefits

- ▶ Deploy storage systems in minutes – The first-time setup wizard makes it simple to connect to the network, configure email alerts, set administrator password, etc.
- ▶ Plug-and-play – Integrates easily into VAR/OEM's hardware platform, ensuring smooth storage system integration. Comes as a bootable flash memory or loadable software. Fully compatible with most AMD64/EM64T multi-processor systems.

- ▶ Ease-of-use in deploying and reconfiguring of the storage array – For system administrators/network administrators in small- and medium-sized businesses. Unified Storage has an intuitive web-based management interface, which is accessible in any compatible web browser, over an encrypted secure connection, providing ease-of-use and requiring minimum training
- ▶ Low TCO - Chelsio's Unified Storage makes possible some of the lowest cost per gigabyte solutions in the industry. Consolidating multiple file servers and iSCSI SAN onto a single device reduces server management overhead and associated IT staff costs. Network storage can be remotely managed using a Web-based user interface, simplifying maintenance and providing centralized control of processes like backups, restores, and upgrades
- ▶ Flexible branding capability feature for OEMs

The Unified Storage Server product is designed to provide the following features to the end-user, utilizing the least amount of time and effort in deploying and managing the appliance.

- ▶ iSCSI SAN target/initiator services.
- ▶ FC SAN target/initiator services.
- ▶ NFS, CIFS, FTP, HTTP, Lustre file sharing services.
- ▶ Dynamic storage allocation and management.
- ▶ Snapshots of storage data, for zero-downtime backups, and snapshot scheduling.
- ▶ Replication of data to a peer Unified Storage Server appliance, for disaster recovery.

- ▶ Extensive hardware support for various network and storage devices.
- ▶ Optimized for the AMD64 / EM64T 64-bit multi-processor architecture.
- ▶ Plug-n-play setup of the Unified Storage Server product, allowing for quick deployment and ease of use.
- ▶ Migration of data between disk volumes on the fly.
- ▶ Boost performance with data caching.

This is achieved by a combination of Chelsio's high-performance, scalable storage stack, and the Linux operating system, and management software. The product fully supports Chelsio T3 and T4 based Unified Wire adapters, for full protocol offload, delivering maximum performance.

In the following sections, the configuration and management of the Unified Storage Server appliance will be explained, with deployment scenarios illustrated, to get the appliance integrated into your environment smoothly.

2. Hardware

2.1 Recommended hardware

Depending on the type of applications / clients using the Unified Storage Server appliance, the hardware subsystem needs to be able to support the bandwidth and latency required. Following is a recommended hardware platform for typical usage scenarios.

- ▶ 1st usage model: Low-cost, high-capacity storage, with workgroup-level workload / applications

CPU:	1 dual core processor, AMD Opteron / Intel EM64T Xeon
Memory:	4GB+
Network:	2 x 1Gbps Ethernet (Chelsio T422-CR recommended)
Storage:	SATA 3Gbps RAID controller, with high-capacity SATA 3Gbps hard disks

- ▶ 2nd usage model: Low-latency, enterprise-level performance, with departmental / data center workload

CPU:	2 dual core processors, AMD Opteron / Intel EM64T Xeon
Memory:	8GB+
Network:	2 x 10Gbps Ethernet
Storage:	SAS dual channel RAID controller, with 15000 RPM SAS hard disks.
HA:	Intel(R) Xeon(R) CPU E5520 @ 2.27GHz, SuperMicro X8DTS-F motherboard, 12GB DDR2 1333MHz Memory

The performance of the storage subsystem is very crucial to the overall performance of the appliance. Ensuring that the storage subsystem is capable of meeting your application needs, is critical to a smooth deployment.

2.2 Reserved memory matrix

Memory usage by a USS appliance may differ based on system configuration and can be calculated by referring to the table below:

Volume Management type	Pool without caching	Pool with Caching	Replication
Logical Volume Management (LVM)	No additional memory reserved	Not supported	33MB additional memory reserved per 1 TB
Thin Provisioning (TP)	700 MB additional memory reserved per pool	1 GB additional memory per pool + RAM cache size mentioned	33MB additional memory reserved per 1 TB

- Usage examples:
 - i. If a pool size of 32 TB with 1 GB RAM cache size is created, then the system requires minimum 8GB + 1GB RAM cache + 1 GB = 10 GB memory.
 - ii. On the above pool, if a 32 TB volume is created and replicated, then the system requires minimum 10GB + 1GB=11 GB memory.

2.3 Configuring the hardware

For ensuring maximum uptime and fault tolerance, the storage subsystem needs to be connected and configured in the following manner:

- ▶ Ensure there are failover paths to the hard disks from the storage controller. This is possible by using an enclosure / backplane which allows multiple paths to the hard disk drives.
- ▶ Use a hardware RAID controller or software RAID, instead of a standard SAS / SATA / SCSI controller.
- ▶ RAID (Redundant Array of Independent Disks) allows for failure of 1 or more hard disks, based on the RAID policy chosen for the group of hard disks. It may also improve performance, based on the type of data access. Refer the Storage section of this guide for RAID configurations.
- ▶ Ensure there are multiple network paths to the appliance, using multiple physical connections. Network load-balancing + failover with LACP, and iSCSI MPIO configurations are recommended. Details on iSCSI MPIO are available in the iSCSI Section of this guide.

2.4 Installing the product

If the Unified Storage Server is provided on a CD/DVD ROM media, insert this into your server's CD/DVD drive and select the drive as *first boot device* in your system's BIOS. Now the server will boot from

the CDROM, and provide an option to install the product onto a hard drive or USB flash drive. Refer **Quick Start guide** for first time setup instructions.

3. First boot

Please refer the Quick Start guide for configuring your Appliance to boot to Unified Storage correctly and configuring it for the first time. On powering on the appliance for the first time, you will be greeted with the **System Setup wizard**. The administrative username is '**root**'. The wizard consists of the following steps:

1. Welcome page
2. Chelsio End User License Agreement acceptance
3. Administrative password change
4. Date, time, time zone, and network time synchronization settings
5. Alerting configuration
6. Hostname, Network Devices, DNS servers configuration

This wizard is designed in an intuitive manner, to get the appliance integrated into your computing environment quickly.

4. Web-based Management

4.1 Accessing the Management Interface

The management interface is designed with a goal of being '*accessible anywhere, securely*'. Due to the pervasive nature of the Internet and the World Wide Web, web browsers are the most common application easily available to IT administrators. The Management Interface is fully web-based, and uses secure 256-bit encrypted HTTP, ensuring that authentication and configuration data are protected during transmission from the web browser to the appliance and vice versa.

Currently supported browsers are Mozilla Firefox 2.0+, Microsoft Internet Explorer 7+, Opera 9+, Apple Safari 3.4+. The Management Interface is accessed by typing in the URL `https://<appliance hostname | IP address>`, in a web browser.

Only the administrative user (root) is allowed to login and configure the appliance. This ensures that the configuration cannot be changed by any other user. The password length and complexity should be good enough for ensuring that it cannot be guessed.

The security certificate used by the appliance's web server is a generic one. It can cause the following types of prompts in different browsers. You will need to select the correct option to continue.

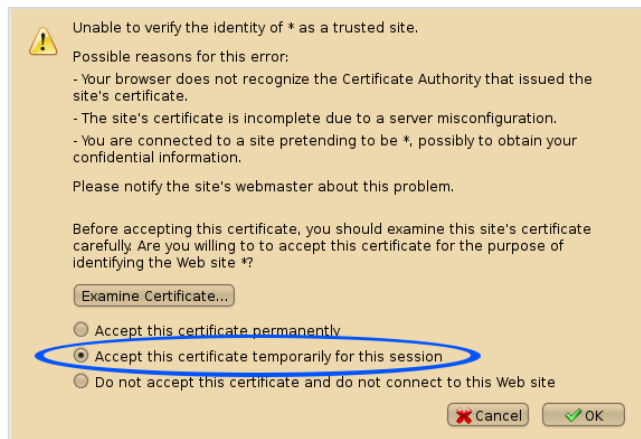


Figure 4.1 (a) – Security Certificate prompt in Firefox 2



Figure 4.1 (b) – Security certificate prompt in Safari 3.2.

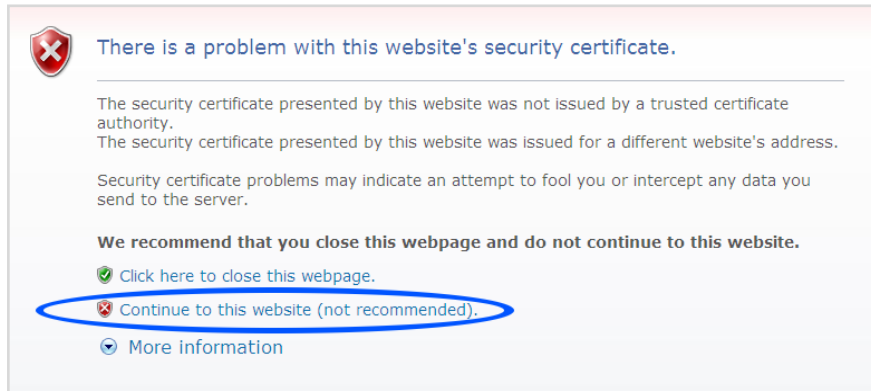


Figure 4.1 (c) – Security certificate prompt in Internet Explorer 7.

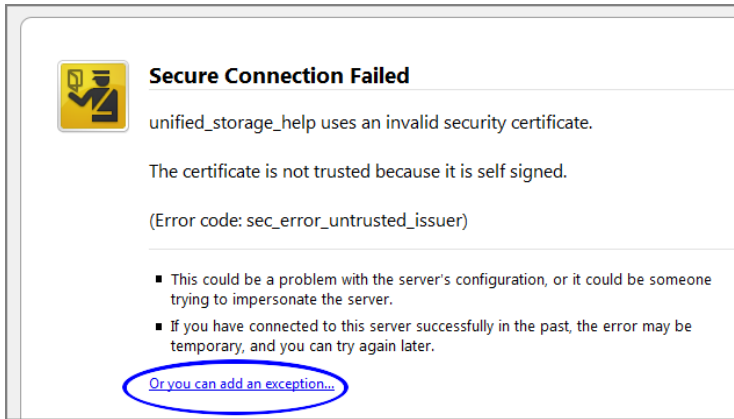


Figure 4.1 (d) – Security certificate prompt in Firefox 3




Figure 4.1 (e) – Add exception




Figure 4.1 (f) – Get certificate from server.




Figure 4.1 (g) – Confirm exception



Welcome to Unified Storage Server

Login to Unified Storage Server on

Administrative password:



Note: Please ensure that this website is added to your trusted sites list, or javascript is enabled for this website.

Figure 4.1 (h) – Login screen for the Management Interface.

4.2 Layout and navigation

The layout of the interface is organized into three panes. The upper area is a banner, with the right corner having a Logout link, a link to this guide (Help), and a support link. Below that, is a navigation menu, and a contents section.

The navigation menu is on the left, with a cascading tree of links to various configuration modules, and the right pane is used to display the content of each menu item.

There are 3 tabs available for each menu item: The 'Configuration' tab is selected by default, and shows the configuration page for the module highlighted in the menu. Switching to the Help tab shows page / context sensitive help text. The 'Event Logs' tab displays Alerts, Errors and Warnings for various Configuration events. Instant notifications for errors will be displayed as *Critical Alerts* under the same tab. The log of notifications can be downloaded.

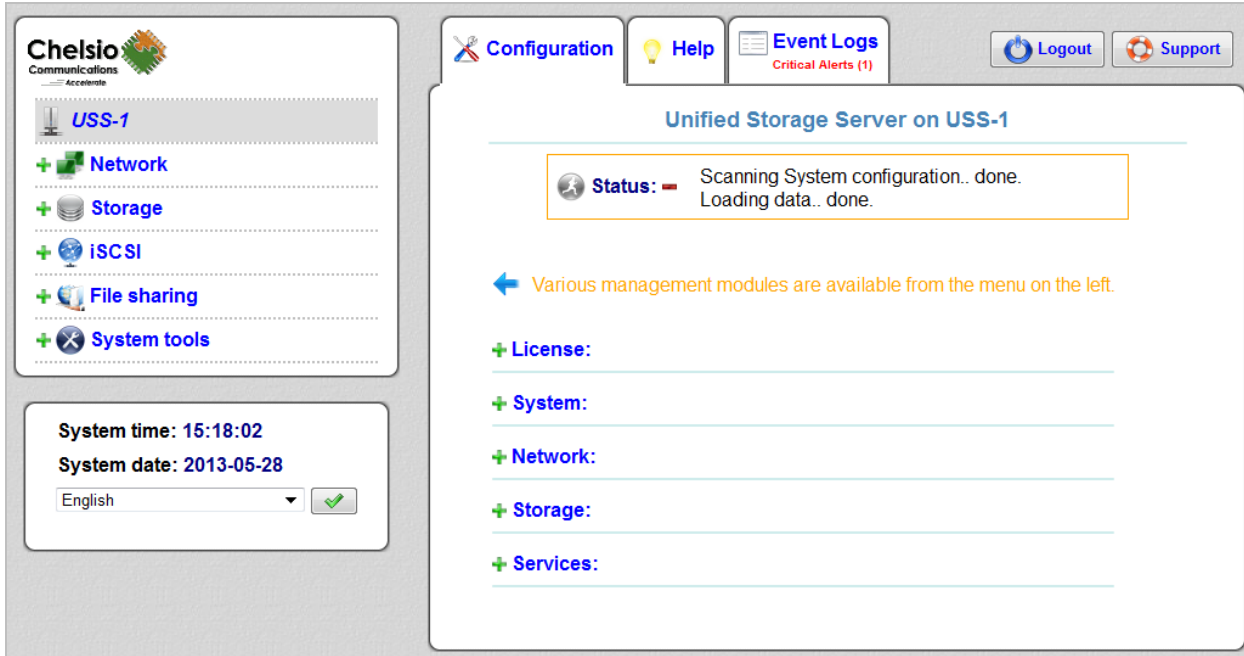


Figure 4.2 – Management Interface Layout

4.3 Licensing

4.3.1 License overview

You can try USS free for 30 days and test drive the product. The trial version is fully functional and all features are available. After the trial period, you will have to apply for a product license to continue using the product.

In addition to Product license, you can also apply for a Maintenance license. This will entitle you to technical help and support from the Chelsio Support team for any product related issues.

If the product was shipped to you with a product key, you can start configuring the Storage, iSCSI and NAS modules of the Appliance. The license details are displayed in the Welcome page, once you complete the System Setup wizard and login.

If you need to apply for a license, please follow the instructions provided in the next section.

Note: To use the full version of USS, please make sure that you have Chelsio unified wire adapter installed in your system, since the product will not work with network adapters of other vendors. Although not recommended, you can however use non-Chelsio network adapters during the trial period for evaluation purpose.

Note: As stated before, USS trial version is for evaluation purpose and hence may not perform to its full potential, especially when used with non-Chelsio network adapters.

License:

iSCSI target:

Evaluation license for 0 years, 0 months, 31 days

Unified Storage Server:

Evaluation license for 0 years, 0 months, 31 days [6 days used]

Network device licensed:



Figure 4.3.1 (a) – Evaluation License details displayed on the Welcome Page

License:

iSCSI target:	Production license
Unified Storage Server:	Production license
iSCSI target:	Chelsio iWARP RDMA transport
Unified Storage Server:	Clustering
Features:	Fibre Channel Target
	File sharing
	IOCACHE
	Replication
	Snapshots, Cloning
	Thin Provisioning
Storage Capacity:	100000 TB
Network device licensed:	00:07:43:05:71:06


 **Licensing wizard**

Figure 4.3.1 (b) – Production License details displayed on the Welcome Page

4.3.2 Obtaining a License

On the **Welcome page** of the Management Interface, under the **License** section, click on the **Licensing wizard** button to obtain a License key. A Chelsio Network Device needs to be selected to associate the license with. So in the **Request new license section**, select a network device among the available Chelsio devices in the list and click Next. This will generate a hardware information file.

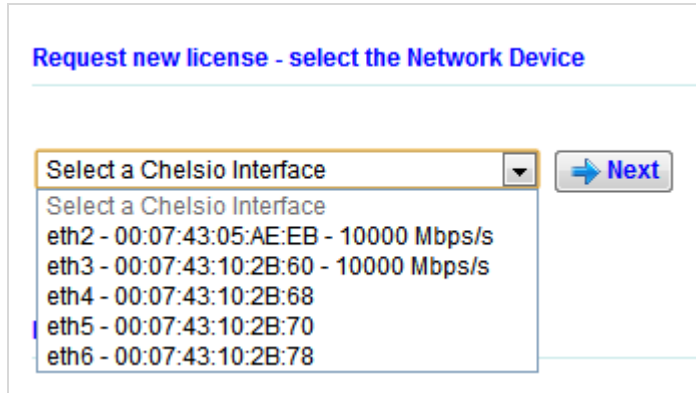


Figure 4.3.2(a) – Select the Chelsio network device to associate with the License from the drop-down menu.

Now, choose one of the two options to contact Chelsio and obtain a License key.

- 1. Download the hardware profile information file:** Download the information file to local drive. Visit <http://service.chelsio.com/licenses> and upload the file along with other information and the key will be sent via email within 24 hrs. The License file can also be obtained by emailing the file to the Chelsio Support Team at support@chelsio.com or your sales representative.


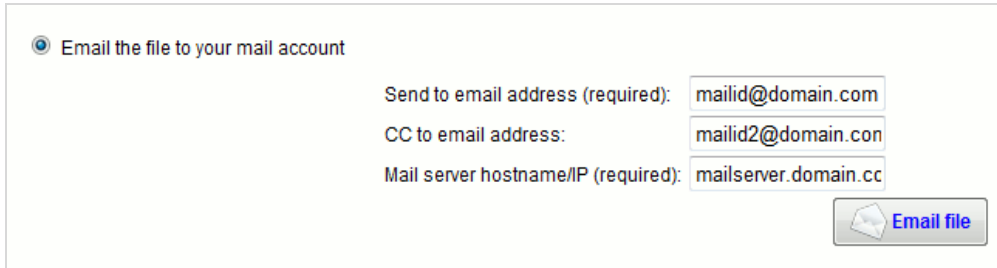


Figure 4.3.2(b) – *Download the hardware profile information file and contact Chelsio to obtain the License*

- 2. Email the file to your mail account:** You can choose the hardware profile information file to be emailed to a specific email account as an attachment. A valid SMTP server IP address or hostname is required. The License key can be obtained by either visiting

<http://service.chelsio.com/licenses> or contacting Chelsio Support Team as mentioned in option 1 above.



The screenshot shows a web form with a radio button selected for "Email the file to your mail account". Below this are three input fields: "Send to email address (required):" with the value "mailid@domain.com", "CC to email address:" with the value "mailid2@domain.com", and "Mail server hostname/IP (required):" with the value "mailserver.domain.cc". A blue "Email file" button with an envelope icon is located at the bottom right of the form.

Figure 4.3.2(c) – Emailing *hardware information file* it to an email account

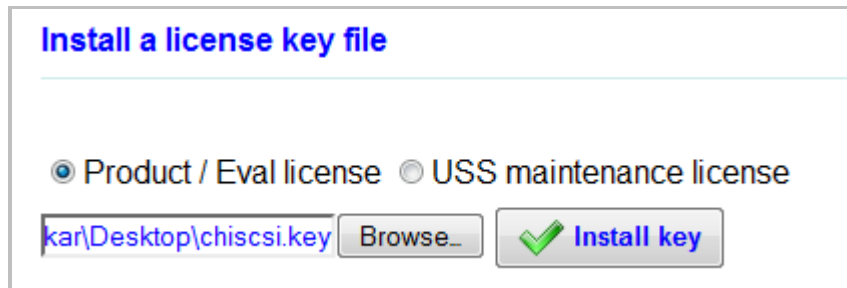
4.3.3 Installing License

The license file sent by Chelsio support team is bound to the network adapter used while generating hardware profile information file. Hence, please ensure that the same adapter is present before installing the license. Now, follow these steps:

1. Save the License key file in your local system.

2. In the **Install a license key file** section in the **Licensing wizard**, select the type of license to be installed: Product or USS Maintenance license.
3. Click **Browse** and locate the file.
4. Click **Install key**.
5. The license will now be installed. After successful installation, the navigation menu and configuration tab will be refreshed.

The **License** section will display details such as type of license installed, features, maintenance license validity (if installed), storage capacity and MAC address of the Chelsio network card to which the license key is bound. The navigation menu will display configuration modules depending on the hardware present (Hardware RAID, Fibre Channel, etc) and license installed.



Install a license key file

Product / Eval license USS maintenance license

kar\Desktop\chiscsi.key

Figure 4.3.3(a) – Uploading license key to USS appliance



Warning: Licensing is designed to capture changes in time and will deny further operation, if time zone is changed multiple times.

4.3.4 Dependencies

- Chelsio Thin Provisioning (TP) License

USS provides Thin Provisioning (TP) and Logical Volume Management (LVM) features to manage your SANs. However, you can license your USS Appliance with either TP or LVM. By default, the appliance will be licensed with LVM. To enable the TP feature, please contact Chelsio Support Team at support@chelsio.com or your sales representative.

- Lustre Support

- i. Lustre over RDMA is currently supported on Chelsio T4 Network Interface Cards only.
- ii. Lustre over RDMA is not supported on bond interface.

- High Availability (HA)

To use USS's HA feature, the appliance needs to be licensed with Chelsio Thin Provisioning (TP), since HA with Logical Volume Management (LVM) is not supported. Hence, make sure you apply for a TP license along with a HA license. HA is currently not supported on FC and Hardware RAID Adapters.

- Replication

Replication of data is possible only if the volume type of systems involved is same. For example, if the local system has TP pool/volume, the peer system should also have TP pool/volume. The same applies for LVM. Combining different kinds of volume types is not supported.

5. Hardware & Software features matrix

5.1 Installation and Boot

Features	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
Bootable CDROM	Y	Y	Y	Y	Y
USB CD/DVD drive	Y	Y	Y	Y	Y
Space required on HDD	1GB	4GB	4GB	4GB	4GB
USB flash boot device	Y	Y	Y	Y	Y
SATA flash boot device	Y	Y	Y	Y	Y
IDE disk boot device	Y	Y	Y	Y	Y
LSI MegaRAID boot device	Y	Y	Y	Y	Y

3ware boot device	Y	Y	Y	Y	Y
LSI SAS boot device			Y	Y	Y
S/W RAID 1 mirror boot device			Y	Y	Y
Upgrade install from CD from previous versions		Y	Y	Y	Y

5.2 Network controller support

Features	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
T3 adapters support	Y	Y	Y	Y	Y
T3 RDMA support	Update pack	License opt	License opt	License opt	License opt
T4 adapters support			Y	Y	Y
T4 RDMA support			License opt	License opt	License opt
Support for	Y	Y	Y	Y	Y

iSCSI/NFS/CIFS/FTP/HTTP on Chelsio NIC					
Support for iSCSI/NFS/CIFS/FTP/HTTP on 3rd party NIC	Y	Y			Trial version only
Bonding support	Y	Y	Y	Y	Y
T3 Offload bonding support	Y	Y	Y	Y	Y
T4 Offload bonding support			Y	Y	Y
Offload bonding modes	Active-Standby, 802.3ad	Active-Standby, 802.3ad	Active-Standby, 802.3ad	Active-Standby, 802.3ad	Active-Standby, 802.3ad

5.3 Storage controller support

Features	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
Hardware RAID					
LSI MegaRAID SAS 3Gb/s adapters	Y	Y	Y	Y	Y
LSI MegaRAID SAS 6Gb/s adapters		Y	Y	Y	Y
Adaptec/PMC SAS 6Gb/s adapters		Y	Y	Y	Y
3ware SATA/SAS 3Gb/s adapters	Y	Y	Y	Y	Y
3ware SATA/SAS 6Gb/s adapters			Y	Y	Y
LSI MegaRAID firmware update			Y	Y	Y

support					
Non-RAID					
LSI SAS 3Gb/s onboard/adapter	Y	Y	Y	Y	Y
LSI SAS 6Gb/s onboard/adapter		Y	Y	Y	Y
FC					
Emulex 8Gbps LPE12002 initiator mode	Y	Y	Y	Y	Y
Emulex 8Gbps LPE12002 target mode		Y	Y	Y	Y
Qlogic 8Gbps initiator mode		Y	Y	Y	Y
Emulex adapter firmware update support			Y	Y	Y
Qlogic adapter firmware update support			Y	Y	Y
iSCSI Initiator					
Chelsio T3 based adapters			Y	Y	Y
Chelsio T4 based adapters			Y	Y	Y

SSD controller support					
Micron SSD					Y
HDD support					
SAS: Seagate Cheetah 36GB 15K RPM 3.5"	Y	Y	Y	Y	Y
SAS: Seagate Cheetah 73GB 15K RPM 3.5"	Y	Y	Y	Y	Y
SAS: Seagate Cheetah 146GB 15K RPM 3.5"	Y	Y	Y	Y	Y
SAS: Fujitsu 146GB 15K RPM 3.5"		Y	Y	Y	Y
SAS: IBM/Hitachi 500GB 7.2K RPM 2.5"		Y	Y	Y	Y
SATA: Seagate Barracuda 1TB 7.2K RPM 3.5"	Y	Y	Y	Y	Y
SATA: Seagate Barracuda 2TB 7.2K RPM 3.5"		Y	Y	Y	Y
SATA: Seagate Barracuda 3TB 7.2K RPM 3.5"			Y	Y	Y
Maximum HDDs connected to 1 LSI RAID controller	240	240	240	240	240

Maximum HDDs connected to 1 Adaptec RAID controller	256	256	256	256	256
Maximum HDDs connected to 1 LSI SAS-only controller	512	512	512	512	512

5.4 Storage management/virtualization support

Features	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
Linux Volume Management (deprecated)	Y	Y	Y	Y	Y
Chelsio Thin Provisioning Volume Management (recommended)		Y	Y	Y	Y
Manual snapshots	Y	Y	Y	Y	Y
Instant snapshots (requires Chelsio TP)		Y	Y	Y	Y
Instant restore from snapshot		Y	Y	Y	Y

(requires Chelsio TP)					
Instant volume clone (requires Chelsio TP)		Y	Y	Y	Y
Snapshot scheduling (requires Chelsio TP)		Y	Y	Y	Y
Max snapshots per volume LVM	1	1	1	1	1
Max snapshots/clones per volume Chelsio TP		254	254	254	254
Max storage pool size (requires Chelsio TP)		1PB (1024TB)	1PB (1024TB)	1PB (1024TB)	Unlimited (2 ⁶⁴ bytes)
Max volume size		1PB (1024TB)	1PB (1024TB)	1PB (1024TB)	Unlimited (2 ⁶⁴ bytes)
Max filesystem size - LVM	8 EB (1 million TB)	8 EB (1 million TB)	8 EB (1 million TB)	8 EB (1 million TB)	8 EB (1 million TB)
Max filesystem size - Chelsio TP		1PB (1024TB)	1PB (1024TB)	1PB (1024TB)	Unlimited (2 ⁶⁴ bytes)
Max open files per client	64k	64k	64k	64k	64k
NAS backing filesystem	XFS	XFS	XFS	XFS	XFS

Remote replication/mirroring	Y	Y	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)
Remote replication - bandwidth management	Y	Y	Y	Y	Y
Remote replication - compression		Y	Y	Y	Y
Remote replication - scheduling		Y	Y	Y	Y
S/W RAID support	Y	Y	Y	Y	Y
S/W RAID alerting and auto-rebuild on drive restore		Y	Y	Y	Y
S/W RAID levels	0,1,5,10	0,1,5,6,10	0,1,5,6,10	0,1,5,6,10	0,1,5,6,10
RAM caching support (requires Chelsio VM / TP)					Y
SSD caching support (requires Chelsio VM / TP)					Y

5.5 High availability / Clustering

Features	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
HA support (requires Chelsio TP)		Y	Y	Y	Y
HA support for iSCSI/NFS/CIFS/FTP/HTTP		Y	Y	Y	Y
HA shared storage controller: SAS		LSI SAS-only	LSI SAS-only	LSI SAS-only	LSI SAS-only
HA shared storage controller: iSCSI			Chelsio T3,T4	Chelsio T3,T4	Chelsio T3,T4
Cache support on HA			Y	Y	Y
Cache coherency/replication over TCP for HA			Y	Y	Y
HA reboot failed node support: IPMI-LAN		Y	Y	Y	Y

5.6 System support

	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
IBM x3650	Y	Y	Y	Y	Y
IBM x3650m2		Y	Y	Y	Y
IBM x3650m3		Y	Y	Y	Y
Dell 2950	Y	Y	Y	Y	Y
Dell R710		Y	Y	Y	Y
Supermicro X7 series	Y	Y	Y	Y	Y
Supermicro X8 series		Y	Y	Y	Y
Supermicro SBB		Y	Y	Y	Y

5.7 Clients supported

	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
iSCSI - Microsoft initiator	Y	Y	Y	Y	Y
iSCSI - RHEL 5.x open iSCSI initiator	Y	Y	Y	Y	Y
iSCSI - RHEL 6.x open iscsi initiator		Y	Y	Y	Y
iSCSI - Chelsio T3 initiator		Y	Y	Y	Y
iSCSI - Chelsio T4 initiator			Y	Y	Y
iSCSI - Qlogic 1GbE initiator		Y	Y	Y	Y
iSCSI - Mac OS global san initiator		Y (no CHAP)	Y (no CHAP)	Y (no CHAP)	Y (no CHAP)
iSCSI - VMware ESX 4.x 5.0 initiator		Y	Y	Y	Y
FC - Emulex LPE12002 8Gb/s initiator		Y	Y	Y	Y

FC - Qlogic 8Gb/s initiator		Y	Y	Y	Y
NFS - RHEL 4.x	Y	Y	Y	Y	Y
NFS - RHEL 5.x	Y	Y	Y	Y	Y
NFS - RHEL 6.x		Y (requires sysctl mod)	Y (requires sysctl mod)	Y (requires sysctl mod)	Y (requires sysctl mod)
NFS - SLES 10.x	Y	Y	Y	Y	Y
NFS - SLES 11.x		Y	Y	Y	Y
CIFS - Win XP	Y	Y	Y	Y	Y
CIFS - Win Vista		Y	Y	Y	Y
CIFS - Win 7		Y	Y	Y	Y
CIFS - Win 2003	Y	Y	Y	Y	Y
CIFS - Win 2008R2		Y	Y	Y	Y
Lustre - RHEL 5.x with Lustre 1.8.x		Y	Y	Y	Y

5.8 Client Protocols support

	Version				
	v1.10.49	v2.0.0-120	v2.1.0-xxx	v2.2.0-xxx	v3.0.0-xxx
iSCSI target	Y	Y	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)
NFS v3 over TCP	Y	Y	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)
NFS v3 over RDMA	Y (Chelsio T3 only)	Y (Chelsio T3 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)
CIFS over TCP	Y	Y	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)
FTP	Y	Y	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)

HTTP		Y	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)
Lustre over TCP		Y	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)	Y (Chelsio T3/T4 only)
Lustre over RDMA			Y (Chelsio T4 only)	Y (Chelsio T4 only)	Y (Chelsio T4 only)

6. Networking



6.1 Networking overview

Generally, the appliance should be connected to high-speed Ethernet networks of 1GbE or 10GbE speeds. It is better to have the **appliance connected to the network segments directly, for the clients that it will be serving most often**. For example, if you have iSCSI initiator clients on network segment 192.168.1.0/24 and CIFS clients on network segment 10.1.2.0/8, you would want to have the appliance connected to the 192.168.1.0/24 and 10.1.2.0/8 network segments over Ethernet switches directly, without any routers in between. This ensures that there is sufficient network bandwidth between the clients and the appliance for effective data transfers. Refer to figures 6.1 a, b for typical network topologies.

DHCP is **not** recommended as a method of configuring the network interfaces, unless you are configuring the DHCP server to assign a specific IP address to the interface every time. iSCSI and NAS services will not work if the IP addresses used in their configuration and by clients keep changing.

The fully qualified hostname and DNS servers should be configured for NAS services to function correctly. The DNS server should be configured with the appliance's IP addresses too, so that clients can access the appliance using the hostname instead of the IP address. DNS and other settings are available in the Global settings page under the Network section. Do not change the TCP, IP, ARP protocol settings in the Global settings page, unless you specifically need to do so. In most cases, the default settings work correctly.

The appliance does not have any network firewall mechanism, since it will generally reside in the core of the network. By default, it will listen on ports or accept any connections for the following services: iSCSI, NFS, CIFS, FTP and remote management over HTTP, HTTPS, and remote login over SSH.

If you wish to firewall the appliance, configure an intermediary firewall device that will be connecting the appliance to the network. Refer the Firewall's documentation to allow iSCSI, NFS, CIFS, FTP, HTTP, HTTPS and SSH traffic through to the Appliance.

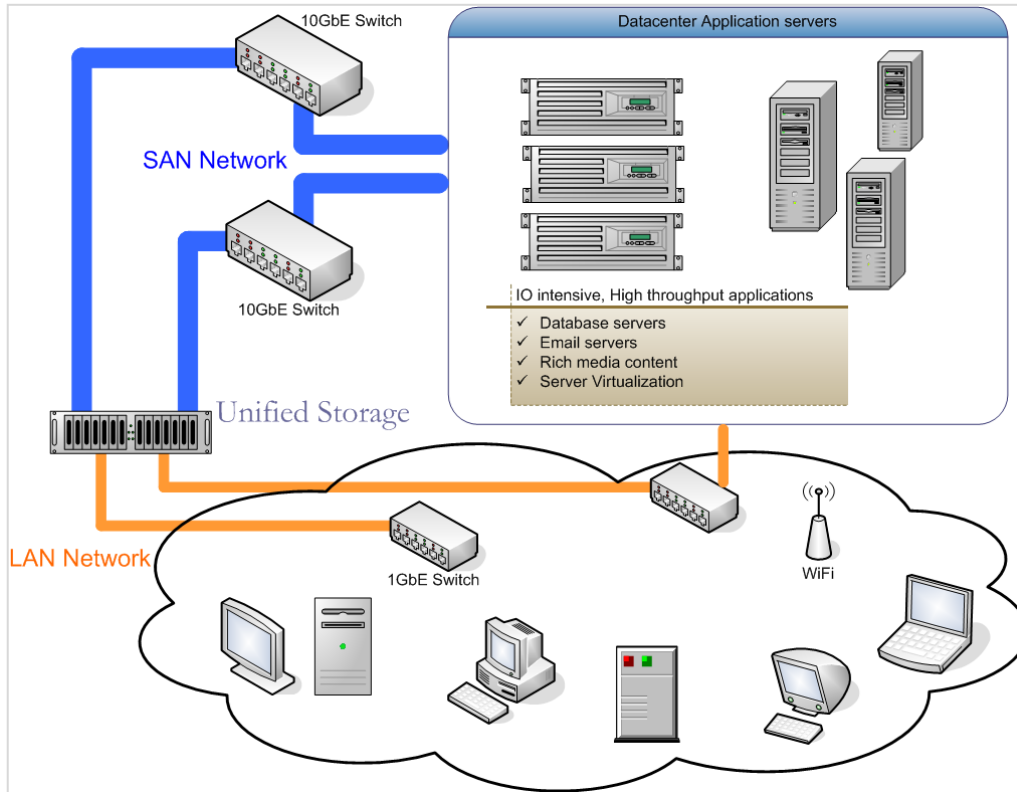


Figure 6.1 (a) – Redundant / Multi-pathed, High performance network topology for Unified Storage Server.

Separate SAN and LAN networks, to ensure optimal SAN performance.

10GbE SAN network provides low latency for high transaction rates, and large bandwidth.

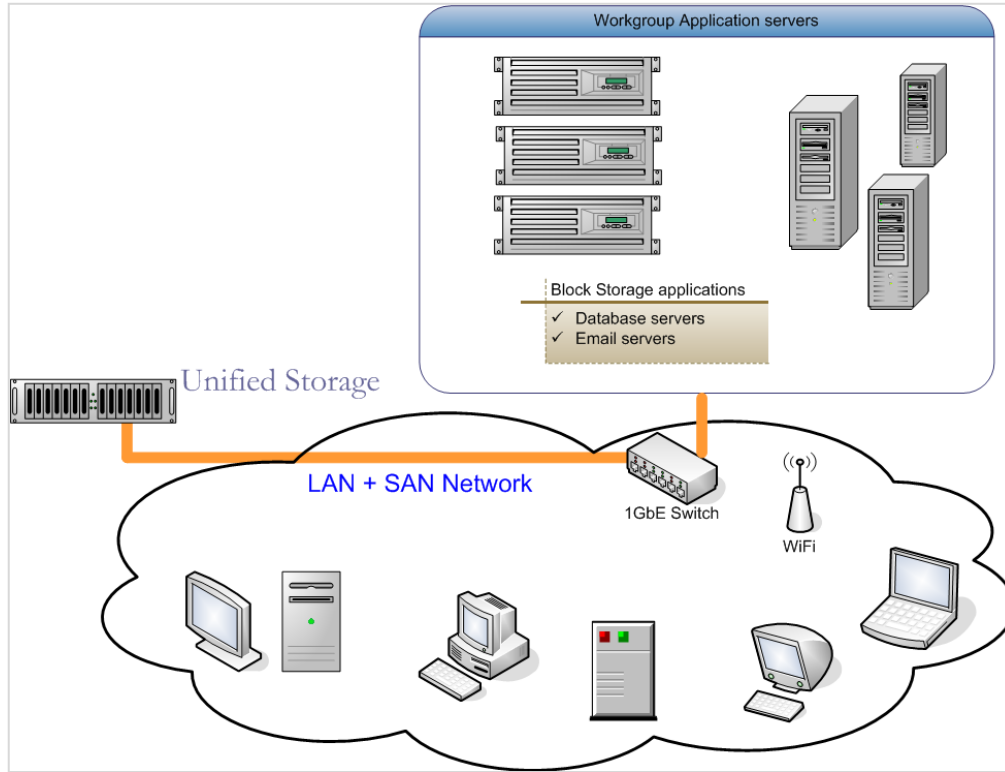


Figure 6.1 (b) – *Low cost, easily deployable network topology for Unified Storage Server.*

Integrated SAN and LAN network, with minimum investment in new hardware.

iSCSI SAN traffic can be isolated from LAN traffic using a VLAN configuration on the switch.

6.2 Configure and control Network Devices and settings

All networking devices that are currently detected and have a valid device driver loaded are listed in the Network summary page under the Devices section. If there is an installed device that is not shown here, check in the system logs if there were any problems loading its device driver, and contact Support, with the most recent log files included in the communication.

Devices that are yet to be manually configured default to DHCP configuration. All devices are enabled at system boot by default. A device that is causing problems can be disabled from starting at boot, so that you can start it when required, for troubleshooting.

The localhost network interface 'lo' is a system interface, which cannot be altered. It is preconfigured and required for the system to function correctly.

Ensure that the default gateway is set for the appropriate interface. **Do not specify a default gateway for multiple interfaces.** It is an invalid configuration. The system can have one default gateway that is used for all unknown networks. For specific non-local networks, add a routing rule in the 'Network Troubleshooting' page.

- **Sections of the interface**

- 1. Summary**

The number of devices present, number of active and inactive devices, IP addresses currently assigned to the Appliance, and bandwidth status is listed here. A count of the active TCP connections to the Appliance and DNS servers summary is also shown. IP addresses can be configured for upto 3 DNS Servers. Hostname can also be configured here.

Summary:

Ethernet Devices details:

Ethernet status: **8 total, 3 active, 5 with no link.**
IP address: **169.254.0.202, 10.193.185.202, 102.1.1.202**
Bandwidth status: **1 x 100 Mb/s | 2 x 10000 Mb/s**
Chelsio Storage Accelerator: **Total 4 ports present, 4 ports offload enabled**

Connections details:

Active TCP connections: **2**
Offloaded TCP connections: **0**

DNS & Hostname:

DNS status: **2 DNS servers configured**
Server 1:
Server 2:
Server 3:

Hostname: **Configured**

Figure 6.2(a) – Summary section with details of the devices, connections and DNS settings

2. Devices

The list of network devices on the system, including Network teams are listed here. The device's name and bandwidth are shown on the left, with an icon indicative of its current state. A text status, any configured IP address, and the physical MAC address is displayed in the center. Control options for the interface, and more information, is available using the buttons on the right. Configuration of the interface can be viewed and changed below, by expanding the configuration link.

In an HA setup with Supermicro SBB systems, altering backplane configuration will result in cluster failure. To prevent this, device configuration for the backplane interfaces will be unavailable after creation of cluster service.



Intel Corporation 82567LF-2 Gigabit Network Connection		
eth0  1Gb/s	Status: Physical Link up, IP configured IP address: 10.193.184.237/22 MAC Address: 00:1C:C0:9F:8A:89	Properties & Statistics Stop Restart
+ Configuration:		
Chelsio Communications Inc S310-CR 10GbE Single Port Adapter		
eth1  10Gb/s	Status: Physical Link up, IP configured IP address: 102.1.1.237/24 MAC Address: 00:07:43:05:C5:D7	Properties & Statistics Stop Restart
+ Configuration:		

Figure 6.2(b) – Device listings in ‘Devices’ section, with status, commands and configuration options




Network Team		Properties & Statistics <input type="button" value="Stop"/> <input type="button" value="Restart"/> <input type="button" value="Remove"/>
bond1 	Status: Physical Link up, IP configured IP address: 102.50.50.76/24 MAC Address: 00:07:43:11:7E:40 Mode: fault-tolerance (active-backup)	
Team members:		
Chelsio Communications Inc T422-CR Unified Wire Ethernet Controller		
eth2  10Gb/s	Status: Physical Link up, Device Active Role: Currently primary , designated primary MAC Address: 00:07:43:11:7E:40	Properties & Statistics <input type="button" value="Remove"/>
Chelsio Communications Inc T422-CR Unified Wire Ethernet Controller		
eth3  10Gb/s	Status: Physical Link up, Device Active Role: Currently secondary MAC Address: 00:07:43:11:7E:48	Properties & Statistics <input type="button" value="Remove"/>
+ Configuration:		

Figure 6.2(c) – Network team and members' listings with status, commands and configuration options


Intel Corporation 82576 Gigabit Network Connection (rev 01)		
eth0  100 Mb/s	Status: Physical Link up, IP configured	Properties & Statistics <input type="button" value="Stop"/> <input type="button" value="Restart"/>
	IP address: 10.193.185.76/22	
	MAC Address: 00:25:90:10:CD:2A	
	IPMI BMC LAN interface: yes	
	IPMI BMC LAN IP address: 10.193.185.77/22	
+ Configuration:		

Figure 6.2(d) IPMI LAN interface status, commands and configuration options



Backplane cluster connection 1		
eth0  10Gb/s	Status: Physical Link up, IP configured	Properties & Statistics <input type="button" value="Stop"/> <input type="button" value="Restart"/>
	IP address: 169.254.0.1/16	
	MAC Address: 00:25:90:14:2E:B8	

Figure 6.2(e) Backplane interface status and commands (Supermicro SBB)

2.1. Device commands

- Properties and statistics: This navigates to a new page, which provides details of the device, including firmware version, driver version, and supported options. There are 2 tabs in this page, one for properties of the device, and the other for detailed statistics that the device's driver reports. All network devices may not report the same statistics. Hence device specific statistics are grouped accordingly.



Property	Value
driver:	igb
Interrupt:	
version:	2.3.4
firmware-version:	1.11-5
MAC:	00:30:48:B9:46:97
Type:	Ethernet
MTU:	1500
arp_enabled:	yes
broadcast_enabled:	yes
multicast_enabled:	yes
port_type:	[TP]
port:	Twisted Pair
supported auto negotiation:	Yes
advertised auto negotiation:	Yes
current auto negotiation:	on

Figure 6.2(f) *Properties and Statistics page for a network device*

- **Start / Stop:** This command provides starting or enabling the device on the network, and applying its IP configuration. Ensure that an IP configuration is set, before trying to start a device. If the device is started / enabled, it allows stopping / disabling the device. Note that the device will be automatically enabled on reboot, if its configuration option “Device activated” is set to “on System startup”. This command prompts for confirmation from the user, to avoid being used accidentally.

Note: The option to stop a backplane interface (Backplane cluster connection) in a cluster setup for Supermicro SBB systems is not available, since stopping or disabling it will result in cluster failure.

- **Restart:** This command allows restarting a running device. If the device is not running, it will still try to disable the device first, thus flushing any previously associated IP address, etc., and then enable it. This command prompts for confirmation from the user, to avoid being used accidentally.



Warning: Please be aware that stopping or restarting a network device can cause all connected clients to lose connectivity, and possible loss / corruption of any data that the clients were saving.

2.2. Device configuration

Expand the configuration link to view the device configuration. The settings in the configuration are as below:

- Configuration type: There are two modes for assigning an IP address; through DHCP or a static IP address. For DHCP, a DHCP server should be present on the network and configured appropriately.
- IP address: IP v4 addressing is supported. Enter a valid IP address in dotted decimal notation (e.g.: 192.168.1.10), if you are configuring a static IP address. In DHCP mode, the current address is displayed, but it is not editable.
- Subnet mask: An IP v4 subnet mask is required, to correctly configure the IP address. Enter the subnet mask in dotted decimal notation (e.g.: 255.255.255.0), if you are configuring a static IP address. In DHCP mode, the current address is displayed, but it is not editable.
- Gateway: Configure a gateway on only one of the devices, to reach non-local networks, if required. USS appliance can be configured to function successfully without any internet connectivity.

- **Broadcast Address:** A broadcast address is not required for a valid networking configuration. But it may affect services which depend on broadcast, to locate and advertise the service.
- **Device activation:** This setting allows you to control if the device should be automatically activated on appliance startup (recommended). If a certain device is problematic, or requires special configuration before activation, or is not required, you can set it to start manually.
- **MTU:** The maximum transmission unit controls the amount of Bytes sent across the network in each Ethernet frame (logical segment of data). This setting is usually 1500 Bytes on most Ethernet networks. Changes in the MTU on the device, to a larger size, such as 9000 Bytes, usually referred to as “Jumbo frames”, are supported. But this will cause problems if the switches and clients on the network do not support the higher frame size.
- **Lustre Networking:** Use this setting to enable Lustre Networking over TCP or RDMA. TCP and RDMA modes are available for T4 adapters. Whereas for T3 adapters, only TCP mode is available.
- **VLAN Child Device:** VLANs give the ability to segregate LANs efficiently, by allowing multiple Virtual LANs on a single Ethernet or wireless interface. As VLAN works on OSI Layer 2, it can be used just as any other network interface without any restrictions. VLAN successfully passes through regular Ethernet bridges.

This setting allows you to add a VLAN child device. Using the VLAN device configuration, you can choose to add a child device in two ways:

i) New/blank configuration: With this setting, you can set up a VLAN child device without any values and configure the settings later using the Configuration option.

ii) Migrate <parent device>configuration: Using this setting you can copy the parent device's configuration to the VLAN child device you are creating.

Note: For systems with IPMI interfaces, the option to configure IPMI device settings is available in the *IPMI BMC* section under *System Tools*.

Note: If you're planning to install cluster service using Supermicro SBB systems, please ensure that backplane interfaces (Backplane cluster connection) have been configured correctly in the 169.254.x.x network with subnet mask 255.255.0.0 before attempting to create cluster. This is to ensure that backplane interface is not used as management interface.



Warning: It is not advisable to configure DHCP IP addressing for network devices that will be used for iSCSI traffic. Do so, only if you are reserving an IP address for that particular MAC address on the DHCP server, to ensure that the same IP address is assigned every time. iSCSI initiator clients require the Target IP address to remain constant, to access the disk at all times.

Configuration:

Configuration type:

IP address:

Subnet mask:

Gateway:

Broadcast address:

Device activated:

MTU:

Lustre networking:

Add VLAN child device:

VLAN device configuration:

- New / blank configuration
- Migrate eth4 configuration

Figure 6.2(g) Configuration settings for a network device

6.3 Network team configuration

A team of network devices can be created in the ‘Network Team’ page, which is visible on navigating to the Network section. A Network Team having multiple active Devices allows for load-

balancing of traffic and failover of network connectivity to the appliance. This is a recommended configuration for high availability and uninterrupted service to clients. The 'Link Aggregation' type of Team requires appropriate configuration on the network switch, to which the devices are connected. The Link Aggregation Control Protocol (LACP) has to be enabled on the switch, for the ports used by the appliance for the team. Refer the Network Switch user guide for details on how you can enable LACP, if it is supported. Network Team devices do not support DHCP, and require a static IP address setup.

Devices: (multi-select)

- eth0 [00:07:43:05:E9:93]
- eth1 [00:07:43:05:E9:94]
- eth2 [E4:1F:13:2C:97:98]
- eth3 [E4:1F:13:2C:97:9A]

Create new Team

IP Configuration type:

IP address:

Subnet mask:

Default Gateway:

Broadcast address:

Device activated:

Primary Team member:

Team mode:

Figure 6.3(a)-Devices and Configuration sections of the Create team page.



Warning: Switching from regular Ethernet to a team setup will cause loss of network connectivity till the appliance configuration (and any required switch configuration) is completed.



Warning: Please make sure that no traffic/service (e.g. iSCSI) is running on any of the member devices before proceeding to create a Network team.

Example: Creating a Network Team

Here's an example on how to create a Network Team:

- i. Select two or more network devices in the **Devices** section by clicking on them. This will display the configuration option for the team to be created on the right.

Note: The selected ports should either be of Chelsio devices of the same architecture or other devices excluding all ports of Chelsio devices.

- ii. Choose “Statically assign” if you want to manually configure the team or “Import from member device” to import the IP address configuration from one of the member devices in the team.
- iii. If you have chosen “Statically assign”, then provide the IP address, Subnet mask, Gateway, Broadcast, for the team.
- iv. Specify if you want to manually activate the device or automatically on system startup in the “Device activated” field.
- v. Specify the primary member for the team.
- vi. Select the Team mode and click **Apply**.

The newly created Network team will appear in the **Network Summary** module, under the **Devices** section.

- **Team configuration settings**

1. **Configuration type:**

This setting allows you to import the IP address configuration from one of the member devices in the team, which is useful when migrating from a single network connection to a network, to a team connected to the same network. Optionally you may manually specify the IP address configuration instead of importing it.

Devices: (multi-select)

- ✓ eth0 [00:19:B9:B0:08:65]
- ✓ eth1 [00:19:B9:B0:08:67]
- ✓ eth2 [00:07:43:05:E9:AF]
- ✓ eth3 [00:07:43:05:E9:B0]

Create new Team

IP Configuration type:

IP address:

Subnet mask:

Default Gateway:

Broadcast address:

Device activated:

Primary Team member:

Team mode:

Figure 6.3(b)-Creating a Network Team by statically assigning IP address.

Devices: (multi-select)

- eth0 [00:07:43:05:E9:93]
- eth1 [00:07:43:05:E9:94]
- eth2 [E4:1F:13:2C:97:98]
- eth3 [E4:1F:13:2C:97:9A]

Create new Team

IP Configuration type: Import from member device ▾

Import from member: eth2 ▾

Imported configuration type: Statically assign ▾

IP address: 10.193.184.157

Subnet mask: 255.255.252.0

Default Gateway: 10.193.184.1

Broadcast address: 10.193.187.255

Device activated: on System startup ▾

Primary Team member: eth2 ▾

Team mode: Balance-RR ▾

Apply

Figure 6.3(c)-Creating a Network Team by importing IP address configuration from another member device

2. IP address, Subnet mask, Gateway, Broadcast, Device activation

These settings are the regular settings for network devices, described in the upper level Network section's help.

3. Primary team member:

Certain team modes need to have a primary member, which can be assigned here.

4. Team mode:

There are six teaming modes available. The modes supported by the teaming driver depend on whether a Chelsio storage acceleration device is chosen as a member of the team or not. With a Chelsio device present, the modes supported are LACP / 802.3ad link aggregation, and Active-Standby.

Teaming modes details:

- **Balance - RR** (Load balancing – Tx round robin):

This mode utilizes all the links to transmit data, with even distribution of traffic across all the links. It also provides fault-tolerance.

- **Active Standby:**

Only one of the member devices in the team is active. A different member device will become active only on if the current active member fails. This mode does not utilize the bandwidth of all members of the team.

- **Balance - XOR (Load balancing – Tx XOR):**

Transmit load balancing based on an XOR of the member device’s MAC addresses.

- **Broadcast (duplicate Tx on all):**

This mode broadcasts all outgoing data on all member devices.

- **LACP / 802.3ad link aggregation (recommended):**

Link Aggregation Control Protocol, (also known as 802.3ad) allows for creating teams of network devices, where all devices become part of one logical Ethernet connection to the switch. For this to succeed, the switch needs to support LACP, and needs to be configured to use LACP on the ports that are connected to the team member devices on the appliance.

- **Balance - TLB** (Load balancing – adaptive Tx):

Outgoing data is transmitted across all links, and is evenly distributed across the links, by an algorithm based on the current load on that team member, taking into account the link speed / bandwidth available on that team member.

- **Balance - ALB** (Load balancing – adaptive Tx+Rx):

This mode uses the above algorithm to distribute outgoing transmit data across member devices, and uses an ARP update mechanism to allow peers on the network to transmit to a particular team member, thus achieving receive load balancing.



Warning: Do not configure a team with LACP / 802.3ad link aggregation, before configuring the respective switch ports to use LACP. The switch will refuse to accept LACP traffic if not configured to do so.

- **Apply settings:**

Once all the settings are specified, clicking apply will start the reconfiguration, to create the team device. Any current traffic on the member devices will be stopped, and the new team device will be made active with the configuration specified. It can become active only if atleast one of the member devices has a physical link present. Notice the “Status” box at the top of the page after clicking apply for the progress of the task, and any problems encountered.

6.4 Active Connections monitoring

The current TCP connections and listening TCP / UDP ports are listed in the 'Active Connections' page. This can be filtered to display a subset of the results. Offload status for each connection is also displayed, if Protocol offload hardware is available. This is useful for troubleshooting any connectivity issues for clients to various services.

- **Query options settings**

- 1. Resolve IP address to hostnames:**

If you wish to see the hostnames of the client systems connecting to the appliance, instead of the IP addresses, the hostnames will be resolved using DNS, and displayed. This is a resource-intensive activity if there are hundreds or thousands of connections. Those IP addresses that could not be resolved to a hostname, are shown as is.

- 2. Show listening services:**

There may be many processes / applications on the system waiting to connect to clients on the network. These processes are in a "Listening" state. To view which ports and IP addresses on the appliances are currently active, listening / waiting for a new connection, enable this option. To view which services are active, enable this and the first option "Decode application layer protocol".

3. Layer 4 protocols:

The two Layer-4 protocols are TCP and UDP. You may view listening and active connections for TCP, and only listening services for UDP, since UDP is not connection oriented (does not establish a lasting connection that can be listed here).

4. Layer 7 protocols:

This option allows filtering out other application protocols, and viewing connections only for a particular service, such as iSCSI / CIFS / NFS.

5. Decode application layer protocol:

This option will cause the listing to show a best guess of what application layer (Layer 7) protocol is running on this connection, based on the destination or source port. This depends on a mapping table based on well-known ports that are registered or used for certain application protocols (e.g.: port 80 is used by HTTP).

6. Refresh page duration:

The list of connections to the system is very dynamic, and requires constant updating. You may increase / decrease the frequency of refreshing the list here.

Query options:

Resolve IP addresses to hostnames:

(Note: Offload information cannot be shown with name resolution enabled)

Show listening services:

Layer 4 protocols:

Layer 7 protocols:

Decode application layer protocol:

Refresh page every:


 **Apply**

Figure 6.4 (a) - Query option settings to filter the connections listing.

- **Connections list**

The tabular list of currently active connections to the appliance is shown here. This list is not accurate for a very long period of time, since connections can get established and torn down very

quickly by clients. The page refreshes every 30 seconds, and you can make it refresh even faster if required.


Connections shown below: 7 (refreshes every 30 seconds)						
Layer 4 Protocol	Local IP	Local Port / Protocol	Remote IP	Remote Port / Protocol	State	TCP Offload
tcp	10.193.184.237	39322	10.193.184.187	445	ESTABLISHED	no
tcp	10.193.184.237	443	10.193.191.180	57007	ESTABLISHED	no
tcp	10.193.184.237	443	10.193.191.180	57009	ESTABLISHED	no
tcp	10.193.184.237	443	10.193.191.180	57010	ESTABLISHED	no
tcp	10.193.184.237	443	10.193.191.180	57011	CLOSE-WAIT	no
tcp	10.193.184.237	52765	10.193.184.187	49158	ESTABLISHED	no
tcp	10.193.184.237	55543	10.193.184.187	389	ESTABLISHED	no

Figure 6.4 (b) – Active Network connections list.

6.5 Network troubleshooting tools

- **Ping utility**

This utility allows troubleshooting many network connectivity issues. Specify a hostname or IP address to ping, and it will try to contact that system. The result of the ping will be displayed in the same section.




The screenshot shows a web-based ping utility interface. On the left, the text "Ping a hostname or IP address:" is displayed. To its right is a text input field containing "www.ietf.org". Below the input field, the text "(IP address or Hostname)" is shown. A button with a green checkmark icon and the word "Ping" is positioned below the input field. To the left of the results, the text "Ping result for www.ietf.org:" is displayed. To its right, the results are shown in green text: "success.", "5 packets transmitted, 5 received, 0% packet loss, time 4024ms", and "rtt min/avg/max/mdev = 285.489/287.793/289.508/1.481 ms".

Figure 6.5(a) - Ping utility with results of a ping test.


- **DNS name resolution (nslookup) utility**

This utility allows for troubleshooting name resolution issues. If a hostname is not reachable via ping, try to see if its hostname resolves to a valid IP address. This also allows verifying that the DNS settings you have provided (in the first boot setup wizard or in “Advanced settings” below) are valid.

The current ping response status of the configured DNS servers is also displayed here. (Note: some servers may not respond to ping due to a firewall, but could be reachable on the network for a different protocol such as DNS.)



Resolve hostname to IP address:
(Hostname: eg.: example.mydomain.com)

 **Search**

Host www.ietf.org IP address resolution result: **64.170.98.32**

Current DNS servers status: **1 reachable**

Figure 6.5(b) - Name resolution utility with results of a lookup.

- **Advanced networking settings**

The settings are grouped into DNS settings, ARP settings, IP settings, TCP settings, and Socket settings. Highlighting a setting will display the field where it can be edited and saved.

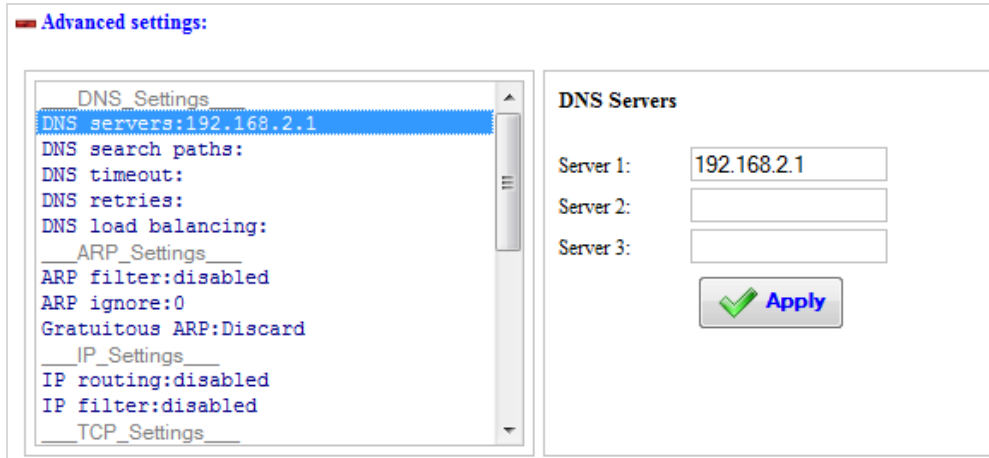


Figure 6.5(c) - Advanced settings list with editable fields displayed on the right for a highlighted setting.

1. DNS Settings

- DNS servers: You may specify upto 3 DNS servers for the system to use for DNS resolution of hostnames to IP addresses.
- DNS search paths: This controls the domain suffixes that are tried for DNS resolution, when only the hostname is available. The DNS domain of the system should preferably be the first search path entry. To edit the DNS domain of the system, edit the system hostname in the “System summary” page.
- DNS timeout: The timeout in seconds for a query to a DNS server.
- DNS retries: The number of retries for a query that is timing out.
- DNS load balancing: Enabling this will evenly distribute DNS queries across all the DNS servers configured.

2. ARP settings

- ARP filter: Enabling this causes ARP responses to only be sent on the device that received the ARP request. By default, ARP responses are sent on all connected network devices.

- ARP ignore: This controls the behavior of the system in responding to ARP requests. By default, there is no restriction on how the system responds.
- Gratuitous ARP: Systems may send ARP updates, without being asked. These can be used to update the ARP cache table, or discarded. The default behavior is to discard it.

3. IP settings

- IP routing / forwarding: Do not enable this setting, unless you are sure of what you are doing. This causes the system to forward IP packets between its network devices, acting like a router. This may cause loss of connectivity for clients, to services, and other network issues
- IP filter: Enabling this causes the system to respond to IP packets only from the received interface. This is recommended for certain networking configurations, but will cause issues with the presence of a Network team.

4. TCP settings

- TCP Recv Mem: Do not alter these settings, unless you are sure of what you are doing. This controls the amount of system memory used for TCP receive buffers. Larger values are

recommended for the minimum and default, (about 256KB = 262144 or 512KB = 524288) for high-bandwidth networks such as 10GbE.

- TCP Send Mem: As above, for TCP send / transmit operations.
- TCP moderate rcvbuf: If enabled, the system will automatically tune the buffer sizes based on the traffic pattern (recommended).
- TCP selective ACK: Reduces the number of acknowledgements sent if enabled. Otherwise, every received TCP segment is acknowledged with a response, which causes additional overhead.
- TCP duplicate ACK: Enable or disable TCP duplicate ACK feature.
- TCP forward ACK: Enable or disable TCP forward ACK feature.
- TCP congestion control: The congestion control algorithm in use by the TCP protocol stack.
- TCP window scaling: Allows the TCP protocol to automatically adjust to larger or smaller data being sent.
- TCP timestamps: Enable or disable the TCP timestamps feature.

5. Socket settings.


- Socket Recv Mem default: The default system memory available to a socket for receiving data, for any transport protocol.
- Socket Recv Mem max: The maximum system memory available to a socket for receiving data, for any transport protocol.
- Socket send Mem default: The default system memory available to a socket, for transmitting data, for any transport protocol.
- Socket send Mem max: The maximum system memory available to a socket, for transmitting data, for any transport protocol.
- Socket max listen backlog (SOMAXCONN): Maximum number of servers waiting to open a listening socket at one time.



Warning: Do not change these settings unless you are sure of what you are doing. These settings affect multiple services, network connectivity and overall system stability, if set to wrong values.

- **Routing table**

The system's routing table decides what happens to incoming and outgoing data on the networking devices and stack, at the IP layer. This is automatically generated when you configure the IP address settings in the "Network" page, for different network devices.



The screenshot shows a web interface for managing a routing table. At the top left, there is a red minus sign icon followed by the text "Routing table:". Below this is a table with the following columns: "Network/Host", "Subnet Mask", "Gateway", "Network Device", "Source IP address", and "Actions".






Network/Host	Subnet Mask	Gateway	Network Device	Source IP address	Actions
192.168.2.0	255.255.255.0		eth0	192.168.2.2	 Edit  Delete
default		192.168.2.1	eth0		 Edit  Delete
<input type="text"/>	<input type="text"/>	<input type="text"/>	---Select--- ▾	---Select--- ▾	 Add Route

Figure 6.5(d) Routing table of the system, with add, edit, and delete options.

1. Routing table editing:

- Network / Host: This is the destination network / host for the route.
- Subnet Mask: The mask decides what part of the address is for the Network, and what part is for the host.

- Gateway (optional): This specifies the next-hop router that will provide passage for traffic to this network or host.
- Network Device (optional): The device on the local system used for sending data using this route.
- Source IP address (optional): The local IP address from which data can be forwarded using this route.

- **ARP table**

The ARP table is a cache of ARP (Address Resolution Protocol) entries that the system has generated, based on the incoming and outgoing traffic. The system maintains the cache, and it usually requires no user intervention. You may need to change this only in extraneous circumstances, where a peer or the network is not functioning correctly with ARP.






ARP table				
IP address	MAC Address	Network Device	State	Actions
192.168.2.3	00:18:de:53:71:d3	eth0	DELAY	 Edit  Delete
192.168.2.1	00:0fb5:e9:4d:40	eth0	STALE	 Edit  Delete
<input type="text"/>	<input type="text"/>	---Select---		 Add entry

Figure 6.5(e) ARP cache table of the system, with add, edit, and delete options.

1. ARP table editing

- IP address: The IP address of the ARP entry.
- MAC address: This is the physical MAC address that the system will resolve the IP address to, and use for Ethernet traffic.
- Network device: The device used for Ethernet traffic to this MAC address.

6.6 DHCP Server

- **Sections of the interface**

1. Configuration Summary

This section allows you to configure DHCP server and iSCSI boot targets (targets which can be discovered and added by USS appliances via DHCP during remote boot over iSCSI).

To configure a DHCP server, provide Cluster IP to be used (HA mode), Subnet, Netmask and scope (range of valid IP addresses) and click Apply. DHCP service will have to be started manually in the “Service Summary” section using the “Enable” button.

To add an iSCSI boot target, provide valid inputs for the following parameters and click Apply:

- **Server name (Required):** IP of DHCP server
- **Port (Optional):** TCP port that the Target should provide iSCSI service on. If not provided, default iSCSI TCP port (3260) will be selected.
- **Lun Number (Optional):** LUN id of the physical SCSI device. If not provided, disk device with LUN id 0 will be selected.
- **Target name (Required):** IQN (iSCSI Qualified Name) address of the target to be added.

Note: Adding iSCSI Boot Target via DHCP is currently supported only on Chelsio's T3 adapters on the client side. For more information on iSCSI boot, please refer <http://www.ietf.org/rfc/rfc4173.txt>

Configuration Summary

Subnet:	<input type="text" value="102.50.50.0"/>
Netmask:	<input type="text" value="255.255.255.0"/>
Start address:	<input type="text" value="102.50.50.1"/>
End address:	<input type="text" value="102.50.50.255"/>
iSCSI Boot Targets :	
Targets 1:	
Server name:	<input type="text" value="102.50.50.233"/> (Required)
Port:	<input type="text"/> (Optional)
Lun Number:	<input type="text"/> (Optional)
(Lun number is starting from zero.)	
Target name:	<input type="text" value="iqn.2012-10.V1:1"/> (Required)
<input type="button" value="✓ Apply"/>	

Figure 6.6(a) *Configuring DHCP server with iSCSI boot target (non-HA mode)*

Configuration Summary

Cluster IP:

Subnet:

Netmask:

Start address:

End address:

iSCSI Boot Targets :

Targets 1:

Server name: (Required)

Port: (Optional)

Lun Number: (Optional)

(Lun number is starting from zero.)

Target name: (Required)

Figure 6.6(b)-Configuring DHCP server with iSCSI boot target (HA mode)

2. Service Summary

The Services Summary section displays the current status of the DHCP service. You can enable, disable or restart the service by using the actions provided.




Figure 6.6(c) DHCP service status with control commands

7. Storage

7.1 Storage overview

The storage subsystem is the single most important factor in the performance and reliability of the appliance. Configuring the storage correctly is critical to ensure that data is always available, irrespective of hardware failures, disruptions and upgrades. There are multiple storage hardware options supported by Unified Storage Server. The scalability and performance of the storage depends on the type of storage controller and hard disk drives (HDDs) used. Following is an overview of the different types of storage hardware that can be used in the appliance.

7.1.1 Integrated storage controllers

These are generally of two types, IDE / ATA, and SATA / SAS.

- a. **IDE/ ATA:** The onboard IDE controller generally has two ports, which allow attaching two HDDs or CD/DVD optical drives per port, totaling up to four drives in the system. IDE / ATA is a lower cost, low performance storage option, that is generally suited for desktop

workloads, and mainly used for booting the operating system. It is not meant for use as the primary storage of the Unified Storage Server appliance.

- b. SATA:** There may be an additional storage controller on the system motherboard, which is usually SATA on lower end systems or SAS on higher end hardware. SATA is an evolution of IDE / ATA, meant for larger capacity hard disks, and desktop / workstation workloads. SATA hard disks generally have a rotational speed of 7200 RPM or 10000 RPM, similar to IDE. This affects certain performance characteristics and responsiveness / time taken for storage operations to complete. SATA hard disks connected to the integrated SATA controller can be used as the primary storage for the appliance, but is recommended only for light to medium workloads or for very large capacity backup applications.
- c. SAS:** Serial-attached-SCSI or SAS is an enterprise-level storage interconnect, which allows for scalability and high performance. SAS hard disks generally have a rotational speed of 10,000 or 15,000 RPM, which allows for good performance with higher workloads. Serial-attached-SCSI allows for dual-ported hard disks, which offers multiple paths between the storage controller and the hard disk drives, allowing for data to be available even with path / hardware failures. If the system has a backplane and disk enclosure, this allows for hot-swap of SAS drives. On certain backplanes, there is support for cascading to further disk enclosures, using SAS expanders, allowing for scaling of the storage to multiple terabytes,

depending on storage controller and expander capabilities. Serial-attached-SCSI is the recommended integrated storage controller to use, if not using a hardware RAID controller.

7.1.2 Add-on storage controllers

An add-on storage controller is generally a SAS / SATA RAID or Fibre Channel controller or Host Bus Adapter (HBA) that sits on a slot on the system motherboard. The slot type may be a PCI-X or PCI-Express.

- a. **RAID controllers:** They have one or more SAS / SATA ports, internally or externally, allowing for connecting multiple hard disks or disk enclosures to the controller. Redundant Array of Independent Disks (RAID) is a mechanism of combining multiple hard disks into a virtual hard disk / 'RAID array', which is accessible to the system. The policy / behavior of reading / writing data to the set of HDDs within a RAID array, is referred to as the RAID level. RAID levels are described in the section 7.2.1 under Storage configuration.
- b. **Fibre Channel controllers:** Fibre Channel is an enterprise-level storage interconnect, which allows for networking of storage controllers with storage arrays, disk enclosures, and other storage devices such as tape libraries. The Fibre Channel protocols allow for many advanced configuration and management features, and complex topologies for deployment.

7.2 Storage management

Storage can be configured in multiple ways, depending on the type of data that will be stored, and the type of access / applications that will use the data. These criteria need to be analyzed before choosing a configuration method. The common configuration options and their usage is detailed below.

7.2.1 RAID arrays

Hardware or software RAID arrays can be configured using real drives, with a RAID policy. The common policies used are:

RAID level 0 – Striping.

This is a non redundant level, since a HDD failure will cause data loss. A minimum of two drives are required. The size of the virtual drive is the total of all the real drives. The data being written to the RAID array, is split into equal chunks, and striped across all the drives. This improves performance, since the storage controller can simultaneously read/write the data chunks from/to all drives at once. But the chances of loss of data are increased in this RAID level, since only one drive out of the entire set needs to fail, to lose the data from the entire set of drives. This level is recommended only for temporary / unimportant data.

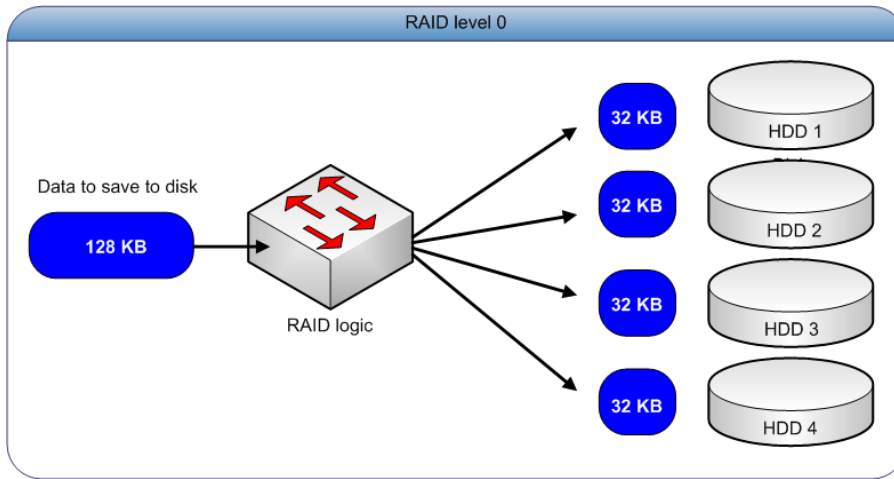


Figure 7.2.1 (a) – RAID level 0.

*Requires a minimum of 2 disks.
Maximum number of disks is
dependant on the RAID controller /
logic.*

RAID level 1 – Mirroring.

This is a redundant level, with one HDD failure out of two drives is allowed. Only two drives are allowed in this type of array. The size of the virtual drive is the size of one real drive within the array. The data being written, is written to both the drives. This ensures that data will be available if one of the two drives fail, but the write performance is equal to performance on a single drive. Read performance may be good on some controllers, which balance the reads across both drives. This level can be used for critical data.

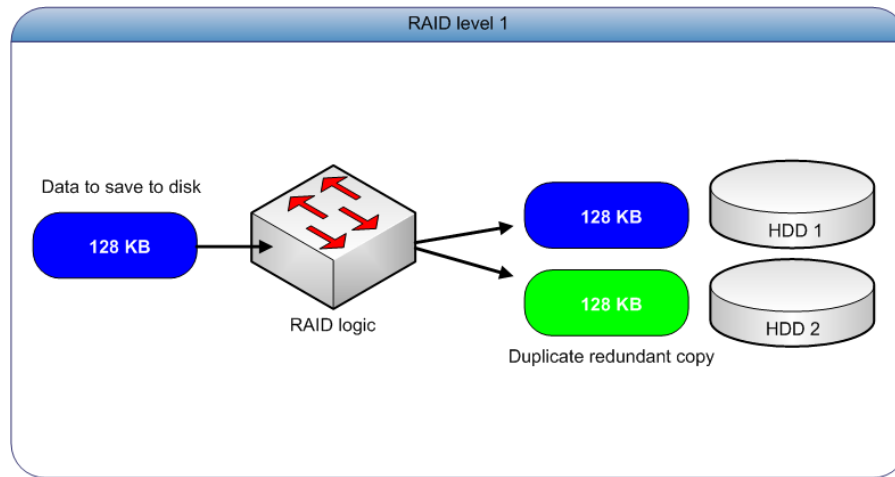


Figure 7.2.1 (b) – RAID level 1.

Requires a minimum and allows a maximum of 2 disks.

RAID level 5 – Striping with parity.

This is a redundant level, with one HDD failure out of all drives allowed. The size of the virtual drive is $(N - 1) \times$ size of real drives, where N is the number of real drives. One drive's space is used by the parity data. The data to be stored is split into equal chunks, but before writing to the real drives, the controller calculates parity data, usually an exclusive OR (XOR) calculation of the actual data, and stores that along with the data chunks. The parity is spread across all the drives along with the data. This is a commonly used RAID level, ideal for most applications requiring performance and redundancy.

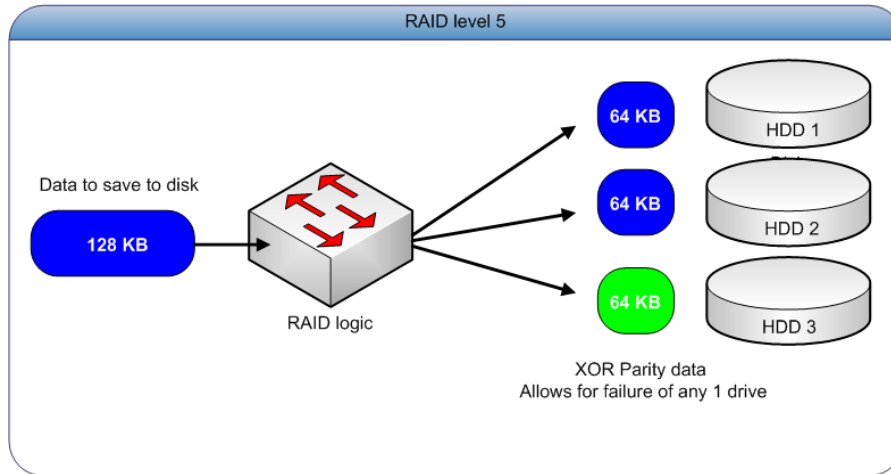


Figure 7.2 .1(c) – RAID level 5.

Requires a minimum of 3 disks. A very commonly used RAID level.

RAID level 6 – Striping with dual parity.

This is a redundant level, with two HDD failures out of all drives allowed. A minimum of four drives are required for this type of array. The size of the virtual drive is $(N - 2) \times$ size of real drives, where N is the number of real drives. Two drive's space is used by the dual parity data sets. The logic is similar to RAID 5, except that two sets of parity data are calculated for the actual data, and stored along with the data chunks. This is ideal for important data, with performance requirements.

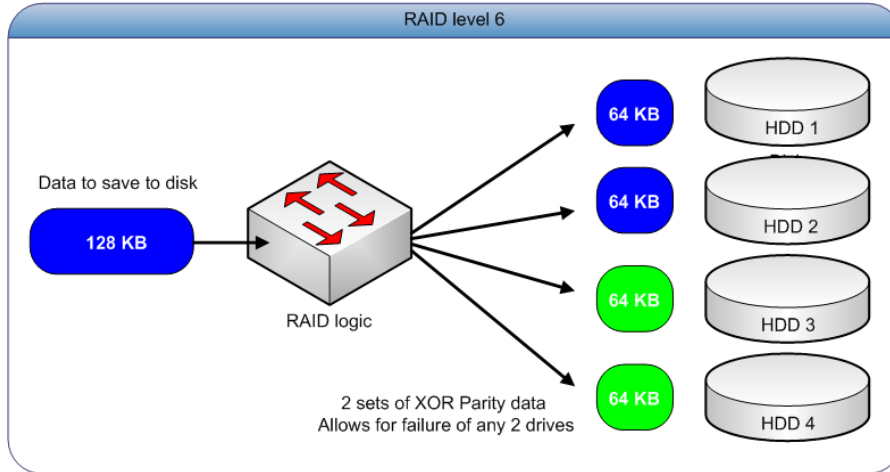


Figure 7.2 .1(d) – RAID level 6.
*Requires a minimum of 4 disks.
Recommended for critical data.*

RAID level 10 – Striping over mirrored / RAID 1 virtual drives.

This is a redundant level, with a HDD failure in each underlying mirror allowed. A minimum of four drives are required for this type of array. The size of the array is half of the total size of all real drives. This combines two RAID levels, with one level stacked above the other. The striping improves performance, and the mirroring improves redundancy. This is probably less preferable to level 6, since half the total drives space is lost in mirroring, and is recommended only for critical data.

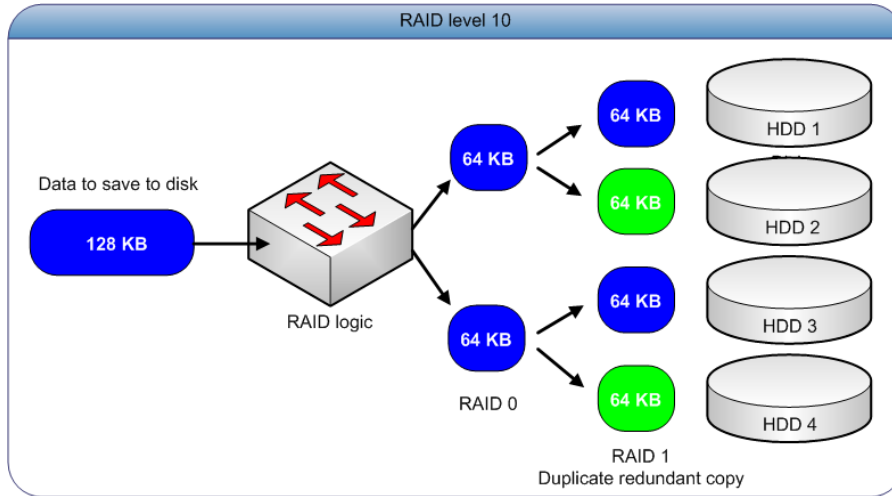


Figure 7.2.1 (e) – RAID level 10.
*Requires a minimum of 4 disks.
Recommended for critical data.*

RAID level 50/60 – Striping over RAID 5/6 virtual drives.

This is a redundant level, with one or two HDD failures in each underlying RAID 5/6 allowed, depending on the underlying level used. A minimum of six drives for level 50 is required, and eight drives for level 60. This is a complex RAID level, and is recommended only if configuring a single RAID Array with a large number of real drives, usually eight or more drives.

Currently, certain popular Hardware RAID controllers can be managed from the user interface.

Software RAID is an option for configuring virtual drives / RAID arrays from real drives accessible to Unified Storage Server. This is useful for systems with only integrated SAS / SATA controllers, with all drives attached to it, and no hardware RAID controller installed. It allows for protection of data, even if a drive fails, depending on the RAID level configured. The RAID algorithms are implemented in software, and utilize the system CPU and memory. It is usually not as efficient as a hardware RAID controller. Software RAID can be configured in the web-based user interface.

7.2.2 Chelsio Volume Management

1. Thin provisioned volumes
 - Volumes only occupy the actual space that was written to. Allocation/usage of disk space is only on demand.
 - Volumes can be dynamically extended.
2. Over provisioning
 - Volumes and Pools can be larger than the physical devices' capacity.
 - Physical devices can be added to the pool as required, when nearing the current devices capacity limit.
 - Alerts are provided for configurable thresholds.
3. Encryption of data to disk
 - AES FIPS-140 compatible encryption is used.
 - Master / Recovery key and per volume keys are supported.
 - Volumes and Pools can be migrated to another system, with the encryption key.

4. Instant snapshot, and instant restore from snapshot

- Snapshots use redirect-on-write, avoiding additional I/O or disk thrashing in copy-on-write implementations.
- Snapshots don't have a separate copy-on-write area, and are always guaranteed to be intact. They will never go invalid due to space constraints.
- The original volume can instantly be restored from any of its snapshots.

5. Cloning of volumes

- This allows creating a new volume with the current data existing on another volume, instantly.

Thin provisioning (TP) is a storage virtualization method to efficiently utilize the storage space. It is the allocation of data blocks as data is written in real-time, hence eliminating almost all whitespace which helps avoid the poor utilization rates that occur in the traditional storage allocation method. Organization can purchase less storage capacity upfront and defer storage capacity upgrade in line with actual business usage and save the operating cost associated with keeping unused disk capacity spinning at lower administrator efforts.

Thin provisioning enables over-allocation or over-subscription. Over-allocation is a mechanism that allows server application to be allocated more storage capacity than has been physically reserved on the

storage array itself. This allows flexibility in growth and shrinkage of application storage volume, without having to predict how much a volume will grow or shrink. Physical space on array is dedicated only when data is actually written by the application and not when the storage volume is initially allocated. Alerts can be set, so that Administrators can respond accordingly when pre defined thresholds are reached.

Features

- Performance: Efficient metadata management. Fewer disk read/write overhead.
- Data consistency and integrity.
- Avoids disk copy of old data to/from snapshot volume to original volume.
- Allows multiple read-write clones of a volume.
- Device specific reference count map: No pre-reserved area for reference map in the beginning of the pool. Size of the storage pool that can be created is 2^{64} sectors.
- Zoning: Address space of each disk in the pool is divided into multiple configurable allocation zones to store the data. The device and zone number for the volume can be user-specified or auto-generated. Data is striped across all the drives (if there are multiple drives in the pool) if only zone number is specified.

- Bulk Allocation: Allocate more than one chunk at a time instead of allocating one chunk every time if there is a new write. Bulk Allocation helps in optimizing metadata IO in HA mode and also Read/Write optimization.

7.2.3 Volume Management

This is the preferred method of configuring the real or virtual drives (RAID arrays) that Unified Storage Server can access. Volume Management allows for grouping of drives into Pools, from which space can be allocated dynamically for different purposes, such as a shared folder or an iSCSI disk. Volume Management is structured in the following manner:

Physical Volumes –real or virtual drives (RAID arrays).

This signifies the actual storage device that Volume Management is using to store data in. The total capacity available to allocate will depend on the number of physical disks in the pool.

Storage Pool – grouping of physical volumes / storage devices.

The storage pool is a grouping of the actual devices, and acts as a container or bucket, from which space can be allocated. It allows for management of allocated space, and the actual devices in a convenient manner.

Logical Volumes – allocated space from the storage pool.

When a certain amount of space is required for a particular use, it can be allocated from the storage pool, and this becomes a logical volume device, that can be formatted and attached to a folder path, or used as an iSCSI LUN (refer the iSCSI section of the user guide for details).

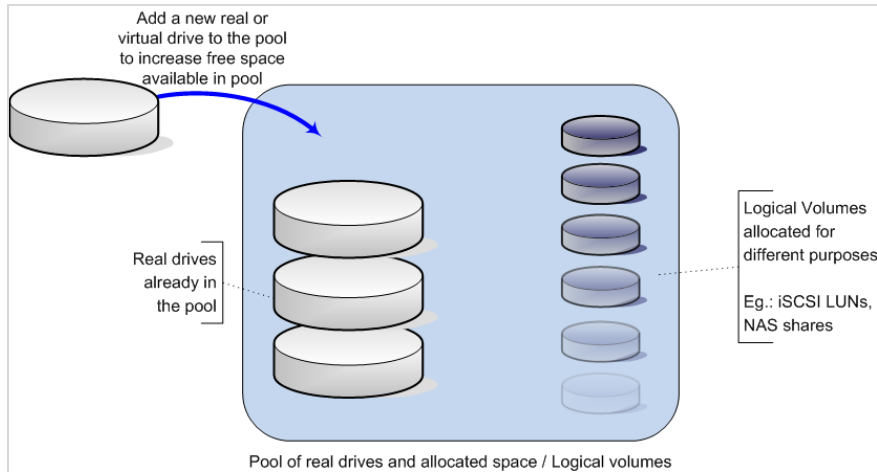


Figure 7.2.3(a) – Volume Management.

Snapshot Volumes – point-in-time view of a logical volume.

A snapshot allows for backing up the data of the logical volume at leisure, without interrupting the current I/O operations on the actual logical volume. The snapshot stores the data that is being overwritten or updated, after the point of time, when the snapshot was instantiated. If a large amount of data on the volume is being updated, the snapshot requires sufficient space to store all the prior data being updated. Lack of space to maintain the prior data, invalidates a snapshot. If the snapshot size is equal to the size of the Logical Volume, then the snapshot will be always valid, till the original volume exists.

On accessing the snapshot, either by attaching to a folder, if the volume has a valid file system, or by sharing it over iSCSI, the entire volume is available, but with the data at the point when the snapshot was taken. Any changes to the data after that point in time, do not reflect in the snapshot. Refer figure 7.2.3(b) for snapshot behavior.

Redirect-on-Write (ROW) is a method of protecting data that needs to be overwritten by new writes after a snapshot has been taken. It preserves the old data in its old location, and instead, redirects the new write to a new location. All subsequent reads and writes of data for the volume are performed at the new location. Snapshot reads continue to be performed from the old location. Redirect-On-Write snapshots feature is more performance optimized than COW snapshots, due to the lower number of I/O operations. Refer figure 7.2.3(c)

All Volume Management devices can be configured using the web-based user interface.

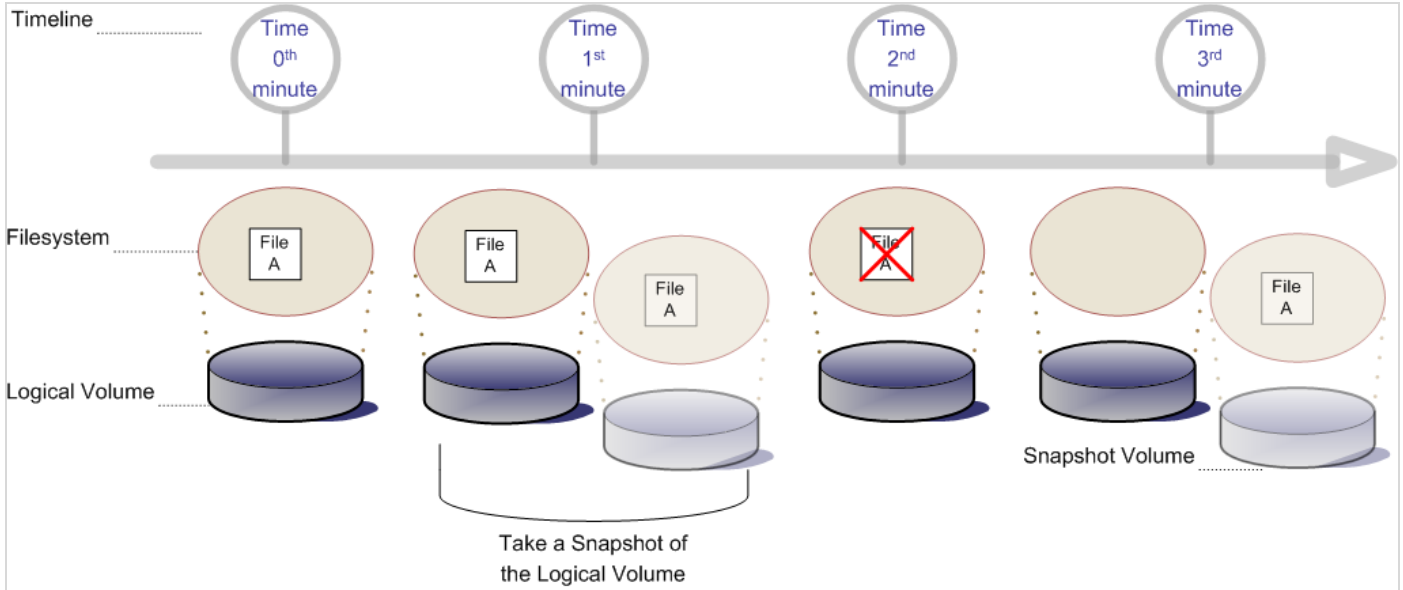


Figure 7.2.3(b) – Snapshot of a Logical Volume.




Figure 7.2.3 (c) – Chelsio Thin Provisioned volumes and pool

7.2.4 Partitioning

This is a legacy method of configuring the real or virtual drives (RAID arrays). Partitioning does not have the benefits of easy online resizing, etc., and is less flexible. It is not recommended to use partitioning to manage the drives attached to the appliance.

There are two types of partition tables supported by Unified Storage Server, i.e. MSDOS and EFI-GPT / GUID partition tables.

▶ MSDOS partition table:

Allows for four primary partitions, and many logical partitions, up to a total maximum of 16 partitions per drive.

Logical partitions need to reside in an extended partition container. The extended partition is part of the four primary partitions count.

▶ EFI-GPT / GUID partition table:

Allows for up to 16 partitions, all primary. It also supports very large drives, with drive size above two terabytes to be partitioned.

7.3 Storage configuration

The Storage configuration section has a few main sections. All currently-detected hard drives or virtual drives such as hardware RAID arrays, are shown in the OS physical disks page. Here you can assign a disk to be managed with Volume Management, or partition it.

This File Systems page lists the various devices and the corresponding folders on which they are mounted (attached).

The Software RAID page displays any software RAID arrays configured on the system, and allows for creating arrays using the free physical disks on the system. Free disks include those that are not partitioned and formatted for direct use, and those that are not assigned to Volume Management.

If a supported hardware RAID controller is detected, its configuration pages are automatically displayed.

Creating, modifying, and deleting RAID arrays on supported hardware RAID controllers are available.

The Volume Management section has a Manage Volumes page, which lists all the storage pools configured, and the free physical devices at the top of the page. Free physical devices can be assigned to an existing pool, or a new pool. From the space available in a pool, a logical volume can be created to allow

for iSCSI or file sharing services to use the storage. This option is available by clicking on 'Edit' at the title of each storage pool.

Logical volumes can be formatted with a file system and attached to a folder, for configuring with file sharing. It can also be directly used by iSCSI as a LUN.

A one-time snapshot can also be taken of the logical volume.

The snapshot scheduling option allows for scheduling snapshots of logical volumes, so that snapshots are taken at regular intervals, from which an administrator can restore the original data.

Replication allows for duplicating the data from a logical volume on the local system, to a peer system, which has sufficient disk space, and the peer volume can be created there. Replication configuration is automated across the two systems involved. Replication requires both systems to be running the same version of Unified Storage Server. It also requires configuring any firewalls between the two Unified Storage Server systems to allow TCP connections to be established between the two systems. The TCP port range being used by the replication connections is displayed in the UI in the summary section.



7.4 OS disk devices configuration

- Sections of the interface

1. List of disks on the system as seen by the OS

This section lists physical disks attached to the system. These may include hardware RAID arrays configured on RAID controllers and LUNs discovered through iSCSI/FC initiators. These disks can be selected and configured for different purposes, such as physical volumes in Logical Volume management or Chelsio Thin Provisioning(TP), partitioning. Options to rescan and remove missing devices are also provided.

Note: Size of disks and partitions here is calculated using the convention 1KB=1024 bytes, 1MB =1024 KB etc.

Select a disk from the list below to view details:

Disk list [Count: 6]

- ✓ 1. HDD: fioa, 74.51 GB
- ✓ 2. DVD/CD: hda
- ✓ 3. HDD: sda [Unified Storage Server Boot Drive], 149.05 GB [Has 4 Partitions]
- ✓ 4. HDD: sdb, ATA ST31000528AS, 931.51 GB
- ✓ 5. HDD: sdc, ATA ST31000528AS, 931.51 GB
- ✓ 6. HDD: sdd, ATA ST33000651AS, 2048.00 GB [Has 1 Partition]





Figure 7.4(a) - List of physical disks attached to the system

2. Devices Details

Select a disk from the Disk list and click on Device Details to view and configure various properties. Device properties like Read and Write cache can be enabled / disabled here.



The screenshot shows the 'Device Details' configuration page for a disk. At the top, it displays the disk's name 'MPIO-enabled: sdc', manufacturer 'FUJITSU MBA3073RC', and capacity '68.37 GB'. A 'SMART' status indicator with a warning icon is visible. Below this, the 'Device properties' section is titled. It shows the location as 'Enclosure 0_0: DELL MD1200 Bay: N/A' and the serial number as 'BLL3P960344Y'. The 'Disk cache' section contains four items: 'Read cache' set to 'Enabled', 'Write cache' set to 'Disabled', 'Parameters stored on device' set to 'Yes', and 'Cache segments count' set to '8'. An 'Apply' button with a checkmark icon is located below these settings. The 'Device layout' section at the bottom shows a horizontal bar with two segments: a small blue segment labeled '1' and a larger light blue segment labeled '2', followed by a large green segment.

Figure 7.4(b) - Device Details

3. Per Disk partition layout

The current partitioning layout is displayed for each disk. The layout is interactive. If you click on a partition, the actions for that partition are shown below. Sections in blue are currently configured, whereas sections in green are free / unconfigured.

sda | Unified Storage Server Boot Drive | Capacity: 232.89 GB | SMART

Device properties:

Serial #: GEK232RS3YJA2R

Disk cache:

Read cache:



Write cache:

Parameters stored on device: No

Device layout:

1 5 6

Figure 7.4(c) - OS boot disk

 **MPIO-enabled: sdc | FUJITSU MBA3073RC | Capacity: 68.37 GB |**  **SMART**

Device properties:

Location: Enclosure 0_0: DELL MD1200 Bay: N/A
Serial #: BLL3P960344Y

Disk cache:

Read cache:

Write cache:

Parameters stored on device: **Yes**

Cache segments count: 8

Device layout:






Figure 7.4(d) - Partitioned disk

 **MPIO-enabled: sdd | FUJITSU MBA3073RC | Capacity: 68.37 GB |**  **SMART**

Device properties:

Location: Enclosure 0_0: DELL MD1200 Bay: N/A
Serial #: B JL3P9C09BAY


Disk cache:

Read cache:

Write cache:

Parameters stored on device: **Yes**

Cache segments count: 8

 **Apply**

Device layout:

Physical volume in Volume management
Pool pool1

Figure 7.4(e) - Disk used in volume management

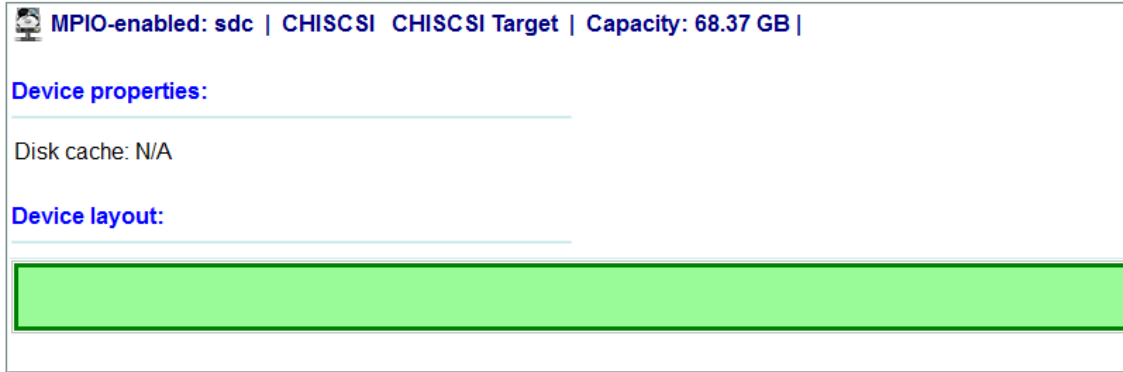


Figure 7.4(f) MPIO iSCSI LUN

Example: Creating partitions

Creating partitions is supported only on disks with existing partitions. You can create partitions on the free/unconfigured sections of a physical device using the following steps:

- i. Select the disk on which partition is to be created and click **Device details**.
- ii. In the Device layout section, click on the green section.

- iii. In the Actions drop down, select *Create partition*
- iv. Enter the partition size in MB, GB or TB
- v. Select the partition type and click **Apply**.

- **Disk free space actions**

- 1. Remove from Volume management**

If a disk is a part of Volume Management but not used in any storage pool, you can re use the disk by using this option.

- 2. Manage with Volume management**

Note: Volume Management Type is LVM by default if the USS appliance is not licensed. Chelsio TP is enabled only after licensing the appliance.

- 2.1. Logical Volume Management (legacy, not recommended)**

When not licensed with TP, user can manage the space on the selected disk with Logical Volume Management.

Device layout:

Free space:
931.51GB

Actions:

Volume Management usage:

New Pool name:

Volume Management Type:

Figure 7.4(g) - Creating a new LVM Pool

There are two ways of configuring the disk for Logical Volume Management:

- i) Create a new pool: Storage Pools are containers for many disks to reside in, and space can be allocated for different purposes from the Pool. If there are no storage pools currently configured on the system, user needs to specify the new pool name and select the Volume Management Type as LVM. Now the disk will be assigned to the newly created Pool.
- ii) Assign to existing pool: Using this option, user can add the selected disk to any of the existing LVM pools.

2.2. Chelsio Volume Management (Thin Provisioned, recommended)

Using this option, User can manage the space on the selected disk with Chelsio Thin Provisioning. This is the recommended mode of managing disk space on the system.

Free space: 68.37GB

Actions:

Volume Management usage:

New Pool name:

Volume Management Type:

Pool Size: (>= 32 GB)

Chunk size: 128K (recommended) 64K 32K

Disk for Cache:

SSD disks:

SSD Cache Mode:

RAM Cache:

Allocation Size: chunks

Allocation Zones per disk:

Max volumes in a pool: (less than or equal to 32767)

Note: "Maximum volumes in a pool" cannot be changed after creating the pool.

Figure 7.4(h) - Creating a new TP Pool with SSD Cache

Free space: 68.37GB

Actions:

Volume Management usage:

New Pool name:

Volume Management Type:

Pool Size: (>= 32 GB)

Chunk size: 128K (recommended) 64K 32K

Disk for Cache:

RAM Cache :

(RAM Cache size: Minimum: Atleast chunk size selected for pool, Maximum of 2048 MB)

Allocation Size: chunks

Allocation Zones per disk:

Max volumes in a pool: (less than or equal to 32767)

Note: "Maximum volumes in a pool" cannot be changed after creating the pool.

Figure 7.4(i) - Creating a new TP Pool with RAM Cache

There are two ways of configuring the disk for Chelsio TP.

i) Create a new pool: If you do not have any storage Pool currently on the system, you need to select the Volume Management Type as Chelsio TP and specify the new Pool name, along with few other settings mentioned below, depending on which the pool will be created, and the disk will be assigned to the newly created TP Pool.

- Pool Size: Allowed pool size, which should be greater than or equal to 32 GB. Once the pool is completely occupied by volumes, new pool needs to be created. Please note that the aggregate size of all the pools configured on the appliance cannot exceed the storage capacity with which the appliance is licensed.
- Chunk Size: The chunk size has to be in power of 2. The allowed chunk sizes are 32, 64 and 128K (default value).
- Disk for Cache: Using this option you can enable/disable SSD Cache. Allowed options are Internal or External SSD Drives.
- SSD Disks: List of SSD disks which can be used as cache disk.
- SSD Cache Mode: Write back mode is available for SSD cache. In this mode, the application data is written to SSD with immediate I/O completion. The data in SSD is flushed to disk at some later time. SSD in this mode has higher write and read performance than disk device.

- RAM Cache: In RAM cache, application writes are cached onto physical memory with immediate I/O completion with no processing overheads of storage stack and latency of storage device exposed to user. RAM cache is the preferred mode for faster write.
- Allocation size: Number of chunks allocated for every new write.
Note: If SSD cache is enabled, the allocation size will be set to 1 and cannot be altered.
- Allocation Zones per disk: Total size of each disk in the pool will be divided into multiple allocation zones as specified here by the user.
- Max volumes in a pool: The Maximum number of volumes that can be configured in a pool. The limit is 32767.

Note:

- "Maximum volumes in a pool" cannot be changed after creating the pool.
- Pool initialization may take some time if the physical disk is slow. The status of the process can be seen in "storage volume management" page.
- You can create a pool with either SSD cache or RAM cache but not both.

ii) Assign to existing pool: Using this option, the user can add the selected disk to any of the existing Chelsio TP pools.

- **Per Partition details**

The details of a partition such as its size, file system type if any, or usage in volume management is displayed, when a partition is selected. Any actions for the partition are also displayed, based on its current usage.

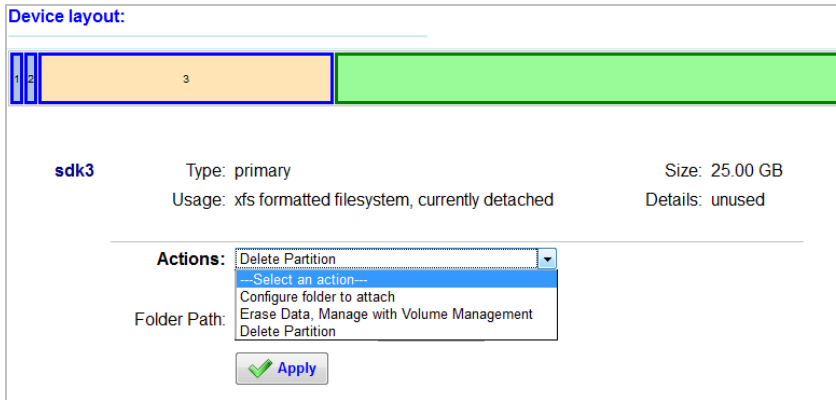


Figure 7.4(j) - Partition details and actions

1. Partition actions

- Erase data, manage with Volume management:
This option allows assigning the partition to be used in Volume management.
- Delete partition:
This option will permanently delete the partition.
- Configure folder to attach:
Using this option, user can mount a partition to a folder.

Note: This option is available only when the partition is formatted with a file system.

- Detach from folder:
Using this option, user can unmount a folder previously mounted on a partition.

Note: This option is available only when a partition (formatted with a file system) is mounted to a folder.



Warning: Please ensure that the partition does not contain any required data before selecting these options. These actions cannot be undone.

S.M.A.R.T (Self-Monitoring, Analysis and Reporting Technology) is a monitoring technology for storage devices that provides information about the status of a drive as well as the ability to run self tests. It can be used to detect and report on various indicators of reliability in the hope of anticipating failures.

- **Sections of the interface**

- 1. Summary**

The Summary section displays disk details like device type, serial number etc. SMART support can be enabled/ disabled here.

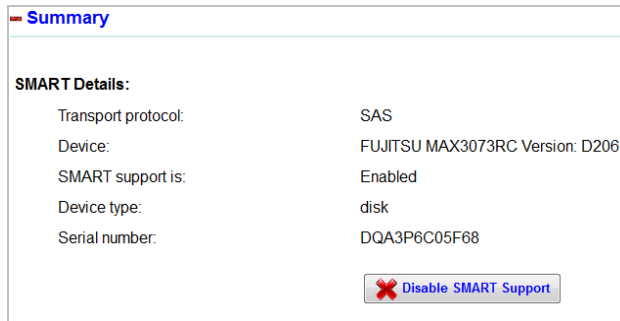


Figure 7.4(k) - Summary section with disk details

2. Smart Test

Using SMART Test, users can run a number of self-tests and also view logs of previously run tests.

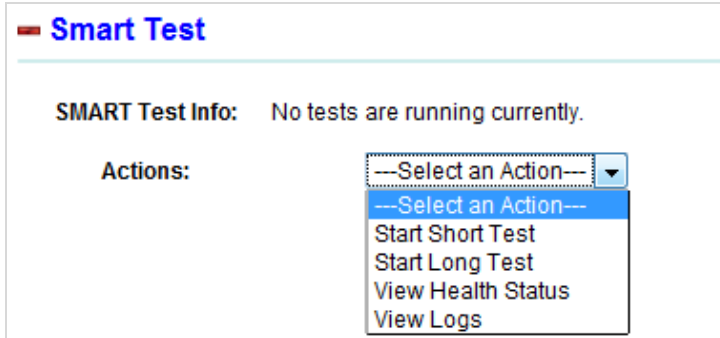


Figure 7.4(l) - SMART test actions

2.1. Smart test actions

- Start Short/Long Test: These are series of self-assessment tests performed to detect any impending drive failure. The exact tests vary by hardware manufacturer and can include Power-On Hours, Temperature, Seek Error Rate and many others. A Short Test usually runs for under ten minutes while Long Test may run for tens of minutes.

- **View Health Status:** Displays information on various hard disk related parameters.
- **View Logs:** Information about the most recent errors that the drive has reported and previously run tests are reported here.

Example: Running S.M.A.R.T test

- i. In the **Disk devices** section, click the disk for which you want to run the test in the **Disk list** and click **Device details**. This will navigate to a new page.
- ii. Now, click on the **SMART** button. This will navigate to the SMART test page.
- iii. Click on **Enable SMART support** in the **Summary** section, if it's disabled.
- iv. In the **Smart Test** section, select the type of test (Short,Long), view health status or logs using the **Actions** drop down.
- v. Click **Apply**.

7.5 Software RAID Array configuration

- Sections of the interface:

1. Free / unassigned disks

This section lists physical disks attached to the system that are unused / free. These disks can be selected and configured for creating a new RAID array.



Figure 7.5(a) - Free physical disks section

1.1. Free / Unassigned Physical devices actions:

1.1.1. Single Select:

- Assign as a dedicated Hotspare: Using this option a disk can be configured as a hot spare for a particular array. If any of the disks in that array fails, the hot spare disk replaces the failed physical disk.

Note: The above option is not applicable for RAID level 0. Also, this option will not be available if there are no arrays created.

1.1.2. Multiple Select

- Create a new RAID Array:

This action is available for multiple disks only. Hence you will need to select multiple disks in the list on the left by clicking on them. The settings are:

- RAID level: The type of RAID array to create. This decides many factors of the final array created. The RAID levels are 0, 1, 5, 6, 10, 50, 60. Please refer to the RAID levels explanation at the end of this page for further details.
- Stripe size: The RAID algorithm splits incoming data to smaller chunks and distributes those chunks to the disks in the array, if the RAID level has a striping

requirement. This setting specifies the size of chunks of data to write to each disk.

Example: How to create a Software RAID Array

- i. In the **Physical Disks** list, select multiple disks by clicking on them. The **Properties/Actions** section (on the right) will change to display related options/actions for creating a RAID array. Clicking on a single disk will display its properties.
- ii. In the **Actions** drop down, select *Create a new RAID Array*.
- iii. Select the RAID level (available levels will differ depending on the number of disks selected).
- iv. Select the Stripe size (if available).
- v. Click **Apply**.

If the array was created successfully, it will be displayed below the **Free / Unassigned Physical devices** section, with related properties and actions.

2. Arrays

The RAID arrays configured are displayed here. The disks used by this RAID array are listed on the left side, and the status and actions for the Array are listed on the right. If a disk is selected, including hot spare, any available actions for the disk are shown on the right.

4. md4: RAID 1 | Size: 2.00 GB | Edit

Physical Disks

- ✓ MPIO (sdg1) - 2.00 GB
- ✓ MPIO (sdh1) - 2.00 GB

Properties / Actions:

Array ID: **md4**

Status: **active**

Actions:

Apply

Figure 7.5(b) - Arrays listing

2.1. Array actions

- 2.1.1 Remove Hotspare/Standby: The selected disk will be removed from the list and no longer be assigned as hotspare for the particular array. It will reappear in the Free/Unassigned Physical disks section.

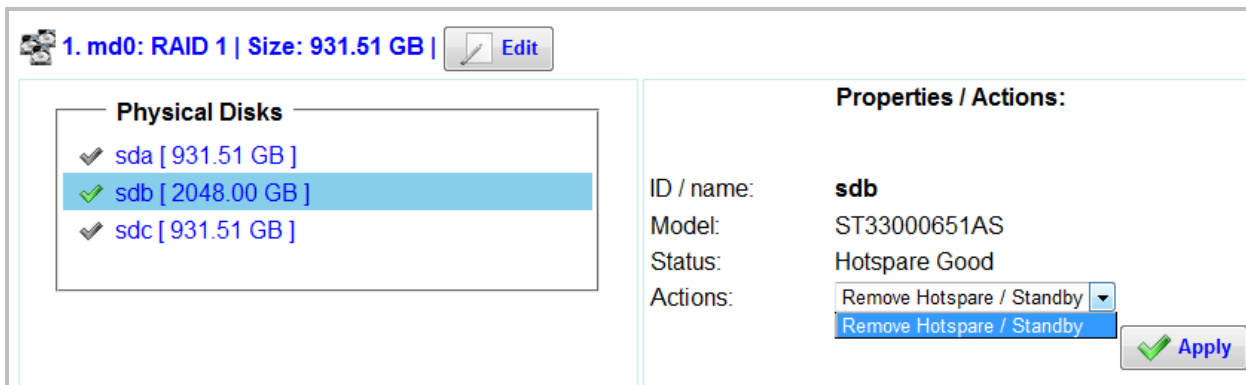


Figure 7.5 (c) - Hotspare disk with related Properties and Action

2.1.2 Manage with Volume management: This option allows assigning the array to be used in Volume management.

Note: Volume Management Type is LVM by default if the USS appliance is not licensed. Chelsio TP is enabled only after licensing the appliance.

1. Logical Volume Management (legacy, not recommended)

Using this option, user can manage the space on the selected disk with Logical Volume Management.

Actions:

Volume Management usage:

New Pool name:

Volume Management Type:


Figure 7.5(d) - Creating a new LVM Pool

There are two ways of configuring the disk for Logical Volume Management.

- a. Create a new pool: Storage Pools are containers for many disks to reside in, and space can be allocated for different purposes from the Pool. Specify the new Pool name and select the Volume Management Type as LVM. Now the disk will be assigned to the newly created Pool.
- b. Assign to existing pool: Using this option, user can add the selected disk to any of the existing LVM pools.

2. Chelsio Volume Management (Thin Provisioned, recommended)

Chelsio Thin Provisioning (TP), in a shared storage environment, is the allocation of data blocks as data is written real-time. This methodology eliminates almost all whitespace which helps avoid the poor utilization rates that occur in the traditional storage allocation method where large pools of storage capacity are allocated to individual servers but remain unused. Using this option, User can manage the space on the selected disk with Chelsio Thin Provisioning. This is the recommended mode of managing disk space on the system.

4. md4: RAID 1 | Size: 2.00 GB |  Edit

Physical Disks

- MPIO (sdg1) - 2.00 GB
- MPIO (sdh1) - 2.00 GB

Properties / Actions:

Array ID: **md4**

Status: **active**

Actions: Manage with Volume Management ▾

Volume Management: Assign to Pool: Create new pool ▾

New Pool name:

Volume Management Type: Chelsio TP (recommended) ▾

Pool Size: TB ▾ (>=32 GB)

Chunk size: 128K (recommended) 64K 32K

Disk for Cache: External Disks (SSD) ▾

SSD disks: fioa, size:74.51 GB ▾

SSD Cache Mode: Write back ▾

RAM Cache: Disable ▾

Allocation Size: chunks

Allocation Zones per disk:

Max volumes in a pool: (less than or equal to 32767)

Note: "Maximum volumes in a pool" cannot be changed after creating the pool.


 Apply

Figure 7.5(e) - Creating a new TP Pool

There are two ways of configuring the disk for Chelsio TP.

- a. Create a new Pool: Select the Volume Management Type as Chelsio TP and specify the new Pool name, along with the few other settings mentioned below, depending on which the pool will be created, and the disk will be assigned to the newly created TP Pool.
 - i. Pool Size: Allowed pool size, which should be greater than or equal to 32 GB. Once the pool is completely occupied by volumes, new pool needs to be created. Please note that the aggregate size of all the pools configured on the appliance cannot exceed the storage capacity with which the appliance is licensed.
 - ii. Chunk Size: The chunk size has to be in power of 2. The allowed chunk sizes are 32, 64 and 128K (default value).
 - iii. Disk for Cache: Using this option you can enable/disable SSD Cache. Allowed options are Internal or External SSD Drives.
 - iv. SSD Disks: List of SSD disks which can be used as cache disk.

- v. SSD Cache Mode: Write back mode is available for SSD cache. In this mode, the application data is written to SSD with immediate I/O completion. The data in SSD is flushed to disk at some later time. SSD in this mode has higher write and read performance than disk device.
- vi. RAM Cache: In RAM cache, application writes are cached onto physical memory with immediate I/O completion, no processing overheads of storage stack and latency of storage device exposed to user. RAM cache is the preferred mode for faster write.
- vii. Allocation size: Number of chunks allocated for every new write.
Note: If SSD cache is enabled, the allocation size will be set to 1 and cannot be altered.
- viii. Allocation Zones per disk: Total size of each disk in the pool will be divided into multiple allocation zones as specified here by the user.
- ix. Max volumes in a pool: The Maximum number of volumes that can be configured in a pool. The limit is 32767.

Note:

a) "Maximum volumes in a pool" cannot be changed after creating the pool.

b) Pool initialization may take some time if the physical disk is slow. The status of the process can be seen in "storage volume management" page.

c) You can create a pool with either SSD cache or RAM cache but not both.

b. Assign to existing pool: Using this option, the user can add the selected disk to any of the existing Chelsio TP pools.

2.1.3 Disable / deactivate array: This option will disable the array. After that, the array will not be accessible by the system. The array will be activated again automatically on the next reboot.

2.1.4 Delete array: You can permanently delete the array using this option.

7.6 iSCSI Initiator

- Sections of the interface:

1. Summary

The Summary section displays the Initiator IQN and service details. You can change the IQN name and start, stop or restart the service. The IQN name should be in the following format: `iqn.<date>.<Naming Auth>:<optional string>`.



Figure 7.6 (a) - Initiator summary and control actions.

1.1. Initiator control actions

- **Enable/Disable:** You can choose to start/stop iSCSI Initiator service using this button. Enabling the service also configures it to start automatically on system bootup. The default is to start the service automatically.
- **Restart:** This command will stop and start the iSCSI Initiator service.

2. Global Settings

You can set Global settings like CHAP username and password, Header digest, Checksum for iSCSI initiator here. You can also restore the settings to their default values using the Restore button.

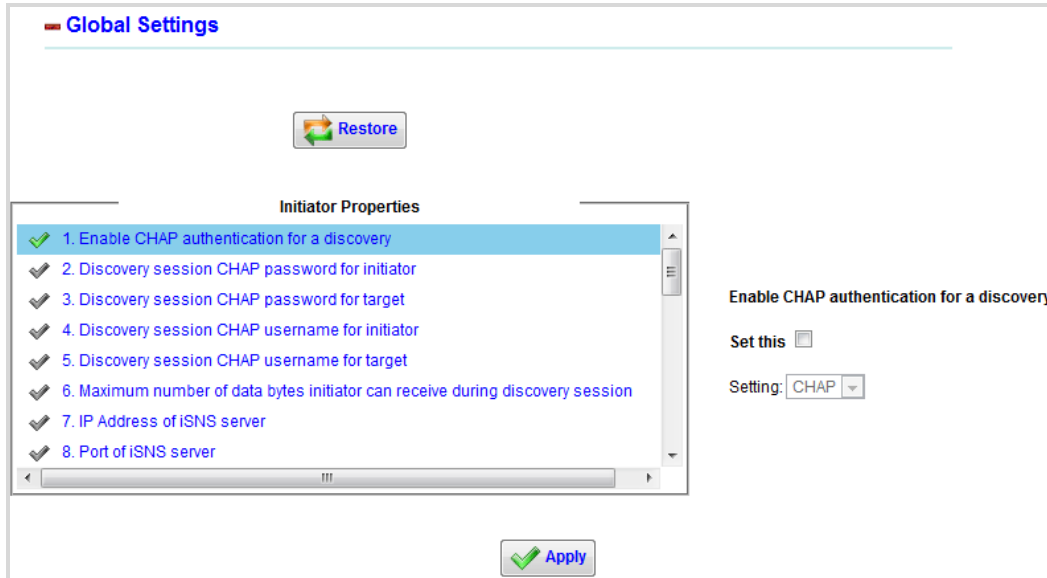



Figure 7.6 (b) - Initiator Global Settings.

3. Chelsio Adapter configuration

This section lists Chelsio adapter details like IP address, status, driver and Mac address.



Chelsio Adapter configuration

iSCSI network ports

- ✓ 1. eth0
- ✓ 2. eth7
- ✓ 3. eth3
- ✓ 4. eth6
- ✓ 5. eth5
- ✓ 6. eth4

Interface: eth0

Status: Physical Link Up

Iface Configuration:

IP Address:	102.77.77.147
Driver:	cxgb3i
HW Address:	00:07:43:05:8e:d2

Figure 7.6 (c) - Chelsio Adapter configuration

7.6.1 Remote iSCSI Targets

- **Sections of the interface:**

1. Summary

The Summary section displays the number of targets discovered, saved, connected and the total number of LUNs configured on the target.



Summary	
Total targets discovered:	2
Targets enabled on startup:	2
Total targets Connected:	1
Total Luns from iSCSI targets:	2

Figure 7.6.1(a) - Remote iSCSI Target summary

2. Target Details

The Target Details section displays details of discovered and connected targets.

- Target Details

Targets List

✓	1.	🌐 Portal: 102.20.20.165:3260 Target: iqn.2013-04.dzongri:1 102.20.20.165
---	----	--

Enabled on startup: **Yes**
Status: **Connected**
Session ID: 1
Connection: **102.20.20.202 -> 102.20.20.165:3260**
lface Name: **cxgb4i.00:07:43:10:69:20**

Actions: Rescan this target ▾


 Apply

Figure 7.6.1 (b) - Remote iSCSI Target details

1.1. Target Details actions

- Rescan this target: Rescan will perform a SCSI layer scan of the session to find new LUNs.
- Logout from this target: This action will log out from the connected target and close the session.
- Delete: The Delete option is available only if the target is not connected. You can delete the selected target record from the Discovery table and Node table using this option.

3. LUN List

The LUN list displays all the discovered LUNs from the iSCSI target and their status.

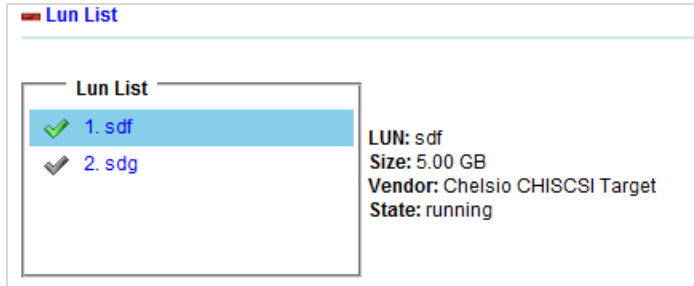


Figure 7.6.1(c) - LUN List

4. iSCSI Session Properties

You can set iSCSI session properties like CHAP username and password, Header digest, Checksum etc here. Please note that some parameters however cannot be changed.

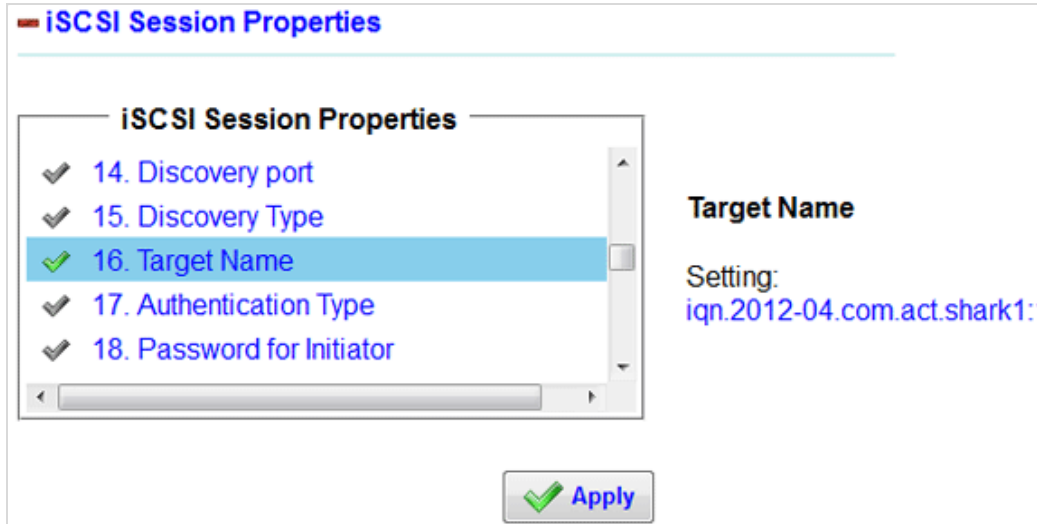



Figure 7.6 .1(d) - iSCSI Session Properties

5. Add a Target

You can add a new Target using two modes: Directly connect to target and Query iSNS Server for targets. In the first mode, you have to provide target IP address, port details and the interface you want to bind the target to. In second mode, mention the IP address of iSNS server and the default interface, and it will discover all the connected targets. In the first mode, only Chelsio interfaces will be listed in the iface field, whereas for the second, only the default interface will be listed. In HA mode, network interfaces of both primary and secondary nodes have to be provided for both methods.



The screenshot shows a window titled "Add Target" with a red minus icon on the left. Below the title bar, there is a radio button selected for "Directly connect to target". Underneath, there are three input fields: "IP Address:" with the value "102.77.77.147", "TCP Port:" with the value "3260", and "iface:" with a dropdown menu showing "eth2 [IP: 102.1.1.62]". At the bottom, there is a button with a magnifying glass icon and the text "Discover Target".

Figure 7.6.1 (e) - Directly connecting to a target (non-HA mode)

Add Target

Directly connect to target

IP Address:

TCP Port:

Iface:

Peer Iface:


 Discover Target

Figure 7.6.1 (f) - Directly connecting to a target (HA mode)

Query iSNS Server for targets

IP Address of iSNS Server:

Iface:

 Discover Target

Figure 7.6.1 (g) - Querying iSNS for targets (non-HA mode)

Query iSNS Server for targets

IP Address of iSNS Server:

Iface:

Peer Iface:

 Discover Target

Figure 7.6.1 (h) - Querying iSNS for targets (HA mode)

Example: How to add a remote iSCSI Target manually

Please ensure that the iSCSI target to be added is running before attempting to add.

- i. If the target IP address and TCP port number is already known, you can add the target directly. Select the **Directly connect to target** radio button.
- ii. Specify the IP address of the target.
- iii. Specify the TCP port. The port specified here must be same as provided on the target. If default was chosen, it will be 3260.
- iv. Choose a Chelsio interface to bind the target to.
- v. Click **Discover Target**
- vi. If the target was discovered successfully, it will appear in the **Target Details** section. Select the newly discovered target node and select *Login to this target*.
- vii. Click **Apply**.
- viii. Select the target node to view its properties.

Example: How to add a remote iSCSI Target using iSNS server

- i. Select the **Query iSNS Server for targets** radio button.
- ii. Specify the IP address of the iSNS server.
- iii. The default interface will be selected by default.

- iv. Click **Discover Target**
- v. All the discovered targets by the iSNS server will be listed in the **Target Details** section. Select a target node to connect and select *Login to this target*.
- vi. Click **Apply**.
- vii. Select the target node to view its properties.

Example: How to access iSCSI target with CHAP authentication enabled

Follow the instructions mentioned above to add an iSCSI target either manually or by using iSNS server. Once the target appears in the **Targets List** in the **Target Details** section, follow these steps:

- i. Select and highlight the target node in the **Targets List**.
- ii. Expand the **Target Properties** section.
- iii. Scroll down to the *Authentication Type* parameter. Select (one-way) *CHAP* or *Mutual CHAP* from the drop-down.
- iv. If (one-way) *CHAP* is selected, enter the username and password for the initiator. If *Mutual CHAP* is selected then enter the username and password for the target along with the initiator. If the Initiator is part of a cluster environment, initiator name of the peer node must also be provided in the *Peer initiator username* field for both one-way and mutual CHAP.
- v. Click **Apply**.

Note: If there are more than one initiator accessing the target, and the CHAP credentials used to login are same, then you can set them in the **Global Settings** section under **iSCSI Initiator** module.

7.7 Hardware RAID Array configuration

7.7.1 Hardware RAID controller Summary

- **Sections of the interface:**

1. **Adapter Summary**

The **Adapter details** sub-section displays RAID controller related details like Serial number, BIOS and Firmware version, different RAID levels the controller supports and Disk Drive types are displayed here. The **RAID Arrays and disk details** sub-section displays Physical Disks and RAID Arrays configuration details like Total RAID Arrays configured, number of Physical Drives present are displayed here.

Note: Adapter details displayed here may differ depending on the RAID controller present in the system.

Adapter details:	
Serial number:	SV12708272
Manufactured date:	07/01/11
Firmware package version:	23.7.0-0035
Boot Block version:	2.05.00.00-0007
BIOS version:	5.33.00_4.12.05.00_0x05160000
Firmware version:	3.190.25-1776
WebBIOS version:	6.1-45-e_45-Rel
Text-mode BIOS version:	05.04-05
Battery Backup Unit:	Present
Alarm:	Present, Enabled
NVRAM:	Present
Memory:	Present
Flash:	Present
RAID levels:	RAID0, RAID1, RAID5, RAID6, RAID00, RAID10, RAID50, RAID60, PRL 11, PRL 11 with spanning, SRL 3 supported, PRL11-RLQ0 DDF layout with no span, PRL11-RLQ0 DDF layout with span
Disk Drive types:	SAS, SATA

Figure 7.7.1 (a) - LSI MegaRAID adapter details

RAID Arrays and disks details:	
Total RAID Arrays:	3
Arrays in degraded state:	0
Arrays in failed state:	0
Physical Disk Drives:	8
Critical Disks:	0
Failed Disks:	0

Figure 7.7.1 (b) - RAID Arrays details

2. Settings

2.1. RAID Controller Actions

- **Enable Alarm:** If a Physical disk is detached from a RAID Array or if the disk fails, enabling alarm action will raise an alarm in the form of short beeps. User can disable or silence the alarm.
- **Upload and Update Firmware:** You can update the RAID controller to the latest firmware using this option.

Note: Configuration options available here may differ depending on the RAID controller present in the system.

Settings

Alarm actions:

Use this option to send the content of cache memory to the virtual disk(s) in adapter.

Maintain PD Fail History: (0-100%)

Rebuild Rate: (0-100%)

Background Initialization Rate: (0-100%)

Consistency Check Rate: (0-100%)

Reconstruction Rate: (0-100%)

Cache Flush Interval: (0-255 Seconds)

Upload and Update Firmware:


Figure 7.7.1 (c) - Settings available for LSI MegaRAID controller

7.7.2 Physical Devices and RAID Arrays

- **Sections of the interface:**

- 1. Free / unassigned disks**

This section lists physical disks attached to the controller, which are unused / free. These disks can be selected and configured for different purposes, such as creating a new RAID array, or as a hot spare.

 **Free / unassigned Physical disks**

Physical disks:

Physical Disks (multi-select)

- Enclosure ID 252 -> Device 4 [1.818 TB]
- Enclosure ID 252 -> Device 5 [1.818 TB]

* Select multiple disks in the list above to create a RAID array, or select a single disk to modify the disk.

Properties / Actions:

ID / Slot: **4**

Model: JK1174YAHL58TWHita

Serial: HDS722020ALA330

Foreign State: **None**

Status: **Free, Unconfigured(good), Spun Up**

Actions: --Select an action--

--Select an action--

Assign as Global Hotspare

Assign as Dedicated Hotspare

Clear Disk

Add to existing RAID Array

Locate / Blink Disk

Stop Blinking

Apply

Figure 7.7.2 (a) - Properties and actions displayed for a free physical disk

1.1. Free/unassigned Physical disks actions

1.1.1. Single Select


- **Assign as Global Hotspare:** Using this option a drive can be configured as a hot spare drive so that if any physical drive in any array fails, it automatically replaces the failed physical drive. Subsequently, the logical drive can be rebuilt automatically.
- **Assign as a dedicated Hotspare:** Using this option a drive can be configured as a hot spare for a particular array. If any of the drives in that array fails, the hot spare drive replaces the failed physical drive.
Note: The above two options are not applicable for RAID level 0.
- **Clear Disk:** Erases all the data from the disk. The disk will now appear in the **Bad/Busy/Foreign Physical Disks** list and clicking on it will display the progress of the ongoing Clear Disk process.
- **Add to existing RAID Array:** Adds the selected disk to one of the existing array. User can also change the Array RAID level here, depending on the numbers of drives present in the array.
- **Stop Clear Operation:** Stops/Aborts the ongoing Clear Disk operation.

- Locate/Blink Disk: This action will cause the corresponding Disk LED indicator to blink.
- Stop Blinking: This action causes the LED indicator to stop blinking.

1.1.2. Multiple select

The following options are available only on selecting/highlighting two or more disks in the free disks list. The user can choose to either create a new RAID Array or add the unassigned disks to an existing one.

- Create a new RAID Array: Using this option user can select multiple disks to create a new RAID Array. Depending on the number of disks selected, user can chose which RAID level the array needs to be configured as.

 **Free / unassigned Physical disks**

Physical disks:

Physical Disks (multi-select)

- Enclosure ID 252 -> Device 4 [1.818 TB]
- Enclosure ID 252 -> Device 5 [1.818 TB]
- Enclosure ID 252 -> Device 6 [1.818 TB]
- Enclosure ID 252 -> Device 7 [1.818 TB]

* Select multiple disks in the list above to create a RAID array, or select a single disk to modify the disk.

Properties / Actions:

Actions: Create a new RAID Array ▾

RAID level: RAID 1 ▾

Name: raid_1 (required)

Write-Back (enabled) ▾

Write caching: Note: depends on presence of Battery backup unit and caching policy on no battery.

Write caching without Battery Backup: Disabled (recommended) ▾

Read Ahead: Adaptive Read Ahead ▾

Caching: Enabled ▾

Stripe Size: 1024 KB ▾




Figure 7.7.2 (b) - creating a new RAID Array

The following parameters need to be configured while creating a new RAID Array.

- i. RAID level: Select the RAID level the array needs to be configured as.
- ii. Name: Choose a name for the RAID Array to be created.

- iii. **Write Caching:** Write Caching is a component that transparently stores data in fast and volatile memory so that future requests for that data can be served to improve performance. There are two kinds of Write Caching available:
 - (a) **Write-Back:** In a write-back (or write-behind) cache, writes are not immediately mirrored to the store. The data is written back to the backing store when evicted from the cache.

Note: It is recommended that battery backup unit be present to avoid data loss in case of system failure.
 - (b) **Write-Through:** In a write-through cache, every write to the cache causes a synchronous write to the backing store.
- iv. **Write caching without Battery Backup:** User can choose to perform Write caching without battery backup using this option. However this is not recommended since it may lead to data loss in case of system failure.
- v. **Read Ahead**
 - (a) **Adaptive Read Ahead:** In this policy, the controller initiates read-ahead only if the two most recent read requests accessed sequential sectors of the disk.

(b)Read Ahead: This policy facilitates reading multiple data records into the cache which in turn increases the chances another request can be satisfied with the data already in the cache.

(c)No Read Ahead: The controller does not use any Read Ahead Policy.

vi. Caching: User can choose to enable/disable caching.

vii. Stripe Size: Choose one of the pre defined values from the drop-down to set the stripe size of the RAID Array.

- Add to existing RAID Array

This action adds the selected disk to one of the existing array. User can also change the Array RAID level here, depending on the numbers of drives present in the array.

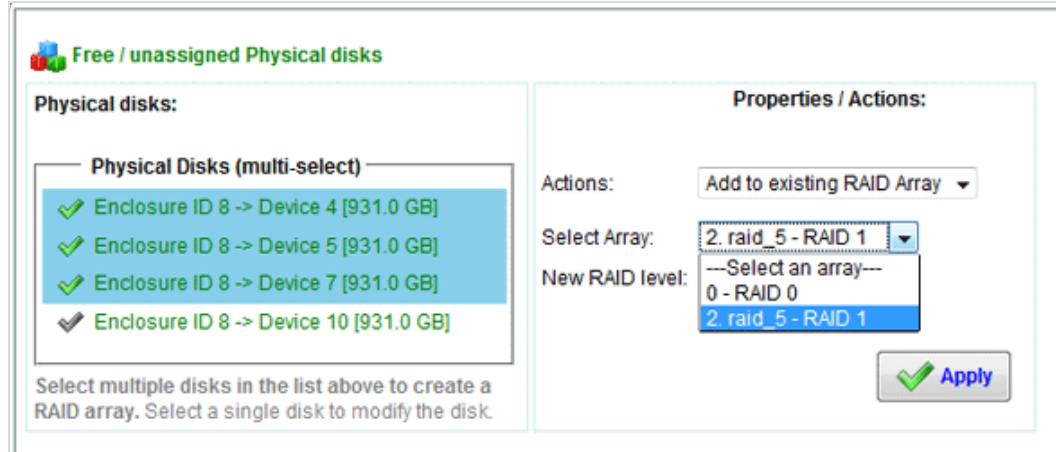


Figure 7.7.2 (c) - adding free/unassigned disks to an existing RAID array

Example: How to create a hardware RAID Array

- i. In the **Physical Disks** list, select multiple disks by clicking on them. The **Properties/Actions** section (on the right) will change to display related options/actions for creating a RAID array. Clicking on a single disk will display its properties.
- ii. In the *Actions* drop down, select *Create a new RAID Array*.

- iii. Select the RAID level (Available levels will differ depending on the number of disks selected).
- iv. Enter a name for the RAID Array to be created.
- v. Select the Write caching method.
- vi. If you want to perform Write caching without battery backup, then enable the *Write caching Without Battery Backup* option.
- vii. Select the Read Ahead policy or disable it completely.
- viii. Enable/disable caching
- ix. Select the Stripe size (if available).
- x. Click **Apply**.

If the array was created successfully, it will be displayed below the **Free / Unassigned Physical devices** section, with related properties and actions.

2. Hotspare/Standby Physical disks

Hot spares are disks that are kept on standby in case a RAID array fails. Then the hot spare drive is used to immediately rebuild the array to an optimal state, if the array is redundant. There are two types of Hot spares: Dedicated and Global. Dedicated hot spares are meant for specific RAID arrays. These disks will be used to rebuild only the specific array or arrays they are assigned to. If another array is degraded, they will not be used. Global hot spares can be used

to rebuild any degraded array. All controllers may not support both types of hot spares. Refer the RAID controller vendor documentation for the exact features supported.

Hotspare / Standby Physical disks

Physical disks:

Physical Disks

- ✓ Enclosure ID 252 -> Device 4 [1.818 TB]
- ✓ Enclosure ID 252 -> Device 5 [1.818 TB]

Select a single disk to modify the disk.

Properties / Actions:

ID / Slot: 5
Model: JK1174Y AHL5M9WHita
Serial: HDS722020ALA330
Status: In use, Hotspare, Spun Up
Type: Global

Actions: Remove Hotspare
Remove Hotspare
Locate / Blink Disk
Stop Blinking

Apply

Figure 7.7.2 (d) - Hot spare disks with and related Properties and Actions

2.1. Hot spare configuration settings

- Remove Hotspare: Removes the selected disk from the hot spare list.
- Locate/Blink Disk: This action will cause the corresponding Disk LED indicator to blink.

- Stop Blinking: This action causes the LED indicator to stop blinking.

3. Arrays

The RAID arrays configured on this controller is displayed here. The disks used by this RAID array are listed on the left side, and the status and actions for the Array are listed on the right. If a disk is selected, any available actions for the disk are shown on the right.

1 array1: RAID 0 | Size: 7.275 TB | [Edit](#)

Physical disks:

Physical Disks

- ✓ Enclosure ID 252 -> Device 0 [Array: 0, Span: 0] [1.818 TB]
- ✓ Enclosure ID 252 -> Device 1 [Array: 0, Span: 0] [1.818 TB]
- ✓ Enclosure ID 252 -> Device 2 [Array: 0, Span: 0] [1.818 TB]
- ✓ Enclosure ID 252 -> Device 3 [Array: 0, Span: 0] [1.818 TB]

Properties / Actions:

Array ID: **0**
 OS Drive: **sdh**
 RAID level: **0**
 Strip size: **128 KB**
 Status: **Optimal**
 Actions: [Edit Array settings](#) ▾
 Name: array1
 Write caching: [Write-Back \(enabled\)](#) ▾
 Note: depends on presence of Battery backup unit and caching policy on no battery.
 Write caching without Battery Backup: [Disabled \(recommended\)](#) ▾
 Note: Write-caching should not be enabled if Battery Backup is not present on the Controller.
 Read Ahead: [Adaptive Read Ahead](#) ▾
 Caching: [Enabled](#) ▾
 Disk Write Caching: [Enabled](#) ▾

[Apply](#)

Figure 7.7.2 (e) - Properties/Actions section for a RAID Array

3.1. RAID Array/Physical Devices configuration settings

3.1.1. RAID Array

The properties/ actions sections for a RAID Array is displayed by default. User can also access this section by clicking on Edit. User can choose to perform various actions on the RAID Array by selecting the appropriate option from the Actions drop down:

- Edit Array settings: Use this option to modify values set while creating the RAID Array.
- Quick Initialize Array: This option initializes the Array for use.
- Fully Initialize Array: This option initializes the complete Array for use.
- Reconfigure Array RAID level: User can change the Array.
- Blink/Identify all drives: This action will cause the corresponding Disk LED indicators of all the disks in the Array to blink.
- Stop Blinking: This action causes the LED indicators to stop blinking.
- Delete Array: This action will remove the RAID Array.

3.1.2. Individual Physical Device

Selecting an individual disk in the Physical Devices list displays related properties and Actions on the right.

The screenshot displays a management interface for a RAID array. On the left, under the heading "Physical Disks", there is a list of four disks, each with a checkmark icon and a blue highlight. The first disk is selected. On the right, under the heading "Properties / Actions:", the selected disk's details are shown. The status is "In use, Online, Spun Up". A dropdown menu for "Actions" is open, showing "Locate / Blink Disk" as the selected option. An "Apply" button with a checkmark icon is located to the right of the dropdown.

Physical Disks	
✓ Enclosure ID 252 -> Device 0 [Array: 0, Span: 0] [1.818 TB]	
✓ Enclosure ID 252 -> Device 1 [Array: 0, Span: 0] [1.818 TB]	
✓ Enclosure ID 252 -> Device 2 [Array: 0, Span: 0] [1.818 TB]	
✓ Enclosure ID 252 -> Device 3 [Array: 0, Span: 0] [1.818 TB]	

Properties / Actions:	
ID / Slot:	0
Model:	JK1181YAHP5XGVHita
Serial:	HDS722020ALA330
Status:	In use, Online, Spun Up
Actions:	Locate / Blink Disk Locate / Blink Disk Stop Blinking

Figure 7.7.2 (f) - Properties/Actions for a Physical Disk of a RAID Array

- **Locate/Blink Disk:** This action will cause the corresponding Disk LED indicator to blink.
- **Stop Blinking:** This action causes the LED indicator to stop blinking.

7.8 Fibre Channel

Fibre Channel, or FC, is a gigabit-speed network technology primarily used for storage networking.

- **Fibre Channel topologies**

There are three major Fibre Channel topologies, describing how a number of ports are connected together. A port in Fibre Channel terminology is any entity that actively communicates over the network, not necessarily a hardware port. This port is usually implemented in a device such as disk storage, an HBA on a server or a Fibre Channel switch.

- **Point-to-Point (FC-P2P).** Two devices are connected back to back. This is the simplest topology, with limited connectivity.

- **Arbitrated loop (FC-AL).** In this design, all devices are in a loop or ring, similar to token ring networking. Adding or removing a device from the loop causes all activity on the loop to be interrupted.
- **Switched fabric (FC-SW).** All devices or loops of devices are connected to Fibre Channel switches, similar conceptually to modern Ethernet implementations.

- **Modes available for the Fibre Channel adapter**

The FC adapter can be used in Target or Initiator mode. It defaults to target mode in Unified Storage Server, allowing Unified Storage Server to be configured as an FC target. Switching between target and initiator mode requires a restart of the appliance.

- **Fibre Channel use cases**

- **Target mode**

Please keep the following points in mind, before attempting to configure USS in FC target mode:

- FC target mode is supported only when USS is used as a standalone array (cannot be HA).
- You can use SAS / iSCSI / PCIe SSD devices as storage and provide LUNs to initiators.
- Emulex LPE12002 8Gbit/s FC HBA supported.
- Targets can be used as shared storage in Microsoft Windows Cluster Service.

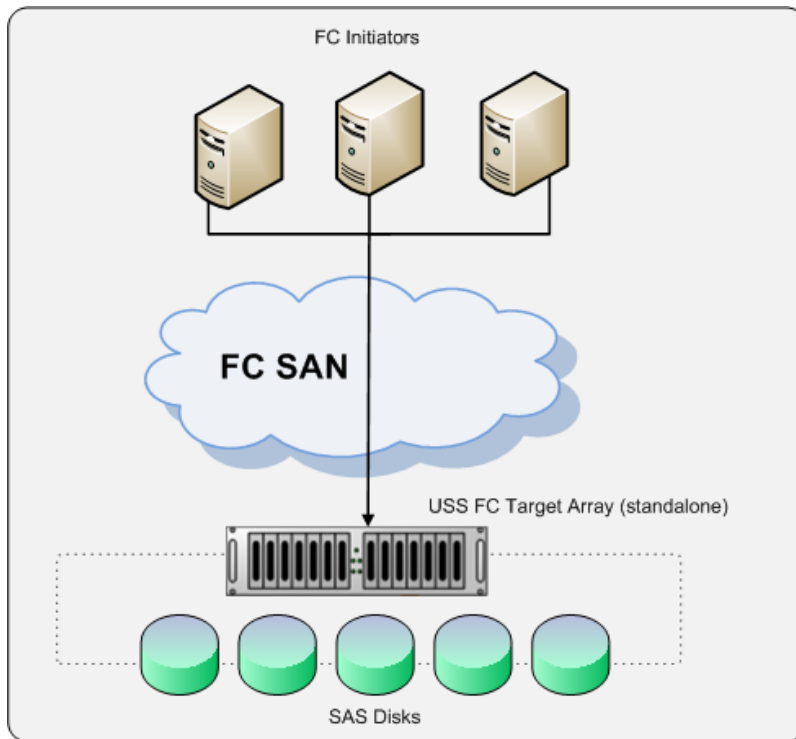


Figure 7.8 (a) - USS FC Target mode

- **Initiator mode**

Please keep the following points in mind, before attempting to configure USS in Initiator mode:

- FC Initiator mode is supported on both standalone and HA enabled array.
- You can use FC storage from another FC array.
- Emulex and Qlogic 8Gbit/s FC HBAs supported.
- NAS or iSCSI services can be provided on volumes using FC storage.

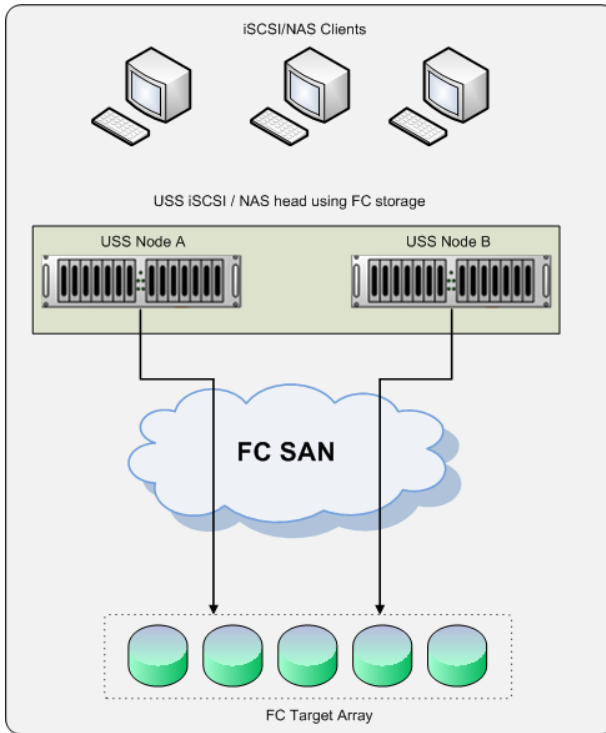



Figure 7.8 (b) - USS FC Initiator mode

7.8.1 Fibre Channel Initiator

- Sections of the interface

1. Summary of devices status:

The status of the FC ports detected on the system, the network topology and bandwidth status is shown here. Global parameters for the device driver are also available. The **Rescan all FC ports** button will rescan all FC targets available, and find all LUNs available. You can also update the HBA to the latest firmware using the **Upload & Update Firmware** option.



The screenshot shows a web-based interface for the Fibre Channel Initiator. At the top, there is a section titled "Summary:" with a minus sign icon. Below this, there is a button labeled "Rescan all FC ports" with a circular arrow icon. Underneath, the "Devices details:" section is displayed. It contains the following information:

- Ports status: 2 total, 1 active, 1 inactive.
- Link types: 0 in Loop, 0 in fabric, 1 Point To Point
- Bandwidth status: 1 x 8 Gb/s
- Model Description: QLogic PCI-Express Dual Channel 8Gb Fibre Channel HBA
- Serial Number: LFC0948R84890
- Firmware Version: 5.09.00

At the bottom of the "Devices details:" section, there is an "Upload & Update Firmware:" section. It includes a text input field, a "Browse_" button, and a green checkmark icon followed by an "upload" button.

Figure 7.8.1(a) -Summary of devices in Initiator mode

1.1. FC Summary properties

- Ports status: Total number of ports on the HBA and whether they are active/inactive is displayed here.
- Link types: It displays the FC Link Topology which can be either of Point-to-Point, Fabric and Loop.
- Bandwidth status: Bandwidth status for each port that is active is displayed here.
- Serial Number: This field displays the serial number of the HBA
- Firmware Version: Firmware version that's currently running on the HBA is displayed here.

1.2. Global parameters for the device driver:

In this section user can change the driver parameters of the HBA. Some parameters take into effect after reboot and some of the parameters are applied instantly. You can also restore the global parameters to their default values using the *Set all Params to Default* button.

Note: Parameters displayed here will differ depending on the FC adapter installed.



Warning: Do not change these settings unless you specifically need to do so. These settings alter the behavior of the device driver and all FC ports on the system. Please refer respective vendor's driver manual if you wish to tune these settings.

Global parameters:

Please select a parameter to view details or change for QLogic adapters:

- q12xextended_error_logging: 0 (default)
- q12xfdmienable: 0 (default)
- q12xloginretrycount: 20 (default)
- q12xlogintimeout: 20 (default)
- q12xmaxqdepth: 32 (default)
- q12xplogiabsentdevice: 0 (default)
- q12xqfullrampup: 120 (default)
- q1port_down_retry: 0 (default)

Option to enable extended error logging, Default is 0 - no logging. 1 - log errors

New setting:

New setting will take affect after an Appliance restart

Figure 7.8.1(b) - Global parameters for QLogic FC adapter

2. Devices list:

The list of Fibre Channel HBAs ports is shown here. Configuration and control options are available for each port.

- Devices:





HBA 0 Port 0  8 Gbps	QLogic QLE2562 Status: Link active, in Point To Point Node WWN: 20-00-00-1B-32-9B-3A-3C Port WWN: 21-00-00-1B-32-9B-3A-3C	
+ Configuration:		
HBA 0 Port 1 	QLogic QLE2562 Status: Link inactive Node WWN: 20-01-00-1B-32-BB-3A-3C Port WWN: 21-01-00-1B-32-BB-3A-3C	
+ Configuration:		

Figure 7.8.1(c) - Devices list with current status, control and configuration options.

2.1. Device command

- **Reset Port:** This will reset the port and cause it to re-initialize itself on the FC network.



Warning: Resetting the port is a disruptive operation, and will cause data loss if any data was being transmitted to an FC target to be saved.

2.2. Device configuration

2.2.1. Link settings

In this section the user can set Link speed and FC topology.

Configuration:

Link settings:

Link speed:

Fibre Channel Topology:

- Auto (Loop preferred)
- Auto (Point-to-Point preferred)
- Loop only
- Point-to-Point only


Figure 7.8.1(d) - FC port link settings for Emulex FC adapter

7.8.2 Fibre Channel Initiator – Target

- Sections of the interface

1. Target summary

The World-Wide-Names (Node WWNs, Port WWNs) for the target, and FC ID is shown here. Total number of LUNs that are exposed from the FC target (Total LUN count) and MPIO status is also shown.



The screenshot displays a 'Target Summary' window with the following information:

Node WWNs:	20:00:00:00:C9:8E:40:DF
Port WWNs:	10:00:00:00:C9:8E:40:DF
FC IDs:	000002
Total LUN count:	4
MPIO-enabled LUNs:	4
Non MPIO-enabled LUNs:	0

Figure 7.8.2(a) - Target information and LUN count.

2. Fibre Channel LUNs/Disks list:

The list of Fibre Channel LUNs shared to this FC initiator port from this target is shown here. The LUN size and any MPIO details of the LUN is shown on the right, when a LUN is selected.

Fibre Channel LUNs / Disks:

LUN 0: 10.74 GB MPIO(sdb [active])	LUN #	0
LUN 1: 10.74 GB MPIO(sdc [active])	Size:	10.74 GB
LUN 2: 10.74 GB MPIO(sdd [active])	MPIO:	Enabled, 1 paths
LUN 3: 10.74 GB MPIO(sde [active])	Path devices:	sdb
LUN 4: 21.47 GB MPIO(sdg [active])	Vendor:	DGC
LUN 5: 5.37 GB MPIO(sdh [active])	Model:	RAID 5
LUN 6: 12.88 GB MPIO(sdi [active])		

Figure 7.8.2 (b) - FC target LUNs list and LUN information.

7.8.3 Fibre Channel Target

- **Sections of the interface:**

- 1. Summary of devices status:**

FC Summary page for target displays the Current Host Bus Adapters (HBAs) mode and other device details such as Ports status, Link types, Firmware version etc. You can also update the HBA to the latest firmware using the **Upload & Update Firmware** option.

Summary:

Current HBA Mode: **Target Mode** 

Devices details:

Ports status: **2 total, 1 active, 1 inactive.**
Link types: **1 in Loop, 0 in fabric ,0 Point To Point**
Bandwidth status: **1 x 8 Gb/s**
Serial Number : VM93132811
Firmware Version : 2.00A4 (U3D2.00A4), sli-3

Upload & Update Firmware: No file chosen 

Figure 7.8.3(a) -Summary of devices in Target mode

1.1. FC Summary properties and actions

- **Current HBA Mode:** This tells the user in which mode the HBA is configured. This can be configured as Target or as an Initiator.
Note: Reboot is necessary for changing the mode.
- **Ports status:** Total number of ports on the HBA and whether they are active/inactive is displayed here.
- **Link types:** It displays the FC Link Topology which can be either of Point-to-Point, Fabric and Loop.
- **Bandwidth status:** Bandwidth status for each port that is active is displayed here.
- **Serial Number:** This field displays the serial number of the HBA
- **Firmware Version:** Firmware version that's currently running on the HBA is displayed here.

1.2. Global parameters for the device driver:

In this section user can change the driver parameters of the HBA. Some parameters take into effect after reboot and some of the parameters are applied instantly. You can also

restore the global parameters to their default values using the *Set all Params to Default* button.



Warning: Do not change these settings unless you specifically need to do so. These settings alter the behavior of the device driver and all FC ports on the system. Please refer respective vendor's driver manual if you wish to tune these settings.

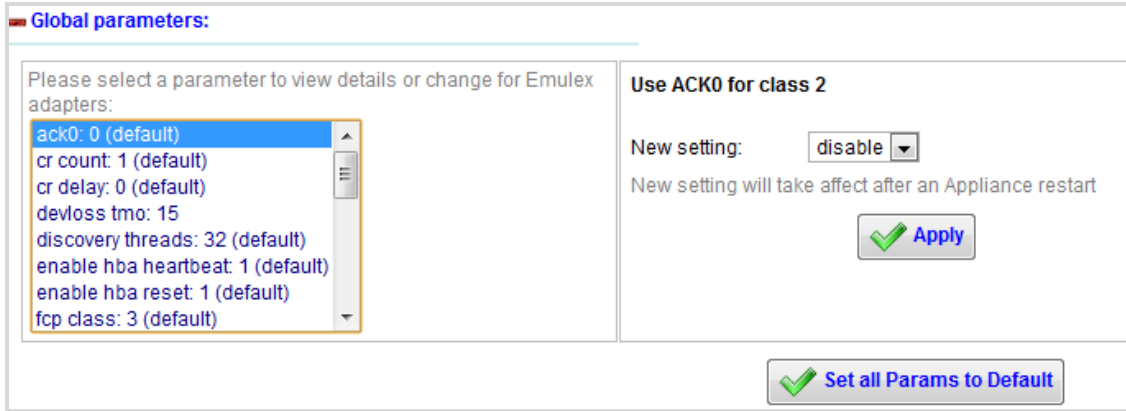


Figure 7.8.3(b) - Global parameters for Emulex FC adapter

2. Devices list:

The list of Fibre Channel HBAs ports is shown here. Configuration and control options are available for each port.



- Devices:	
HBA 0 Port 0	Emulex Corporation LPe12002-M8 Status: Link inactive Node WWN: 20:00:00:00:C9:8E:44:B6 Port WWN: 10:00:00:00:C9:8E:44:B6 
+ Configuration:	
HBA 0 Port 1	Emulex Corporation LPe12002-M8 Status: Link active, in Fabric Node WWN: 20:00:00:00:C9:8E:44:B7 Port WWN: 10:00:00:00:C9:8E:44:B7 8 Gb/s 
+ Configuration:	

Figure 7.8.3(c) - Devices list with current status, control and configuration options.

2.1. Device command

- Reset Port: This will reset the port and cause it to re-initialize itself on the FC network.



Warning: Resetting the port is a disruptive operation, and will cause data loss if any data was being transmitted to an FC target to be saved.

2.2. Device configuration

2.2.1. Link settings

In this section the user can set Link speed and FC topology.

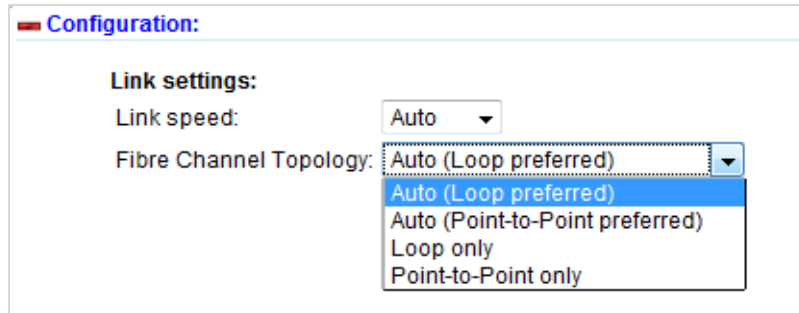


Figure 7.8.3(d) - FC port link settings for Emulex FC adapter

7.8.4 Fibre Channel Target-LUN

- **Sections of the interface:**

1. Current LUN Configuration

This section lists the LUNs that are exposed through the FC target. User can select this LUN and view details and the user can also remove the LUN from the FC target using the Delete LUN button.

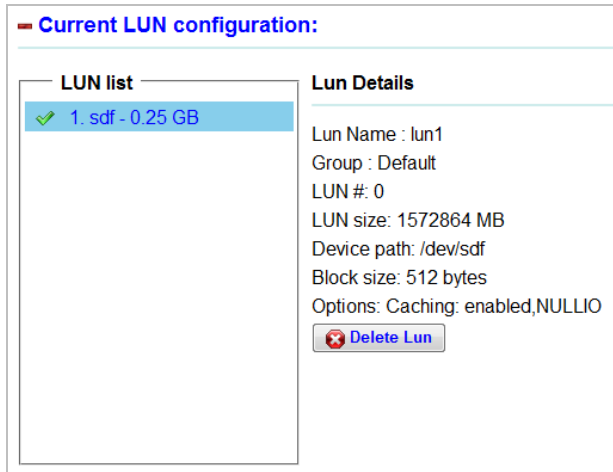


Figure 7.8.4(a) - Current LUN Configuration

2. Add a LUN

In this Section, user can select the existing storage devices that can be added to the FC target. The LUNs selected here will be added to the selected group. User needs to select the LUN from the list, give the FC LUN name and select the group. An option to add the LUN from the existing storage-pool is also available.

- Caching: There are two modes of caching available: disabled (WRITE_THROUGH) and enabled (NV_CACHE). In disabled mode, changes made to cached data are simultaneously made in the original copy. User can also change the attribute of the lun i.e make it read-only/read-write only while adding it to LUNs list.
- Read-write: User can also change the attributes of the disk/LUN i.e. make it READ_ONLY/READ_WRITE. In READ_WRITE, Disk/LUN is exposed with read and write permissions. Whereas, in READ_ONLY, they have only read permissions.

Note: For Linux Initiators, after discovering and connecting to a LUN device, format it with XFS or ext4 file system before using it.

Add a LUN:

Add a LUN by specifying an existing device

Select the Storage device: pool1/vol2: 4.10 GB

FC LUN name: lun1

Caching: enabled

Permissions: read+write

ACL group: group1

Add a LUN by allocating space from a storage pool

Storage Pool for the LUN: -----

New volume name:

New volume size: GB

FC LUN name:

Caching: enabled

Permissions: read+write

ACL group: group1




Figure 7.8.4(b) - Adding a LUN

Example: Adding an FC Target LUN by specifying an existing device

Using this method you can use an existing storage device (logical volume, snapshot device, clone, RAM device, etc) as FC Target LUN. Please note that only free/unassigned devices can be used.

- i. Click the **Add a LUN by specifying an existing device** radio button to enable it.
- ii. Select from the available storage device (listed in black) to use as FC Target LUN. Devices listed in gray are in use and cannot be selected.
- iii. Specify a name for the LUN.
- iv. Enable/disable Caching.
- v. Set read-write or read-only permission for the LUN being added.
- vi. Select the ACL group, the users of which you want to expose the LUN to.
- vii. Click the **Add LUN** button.
- viii. If the device was successfully added, it will appear in the LUN list in the **Current LUN configuration** section.

Example: Adding an FC Target LUN by creating a logical volume.

Using this method, you can create a logical volume and then use it as FC target LUN.

- i. Click the **Add a LUN by allocating space from a storage pool** radio button to enable it.
- ii. Select the storage pool, to create the logical volume.

- iii. Specify a name for the logical volume and size.
- iv. Enable/disable Caching.
- v. Set read-write or read-only permission for the LUN being added.
- vi. Select the ACL group, the users of which you want to expose the LUN to.
- vii. Click the **Add LUN** button.

If the device was successfully added, it will appear in the LUN list in the **Current LUN configuration** section.

7.8.5 Fibre Channel Target – ACLs

- **Sections of the interface**

1. **Currently configured groups**

Here, the user can view the current existing groups and the LUNs belonging to them. User has the option to remove the Security Group, and LUNs and FC initiators WWN for that particular group. The order of LUNs can also be changed.

Note: The option to remove a Security Group is available only when all the LUNs and FC Initiators WWN belonging to that group have been removed.

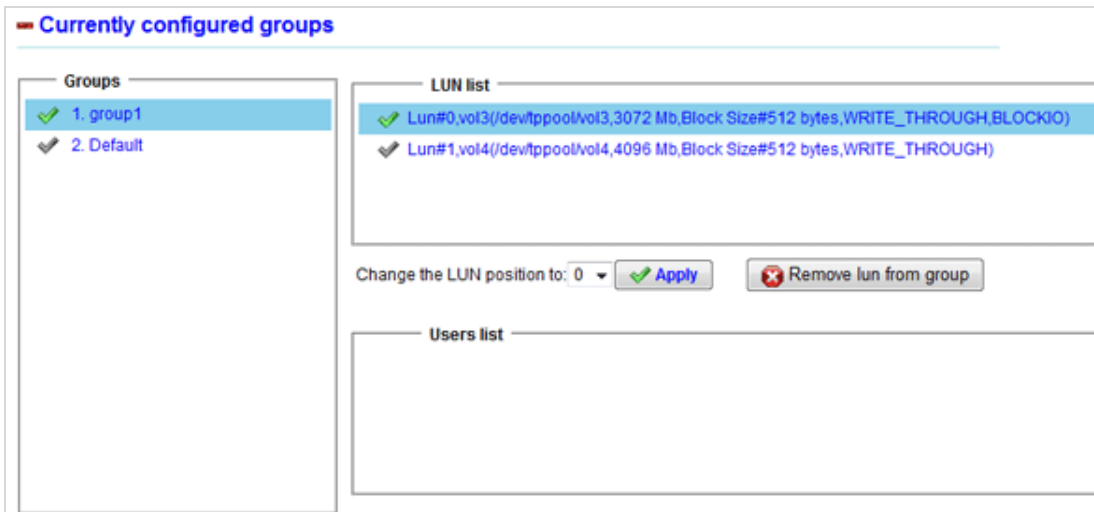



Figure 7.8.5(a) - currently configured groups and LUNs list

2. Add new Group

User can create security group here.



- Add new Group

New Group Name :


 Add Group

Figure 7.8.5 (b) - Adding a new Security Group

3. Add new User

User can add new FC initiators WWN for a particular group here. The initiators that are assigned to particular group will access the LUNs that are present in that group. User has option to remove the assigned initiators WWN from the security group.

7.9 Volume management

- Sections of the interface

1. Storage Pools

This section lists pools created in OS Disk devices section with the name, type and usage details of each pool.

Note: Size of Pools, Volumes and Physical devices listed in this section are calculated using the convention 1KB=1024bytes, 1MB =1024 KB etc.

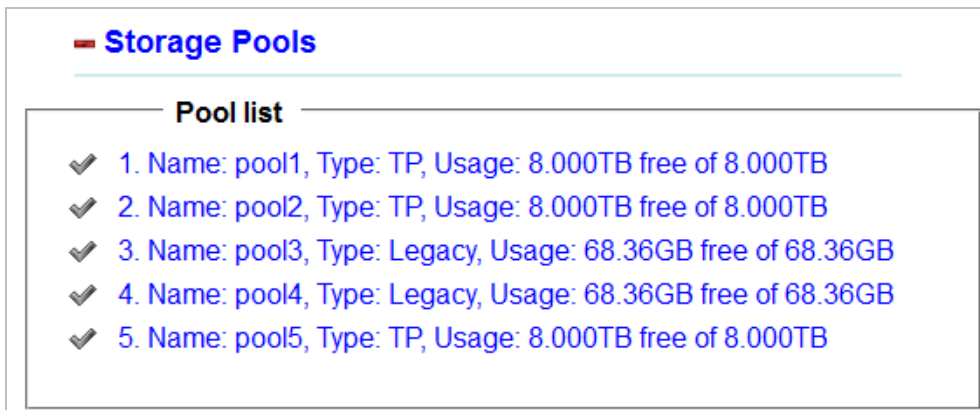







Figure 7.9(a) - pool list summary

1.1. Pool Details

- Pool details: This section on the left displays various properties related to the selected pool. e.g. whether the pool is Writeable or not, Chunk Size, Snapshots count, Logical/Allocated Volume count etc.
- Usage bar: The usage bar indicates the usage status of the pool. Various colors have been used to illustrate the status of the pool. Please refer to the color legend below for more information:


color	percentage usage
	100
	≥ 90
	90< (Allocated Volumes Space)
	90< (Used Physical Disk space)
	free space

Usage bar color legend





1.2. Inactive Pool Actions

You can activate a TP pool by enabling/disabling RAM/SSD cache here. The allowed RAM/SSD cache size will also be displayed, if RAM/SSD cache was selected while activating the pool. You can also delete the pool.

This is particularly beneficial when USS is upgraded from previously released version 2.0 to this version. Or, in case pools and volumes were created without caching in the latest version and then user decides to utilize the feature later. In either scenario, volume/pool configuration remains unaffected.

Pool: **pool1** | Type: **Chelsio TP** | Size: 8.000TB | Free: 7.849TB |  Edit

Inner Devices:

Physical devices :	Logical Volumes :
 sdc	 vol1  vol2  vol3

Details and Actions:

Disk for Cache:

SSD disks :

SSD Cache Mode :

RAM Cache :



 **Activate the pool**  **Delete the Pool**

Figure 7.9(b) - pool list summary Inactive pool actions

1.3. Cache counters - Chelsio Volume Management (TP)

If you have a TP pool configured with RAM/SSD cache, you can view the cache usage statistics here. The *Clear counters* button resets HIT/MISS counters to zero.

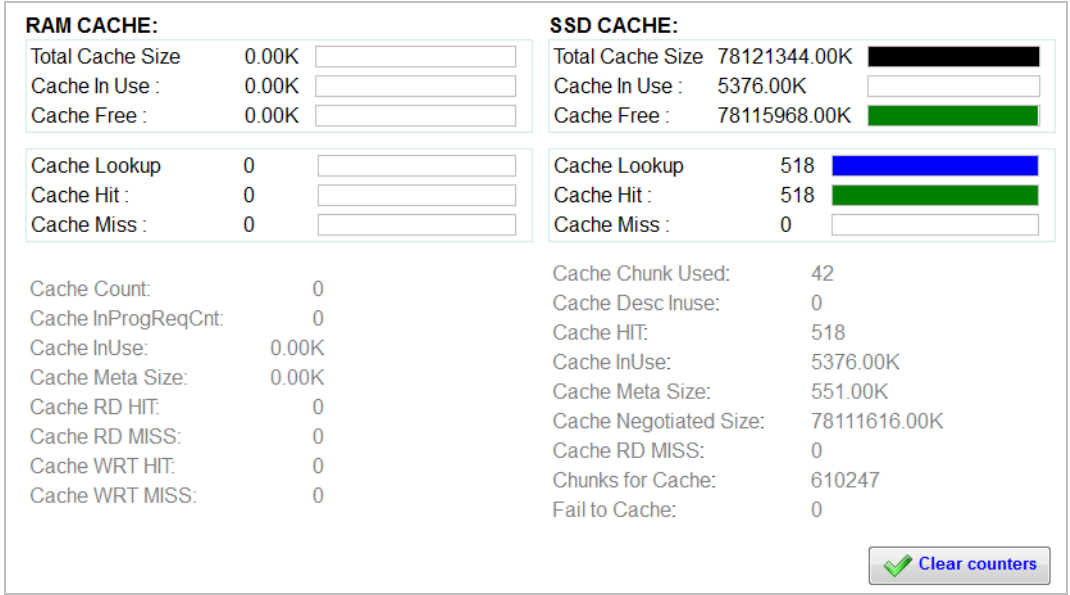


Figure 7.9(c) – Cache counters for a SSD cache enabled pool

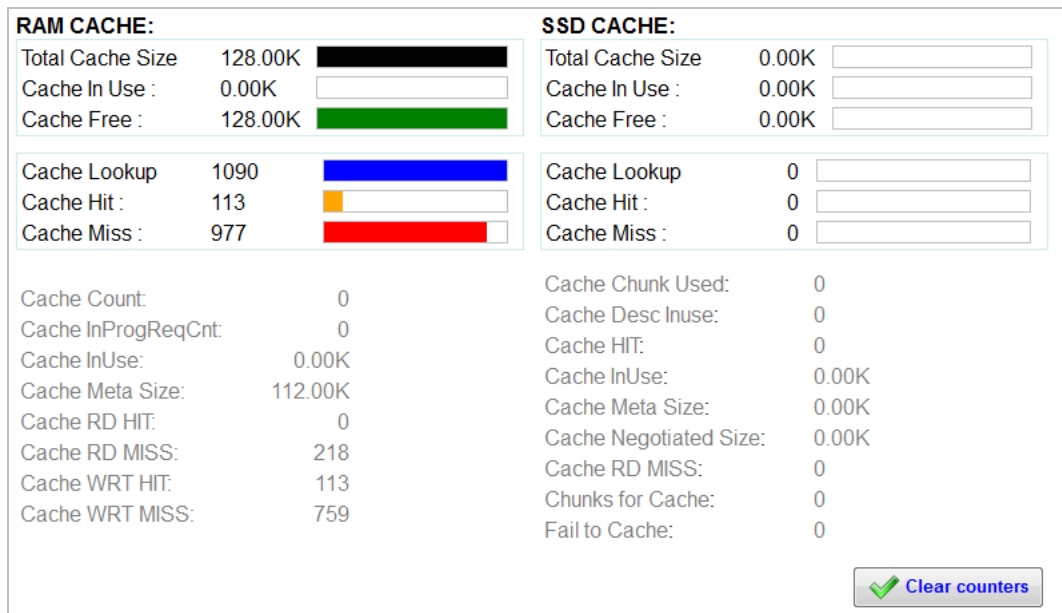
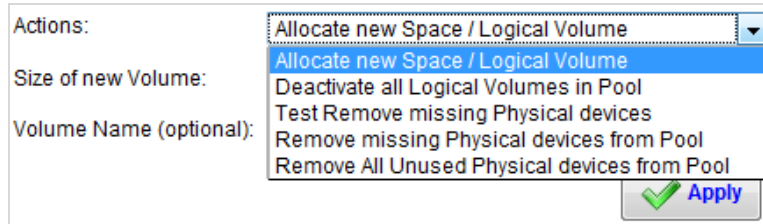


Figure 7.9(d) – Cache counters for a RAM cache enabled pool

1.4. Pool Actions

1.4.1. Logical Volume Management (Legacy)

User can perform various actions on a selected Legacy pool as listed below:



Actions: Allocate new Space / Logical Volume

Size of new Volume:

Volume Name (optional):

Allocate new Space / Logical Volume
Deactivate all Logical Volumes in Pool
Test Remove missing Physical devices
Remove missing Physical devices from Pool
Remove All Unused Physical devices from Pool

Apply

Figure 7.9(e) - Legacy Pool Actions

- Allocate new space/Logical volume: Create new logical volumes on the pool by specifying the name (optional) and size for the logical volume. In case, there is no sufficient free space on the pool, then new volumes can be created after adding new disks to the existing pool or on a new pool.




Figure 7.9(f) - Creating new Logical Volumes

- Activate/Deactivate All logical volumes in Pool: Enable/Disable all logical volumes previously created on the pool.

Note: Only the volumes which are free/ unassigned will be deactivated. The option to activate will be available only if all the volumes in the pool are deactivated/ disabled.

- Rename the Pool
Rename the selected pool.

Note: This option is available if there are no volumes configured on that pool OR if all the volumes are deactivated.

- Test Remove missing Physical devices: This option checks whether the 'Remove missing physical devices' option will succeed on the selected pool or not.

Note: This above option shows failed message, when the missing physical disk in the pool has volumes configured on it. To remove such disks, the volumes have to be removed first.

- Remove missing physical devices from Pool: Remove any physical device which is not present, from the pool.
- Remove All Unused Physical devices from Pool: This option removes the disks from the pool on which there are no volumes created.
- Delete the Pool: Delete the selected pool.

Note: This option is available only after all the Logical Volumes on the selected pool have been deleted.

1.4.2 Chelsio Volume Management (TP)

User can perform various actions on a selected TP pool as listed below:

Actions : Allocate new Space / Logical Volume ▼

Size of volume as % of freespace in pool: Allocate new Space / Logical Volume

Size of new Volume: Resize the pool

Volume Name: Recover the pool

Device for data: Modify Metadata Allocation

Allocation zone: Distributed on all devices ▼

Enable volume encryption: Auto Select ▼

Deactivate all Logical Volumes in Pool

Apply

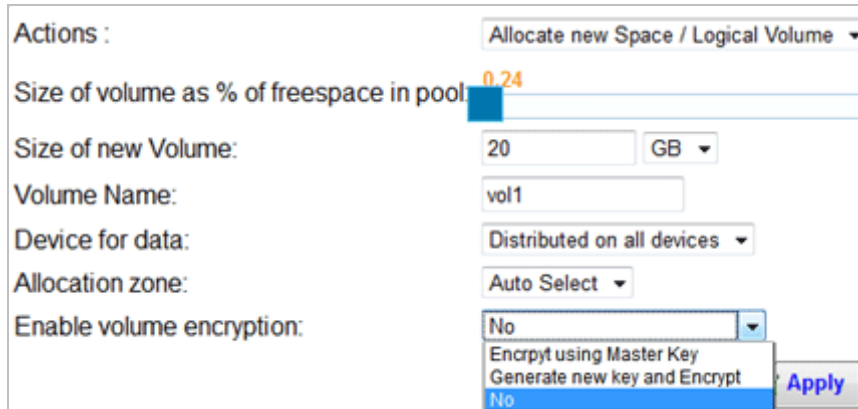
Figure 7.9(g) - TP Pool Actions

- Allocate new space/Logical volume: Create new logical volumes on the pool by specifying the name and size for the logical volume. The user can also choose to encrypt the Volume. If there is no sufficient free space left on the pool, new volumes can be created only on a different/new pool.

The volume can be encrypted in two ways:

- i) Encrypt using Master Key: The volume will be created and encrypted using the Master Key (Pass Phrase) which was generated when the Volume Management section was accessed for the first time.
- ii) Generate new Key and Encrypt: User can choose to generate a new key and encrypt the volume.

Note: Master Key (Pass Phrase) is required for generating a new key.



The screenshot shows a configuration window for creating a logical volume. The 'Actions' dropdown is set to 'Allocate new Space / Logical Volume'. The 'Size of volume as % of freespace in pool' is set to 0.24. The 'Size of new Volume' is 20 GB. The 'Volume Name' is 'vol1'. The 'Device for data' is 'Distributed on all devices'. The 'Allocation zone' is 'Auto Select'. The 'Enable volume encryption' dropdown is open, showing options: 'No', 'Encrypt using Master Key', 'Generate new key and Encrypt', and 'No' (highlighted). An 'Apply' button is visible at the bottom right.

Figure 7.9(h) - Creating a Logical Volume using encryption

- Deactivate the Pool: Using this option User can disable the selected pool.
- Resize the pool: Using this option User can resize the selected pool.
Note: The aggregate size of all the pools configured on the appliance cannot exceed the storage capacity with which the appliance is licensed.
- Activate/Deactivate All logical volumes in Pool: Enable/Disable all logical volumes previously created on the pool.
Note: Only the volumes which are free/ unassigned will be deactivated. The option to activate will be available only if all the volumes in the pool are deactivated/ disabled.
- Rename the Pool: Rename the selected pool.
Note: This option is available if there are no volumes configured on that pool.
- Recover Pool
If any pool/volume operation is complaining of corrupt metadata, the pool/volume metadata can be restored using this option.
Note: Only metadata can be recovered, it has no effect on the data stored in volumes.
- Modify Metadata Allocation

Using this option, you can specify the disk (Metadata device) and/or zone in which volume metadata will be stored or let the appliance automatically choose the appropriate settings (default).

- Delete the Pool

Delete the selected pool.

Note: This option is available only after all the Logical Volumes on the selected pool have been deleted.

- Refresh Pool

Use this option if pool goes to STALE state due to I/O errors detected on the pool when the disk is pulled /reinserted. The future I/O operations to the pool will succeed only after the pool is refreshed.

1.5. Physical Devices actions

1.5.1. Logical Volume Management (Legacy)

Click on any one Physical Device to view the various Details and Actions.

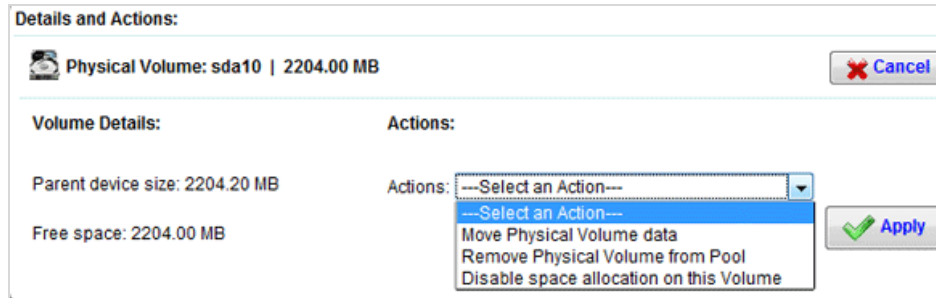


Figure 7.9(i) - Physical Devices Details and Actions

- Move Physical Volume data (not available when only one physical device exists): This option allows user to move the allocated physical extents (PEs) on *SourcePhysicalVolume* to any other physical volume (PV).

- Remove Physical Volume from Pool (not available when one physical device exists): Removes the selected Physical Device from the Pool.
- Disable space allocation on this Volume: This option can be used when user does not want to create any volumes on the selected disk.

1.5.2. Chelsio Volume Management (TP)

- Extend the Physical Volume: This option is to extend the size of the existing Physical Volume (RAID device) added into the Pool.
- Replace Disk
Using this option, a TP disk can be replaced with a new disk with storage capacity greater than or equal to the current disk.

Note: This option is not applicable to SSD cache disk.

Example: How to replace a TP disk (Live Migration)

Before going ahead, please make sure that you have a free/unassigned spare physical disk attached to the system with storage capacity greater than or equal to the disk to be replaced. You will also need a

free/unassigned disk to be used as Log device to store metadata during the live migration process. The selected disk will be free once the process is completed.

- i. In the **Volume Management** module, expand the **Storage Pools** section.
- ii. In the **Pool list**, select the pool, the physical device of which is to be replaced.
- iii. In the **Physical devices** section, click on the disk to be replaced.
- iv. In the **Actions** drop-down, select *Replace disk*.
- v. In the **Physical disks** drop-down, select the disk which will replace the selected disk.
- vi. In the **Log Device** drop-down, select the disk to be used as log device.
- vii. Click **Apply**.
- viii. Disk replacement process will now take place which may take some time depending on the size of the disk being replaced. You can click on the pool again to view the status of the process.
- ix. You can cancel the process anytime by selecting either of the disks and selecting *Abort Disk replacement*.

1.6. Logical Volume actions

1.6.1. Logical Volume Management (Legacy)

After a Logical volume is created, click on the Logical Volume icon to view and perform various actions as mentioned in the figure below.

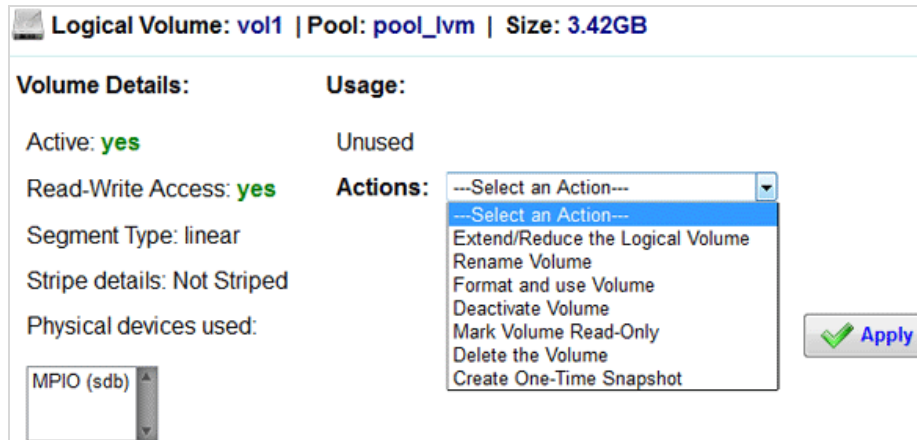


Figure 7.9(j) - Logical Volume Details and Actions – Legacy

- **Extend the Logical Volume:** This option is to increase the size of the existing logical volume. For example, if the size of the logical volume is 5GB and needs to be changed to 6GB, then enter 6GB as the size in the New Volume size box. Logical Volume size can also be decreased.

Note: Logical volume size cannot be decreased if logical volume is part of a target and the target is running currently, if the volume is active or if the volume has been formatted with a file system.

The screenshot shows a web-based interface for managing a logical volume. At the top, it displays the volume name 'vol1', the pool 'pool_lvm', and the current size '3.42GB'. Below this, there are two columns: 'Volume Details' and 'Usage'. Under 'Volume Details', it shows 'Active: yes', 'Read-Write Access: yes', 'Segment Type: linear', 'Stripe details: Not Striped', and 'Physical devices used: MPIIO (sdb)'. Under 'Usage', it shows 'Unused'. To the right of the 'Usage' section, there is an 'Actions:' dropdown menu currently set to 'Extend/Reduce the Logical Volume'. Below the actions menu, there is a 'New Volume size:' input field followed by a 'MB' dropdown menu. At the bottom right of the interface, there is a green 'Apply' button with a checkmark icon.

Figure 7.9(k) - Extending a Logical Volume

- **Rename Volume:** This option can be used to give a different name to the logical volume.
- **Use volume, attach to folder**
Using this option, a logical volume can be mounted to a folder.
Note: This option is only available for volumes which contain file systems.
- **Erase Data, format and use Volume**
Remove existing data, format with xfs file system and mount the volume to the folder selected from browse.
- **Run filesystem check**
This option checks a logical volume for file system errors and provides options to fix them.
- **Format and use Volume:** This option is used to format the volume with a filesystem and use it for NAS or Lustre.
 - a) **Folder to attach (to mount):** Choose the directory on the system to attach the volume to, using the Browse button.
 - b) **Volume Usage:** The volume can be formatted for use in regular NAS, such as CIFS, NFS, FTP, HTTP shares or as a Lustre OST / MDT.

c) Lustre Target Type:

- i. Metadata Server (MGS/MDS) - The MDS makes metadata stored in one or more MDTs available to Lustre clients. Each MDS manages the names and directories in the Lustre file system(s) and provides network request handling for one or more local MDTs.
- ii. MGS+MDT- Using this option you can create a combined MGS/MDT file system.
- iii. Metadata Target (MDT) - The MDT stores metadata (such as filenames, directories, permissions and file layout) on storage attached to an MDS. Each file system has one MDT. An MDT on a shared storage target can be available to multiple MDSs, although only one can access it at a time. If an active MDS fails, a standby MDS can serve the MDT and make it available to clients. This is referred to as MDS failover.
- iv. Object Storage Target (OST): User file data is stored in one or more objects, each object on a separate OST in a Lustre file system. The number of objects per file is configurable by the user and can be tuned to optimize performance for a given workload.

d) File system name: name of the Lustre share.

- e) IP/Hostname of MGS: Machine IP Address/hostname where MDS/MGS is configured.
- f) IP/Hostname of Secondary MGS (if MGS clustered): Machine IP Address/hostname of the secondary node where MDS/MGS is configured (HA mode).
- g) Interface: Interface on which Lustre network is to be configured.
- h) Disable caching: This disables caching on the OS for this volume, and causes every I/O operation to be sent to the physical device. Disabling this will affect performance for the users of the volume.

Note: This setting is not available (by default set to 'No') for encrypted volumes.

Actions:	Format and use Volume
Folder to attach(to mount):	/shares/mdt <input type="button" value="Browse.."/>
Volume Usage:	<input type="radio"/> Regular NAS <input checked="" type="radio"/> Lustre
Lustre Target Type:	<input type="radio"/> MGS <input type="radio"/> MGS+MDT <input checked="" type="radio"/> MDT <input type="radio"/> OST
File system name:	lustre1
IP/Hostname of MGS:	102.50.50.76
IP/Hostname of Secondary MGS (if MGS is clustered):	
Interface:	tcp0(102.50.50.76)
<p>Note: 1) It is recommended to format MGS/MDT and OST in different servers. 2) Reformatting/Mounting should be done in this order: MGS, OST and MDT.</p> <p>1) Any existing data on the volume will be lost!</p>	
<input type="button" value="Apply"/>	

Figure 7.9(l) - Format and use Volume settings

- **Activate/Deactivate Volume:** This option can be used to enable/disable the Volume.

Note: This action is available only if the volume is not attached to a folder or used as target LUN device.

- **Mark Volume Read Only:** This option is selected to set Read Only permissions for the volume.
Note: Permission should not be set to Read Only if it is Read Write and is part of running target.
- **Delete the Volume:** This option will delete the selected volume.
Note: Logical Volume cannot be deleted until all snapshots created on the logical volume are deleted.
- **Create snapshot:** With this option user can create a snapshot volume with the specified size and name which would backup the data at the point of creating the snapshot.

Logical Volume: vol1 | Pool: pool_lvm | Size: 3.42GB

Volume Details:	Usage:
Active: yes	Unused
Read-Write Access: yes	Actions: <input type="button" value="Create One-Time Snapshot"/>
Segment Type: linear	
Stripe details: Not Striped	Snapshot Volume Name: <input type="text"/>
Physical devices used:	Snapshot Volume size: <input type="text"/> MB
<input type="text" value="MPIO (sdb)"/>	This creates a one-time point-in-time image of the Volume pool_lvm_vol1 and allows access to this point-in-time copy, until it is deleted or space allocated to it runs out.
File system: ext4	Warning: No IO should be running to the Volume, to get a consistent, good snapshot of the volume.

Figure 7.9(m) - Create a One-Time Snapshot

1.6.2. Chelsio Volume Management (TP)




Figure 7.9(n) - Logical Volume Details and Actions – TP

- **Extend/Reduce the Logical Volume:** This option is to increase/decrease the size of the existing logical volume. For example, if the size of the logical volume is 5GB and needs to be changed to 6GB, then enter 6GB as the size in the New Volume size box. Logical Volume size can also be decreased in a similar manner.

Note: Logical volume size cannot be decreased if logical volume is part of a target and the target is running currently.



Figure 7.9(o) - Extend/Reduce Logical Volume

- **Rename Volume:** This option can be used to give a different name to the logical volume.
- **Use volume, attach to folder**
Using this option, a logical volume can be mounted to a folder.
Note: This option is only available for volumes which contain file systems.
- **Format and use Volume:** This option is used to format the logical volume with xfs file system and mount the volume to the folder selected from browse.

- Erase Data, format and use Volume

Remove existing data, format with xfs file system and mount the volume to the folder selected from browse.

- Run filesystem check

This option checks a logical volume for file system errors and provides options to fix them.

- Activate/Deactivate Volume: This option is to enable/disable the Volume.

Note:

- This action is available only if the volume is not attached to a folder or used as target LUN device.
- If the volume is encrypted, pass phrase is required to activate.

- Modify data Allocation

Using this option, you can specify the disk and/or zone in which data will be stored or let the appliance automatically choose the appropriate settings (default).

- Create a clone: Using this option, User can replicate a Logical Volume.

- **Schedule Snapshots:** With this option, user can schedule snapshot creation for a volume on an hourly, daily or weekly basis.

Usage:
Unused

Actions:

Hourly Snapshots: Disabled

Daily Snapshots: Enabled
Create a snapshot every day at 00:05
Keep the snapshots of last days.

Weekly Snapshots: Disabled




Figure 7.9(p) - scheduling a daily snapshot

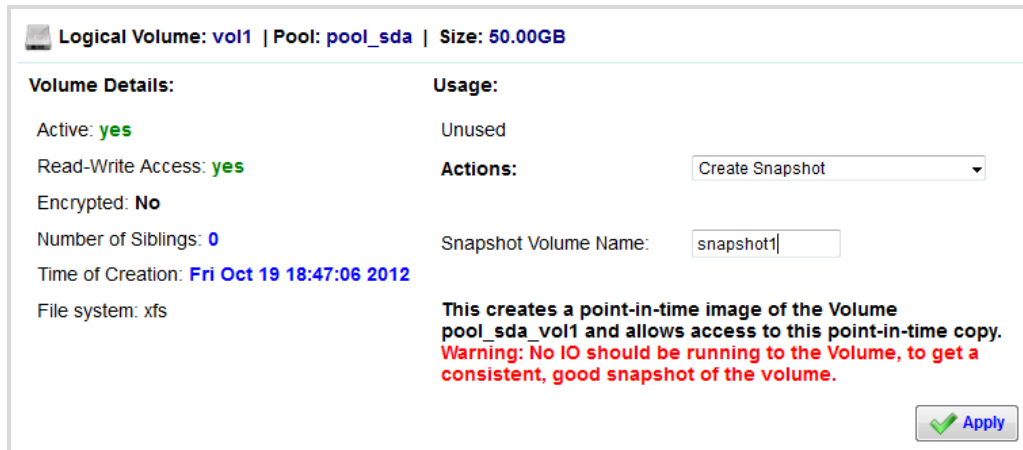
Example: Scheduling Snapshots

- i. Select the unencrypted logical volume for which you want to schedule snapshots. This will navigate to the logical volume properties/actions page.
- ii. In the **Actions** drop-down, select *Schedule Snapshots*.
- iii. Select if you want the snapshot to be created on an hourly, daily or weekly basis by selecting the checkbox for each.
- iv. The *Keep the snapshots of last...* option will preserve the previously created snapshots based on the number selected in the drop-down for each option. For example, if you specify 4 for the hourly snapshots, then as soon as the 6th hourly snapshot is created by the scheduler, it will remove the oldest hourly snapshot.
- v. For weekly snapshots, you can select the days on which snapshots will be created.
- vi. To change the time at which a snapshot is taken, expand the **Settings** section and change the time as need in the *Snapshot Scheduling Settings*.
- vii. Click **Apply**. You can view the details of the scheduled snapshot in the **Settings** section.

viii. To change any settings, click on the volume and choose select *Schedule Snapshots* in the **Actions** dropdown.

- Create Snapshot: With this option user can create a snapshot volume with the specified name which would backup the data at the point of creating the snapshot. Multiple Snapshots can be created.

Note: If the volume is encrypted, pass phrase is required to create a Snapshot.



The screenshot shows a dialog box for creating a snapshot of a logical volume. At the top, it displays the volume's metadata: "Logical Volume: vol1 | Pool: pool_sda | Size: 50.00GB". The dialog is divided into two columns: "Volume Details" and "Usage".

Volume Details:	Usage:
Active: yes	Unused
Read-Write Access: yes	Actions: <input type="text" value="Create Snapshot"/>
Encrypted: No	Snapshot Volume Name: <input type="text" value="snapshot1"/>
Number of Siblings: 0	
Time of Creation: Fri Oct 19 18:47:06 2012	
File system: xfs	

Below the "Usage" section, there is a red warning message: "This creates a point-in-time image of the Volume pool_sda_vol1 and allows access to this point-in-time copy. Warning: No IO should be running to the Volume, to get a consistent, good snapshot of the volume." At the bottom right, there is an "Apply" button with a green checkmark icon.

Figure 7.9(q) - Creating a Snapshot – TP

- Delete the Volume

This option will delete the selected volume.

Note: Logical Volume cannot be deleted until all snapshots created on the logical volume are deleted.

1.7. Snapshot Actions

1.7.1. Logical Volume Management (Legacy)

Once Snapshot is created; it can be extended, renamed or deleted.




Figure 7.9(r) - Snapshot Actions – Legacy

- **Extend the Snapshot Volume:** With this option, the size of the selected Snapshot can be increased. For example, if current size is 5GB and needs to be changed to 6GB, then enter 6GB as the size in the New Snapshot size box.

- Rename the Snapshot Volume: The selected Snapshot can be given a different name using this option.
- Use volume, attach to folder: Using this option, a snapshot can be mounted to a folder. Note: This option is only available for snapshots which contain file systems.
- Delete the Snapshot: This option will delete the selected Snapshot

1.7.2. Chelsio Volume Management (TP)

Once Snapshot is created, it can be renamed, deleted, deactivated or reverted back.

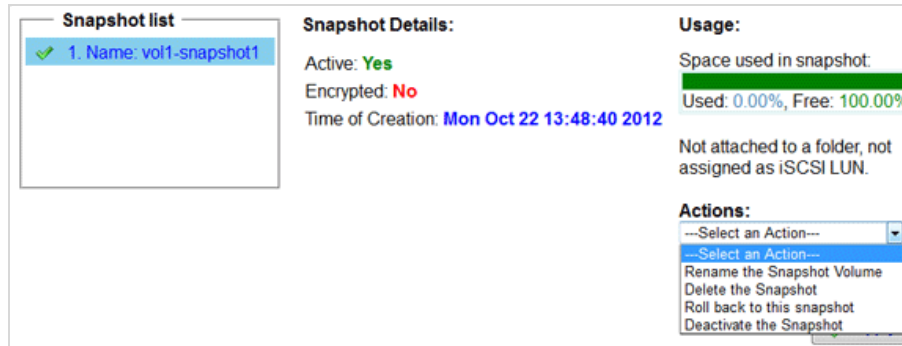


Figure 7.9(s) - Snapshot Actions – TP

- **Rename the Snapshot Volume:** The selected Snapshot can be given a different name using this option.
- **Delete the Snapshot:** This option will delete the selected Snapshot.
- **Use volume, attach to folder:** Using this option, a snapshot can be mounted to a folder.
Note: This option is only available for snapshots which contain file systems.
- **Rollback to this Snapshot:** Using this option, the Volume can be reverted back to a particular Snapshot, displayed in the Snapshot List.

Note:

- Reverting back to a Snapshot will lead to the deletion of all the Snapshots created after the selected Snapshot.
- This option is available only when the Logical Volume and Snapshot are active and free/unused.
- Encryption key is required, if the Logical Volume is Encrypted.
- **Activate/Deactivate Snapshot:** This option is to enable/disable the Snapshot.

Note:

- This action should not be performed if the Snapshot is part of target and is running.
- If the Snapshot is encrypted, pass phrase is required to activate.

Example: How to roll back a logical volume

Here are the steps to roll back a logical volume to a particular snapshot

- i. Before a logical volume is rolled back, please ensure that the volume and the snapshot are free/ unused. For example, the volume and snapshot shouldn't be attached to any folder or used as an iSCSI LUN device. Also, since all the snapshots created after the selected snapshot will be deleted during the process, please ensure that they also are free/unused.
- ii. In the **Snapshot-list**, select the snapshot to which you want the logical volume rolled back to.
- iii. If the snapshot is not active, select *Activate the Snapshot* in the **Actions** drop-down and click **Apply**.
- iv. Click on the snapshot again. This time, select *Roll back to this snapshot* and click **Apply**.

If successful, the volume will be rolled back to the selected snapshot. The volume should be active and functional.

2. Encryption Settings

Here, the user can reset the Master Pass phrase and Volume Pass phrase used during encryption of Volumes and Snapshots. User can also unlock any newly added encrypted volumes.

Encryption Settings:
Note: Corresponding recovery keys are available in the Encryption Settings.

Snapshot Scheduling Settings:

Hourly Snapshots time: 10 minutes past the hour

Daily Snapshots time: 00 : 05 AM

Weekly Snapshots time: 00 : 15 AM

Scheduled snapshots:

Type	Pool	Volume	Options
hourly	pool1	vol1	Preserve the last 3 snapshot.

Figure 7.9(t) - Encryption Settings

2.1. Encryption Settings Actions

- Reset Master Pass Phrase: Using this option, user can reset the Master Key (Pass Phrase).

Note: The new pass phrase has to be of minimum 6 characters.

Reset Pass Phrase

New Pass phrase:

Confirm Pass phrase:

Please paste the unencrypted recovery key here:

```
-----BEGIN RSA PRIVATE KEY-----
MIIBOwIBAAJBAAOjDLS8ujGSj1gLtrtn50AhRBwMaMwsvRogNjFQH0cU8bXNV63Hz
z4jGqmMHAct9rf8857buGGZuJRWUaGG3UzMCawEAAQJATm+7xxAV66mXMH6IgrUN
NwNVQDOJLiFd2YXcTEQpd3WLHybasJTqHiz5/yt5oqnu7q/YIToAn/Kx1s3O4PQ9
YQIhAPfWLckHhmky+q71HVc1CGHqerIEtcDABeedDxVzKcv3A1EA72V0qMUymN8z
1JIEAvHDQYw1aXAVU20e9EeVBk7Gy6UCIQCFVsrIrk23pv+KvCRuz7OUygs8XqBR
lH0YwDxvM8fWvQIge8bMO5rCH3KAXO8znzgbSZ/R2fM+w4TPFWRPRZSjipUCIQDh
tVZr2umDhxTKnANGX8SFkFIIMtute7R8JFjsKUMkFCQ==
-----END RSA PRIVATE KEY-----
```

Note: Make sure you give the valid recovery key.

[← Go Back](#) [→ Change Pass Phrase](#)

Figure 7.9(u) - Resetting Master Pass Phrase

- Reset volume pass phrase: Using this option, user can reset the Volume Pass Phrase. Each encrypted Volume will have a corresponding recovery key, which will have to be entered when that particular volume's pass phrase is reset.

Note:

- The new pass phrase has to be of minimum 6 characters.
- This action is available only if the user has chosen the option to be prompted for generating a key file to encrypt the volume when the Volume Management section was accessed for the first time.

Reset Pass Phrase

New Pass phrase:

Confirm Pass phrase:

Pool:

Encrypted Volume:

Please paste the unencrypted recovery key here:

```

-----BEGIN RSA PRIVATE KEY-----
MIIBOwIBAAJBANsIRhbF4RYkgQhdNHVeH3ASJzDCK8KQ381SOKVbHjtoXMyNQIIU
/c7sho4mNG03t1eQniLAY72i+d2kuFBNg00CAwEAAQJBAIkNtE2rpBxm2xaQu7J1
ItbNhrRhNZ3RnkqA2gpt1WIVy8LFOh2tHrOCKY8dSxFKUIWbj/kMsVOCPCf3xVEUt
lgECIQD0DHZ+C2R0rzddo/1WXQXuJpu6iJX0Y2H/hXJsqT3LBQIhAOXCM08M6+Nc
vtI5cQPJH/jOoMvJc2vGs8cGBaJpc9LJAiASuxo1Fm6+c0hgkZ1DeaRI4QG20WDg
/ve4it97Y4nwmQIhAMBO/ypvJ3bz+Av+w9VI9zg1VXz+hkucYWFYgQWgC1pJAiBq
REFiAC1UzJB9DQxNRIMQ3Z15QSmurkrVGv6yT9Z5dw==
-----END RSA PRIVATE KEY-----

```

Note: Make sure you give the valid recovery key.

Figure 7.9(v) - Resetting Volume Pass Phrase

- Upload old master key: This option will be available when the Master key file for the encrypted volumes is missing, e.g after reinstallation of OS. The user can upload the Master key and unlock all the encrypted volumes, including those which were encrypted with specific volume keys.

Upload old master Recovery Key

This wizard unlocks all the old encrypted volumes, and re-encrypt with current master key.

Current Master pass phrase:

Please paste the unencrypted recovery key here:

```

-----BEGIN RSA PRIVATE KEY-----
MIIBOgIBAAJBAALmBI8hA1lpCPr9mchnfeC6gnw5L6mxykNwjWsmHtdkfXV02DTC
wSs3FuYzXHNp4VfJHWJHGOWSMY1sK/E1H1sCawEAAQJACaoObHxLftWzk9rjjAB9
yPuKpwJ68+hFEAKRvTabV81pVvZR8t2tPHdTeXNwRyQHob8W+6o+t2tt2tvzR/c
AQIhAOTJgqAlpMuf1f4btqicIH3j7ZHnC9nzRmISH6GEZntcbAIEAz7rz2m4uL83T
dyjCLzAZIxYpDgd61Ru/5ntXPPqSOQECICGK1WP84GwZ0VvXpIfDm9/Y8pNZVBt3
1vG5xn1fd4nRAiBeOK/6P/OXH0bRfS1Ehg6UvzrP0cJP16K5Lgem58FWuAQIhAMWw
iCeIYN7JcUCjIdVhFd1ypBMwt/1H91KGmIDbgR3f
-----END RSA PRIVATE KEY-----

```

Note: Make sure you give the valid recovery key.

[➔ Next](#)

Figure 7.9(w) - Uploading old master key to unlock encrypted volumes

- Upload master key for new volumes: This option will be enabled only when new encrypted volumes/disks from other USS appliance are added to this USS appliance.

Note: Every time user wants to move encrypted volumes/disks from a USS appliance to another (called pool migration), they need to provide the respective master key.

Note: After successful pool migration & master key upload, you will be asked to provide volume pass phrase to activate encrypted volumes/snapshots. In this case, provide the master pass phrase of the appliance to which the pool was migrated.

2.2. Snapshot Scheduling Settings

The schedule timing for hourly, daily and weekly snapshots can be configured here globally.

2.3. Scheduled Snapshots

You can view details such as snapshot type, preserve limit, volume name and pool name of volumes for which snapshots have been scheduled.

7.10 Replication

Data can be replicated in synchronous and asynchronous manner. In synchronous mirroring file system on the active node is notified that the writing of the block is finished only when the block has made it to both disks of the replication pair. Synchronous mirroring is the right choice for mission critical environments where we should not lose a single transaction in case of the complete crash of the active primary node. Since Synchronous mirroring involves same amount of data being transferred and IO to remote host as to the local host compressing the data can save the amount of data sent over the network.

Compression squeezes the IO to a smaller value hence saving the network bandwidth. The data received on the remote host is decompressed and written to local disk.

- ▶ Continuous replication of block level data across a TCP/IP network.
- ▶ Synchronization of changed data after an interruption in replication
- ▶ Read and write access on one of the 2 systems to the block device (designated primary system)
- ▶ Compression of I/O data

Note: To know about Replication dependencies, please refer to the **Dependencies** section under **Licensing**.

- **Sections of the interface**

1. Summary

This section displays brief summary information regarding:

- Number of Primary/Source volumes on the system.
- Number of Secondary/Peer volumes on the system.
- Mis-configured volumes if any.
- TCP Port range that would be used for replication.
- Listening port number of the peer system.

Summary:	
Primary / Source volumes:	1
Secondary / Mirror volumes:	0
Mis-configured volumes:	0
Replication TCP incoming port range:	7700 - 7956
Peer configuration service listening port:	9000

Figure 7.10(a) - Volume Replication Summary

2. Add replication configuration

Replication configuration can be created from this section. Below mentioned notes should be taken care of before creating replication.

- Replication can be created only on LVM/Chelsio TP volumes.
- Snapshots cannot be created on the LVM/Chelsio TP volumes that would be part of Replication.
- System from which Replication would be created will be Primary volume by default.

- Replication can only be created over Chelsio networks.

Example: Create Replication

Follow the Steps below to create Replication (Provided Pool and logical volume is created)

Important:



- a. When replicating a volume, system creates a new volume in the same pool to store metadata(33 MB of metadata for 1 TB of volume space) in both peers and part of system memory is reserved for this resource (33 MB of system memory for 1TB of volume space). This memory will be released only when the resource is deleted. Hence make sure the systems have enough free pool space and system memory for successful resource creation.
- b. When the replicated volume expands, the corresponding metadata volume size and reserved memory will also expand proportionately.

- i. Click on Replication on the Left Menu Tab.
- ii. Click on Add Replication Configuration. Screen as seen below appears.

— Add replication configuration:

- Choosing a volume that is currently in use, will disrupt I/O services using the volume, such as an iSCSI target or NAS file sharing, to configure replication. Clustered systems will not allow replication of volumes already in use.
- Replication will create a new volume for storing meta data in the selected pool, ensure that the pool has free space of atleast 33 MB per TB of volume size (Eg: a 10TB volume need 330 MB of free space in the pool).

Source / primary volume:

Peer Hostname / IP address:

Figure 7.10(b) – Add replication configuration

- iii. Select the LVM/Chelsio TP volume from the list and IP address of the system which would be its peer and click Next. Screen as seen below appears.
- iv. Select a Chelsio IP address on the primary/source system which would communicate with the Peer.
- v. Select a Chelsio IP address on the secondary/Peer system which would communicate with source system.
- vi. Select the Pool on the peer system which would be used for storage.
- vii. Choose Synchronization policy.
- viii. If Bandwidth size has to be mentioned, check the Specify bandwidth restriction box and enter the size in the Maximum bandwidth allowed box. Default value will be set which can found after replication creation in case Specify bandwidth restriction box is not checked.
- ix. To save network bandwidth while replication, check the Enable data compression box.
- x. Click on **Add replication configuration** button to create replication


- xi. Adding Replication may take some time depending on the network type, volume size etc.
- xii. Refresh the page to see the updated synchronization status.

Error messages if any during creation of Replication can be seen in the status box.


Note: Replication is considered to be in optimal state only if the Primary and Secondary devices are in sync. System would take time before they are in sync. Synchronization status can be checked by selecting the device from Replicated device.

3. Replicated devices

Devices which are still in the synchronization process would be identified as shown in the following figure.



The screenshot displays a configuration window for a replicated device. At the top, it identifies the device as 'pool_1/vol1' with a size of 5.00 GB, designated as a 'primary / source volume'. The status section indicates that the device is 'Synchronizing (Primary / Source Volume)' and that 'Synchronization complete: 11.4%'. It also shows 'Data yet to be sync'ed: (4540/5120)M and a 'Sync complete ETA: 0:00:46 hours'. Under 'Storage devices', the 'Local device' is 'Up to date / in sync', while the 'Peer device' is 'Inconsistent / needs updating'. Network portals are listed as 'Local Network Portal: 102.1.1.237:7700' and 'Remote Network Portal: 102.1.1.62:7700'. On the right side, there are five control buttons: 'Start', 'Stop', 'Reverse role', 'Delete replicated volume', and 'Delete configuration'. At the bottom, there are expandable sections for 'Configuration' and 'Synchronization schedule'.

1.  Device: pool_1/vol1 | 5.00 GB | primary / source volume

Status: Synchronizing (Primary / Source Volume)
Synchronization complete: 11.4%
Data yet to be sync'ed: (4540/5120)M
Sync complete ETA: 0:00:46 hours

Storage devices:
Local device: Up to date / in sync
Peer device: Inconsistent / needs updating

Local Network Portal: 102.1.1.237:7700
Remote Network Portal: 102.1.1.62:7700

[+ Configuration:](#)
[+ Synchronization schedule:](#)

Start
Stop
Reverse role
Delete replicated volume
Delete configuration

Figure 7.10(c) – Replicated devices in the process of synchronization

Devices which have completed synchronization would be identified as below.

The screenshot displays a configuration window for a replicated device. At the top, it identifies the device as '1. Device: pool_1/vol1 | 5.00 GB | primary / source volume'. The status is 'Connected'. Under 'Storage devices', both 'Local device' and 'Peer device' are 'Up to date / in sync'. Network portals are listed as 'Local Network Portal: 102.1.1.237:7700' and 'Remote Network Portal: 102.1.1.62:7700'. On the right, there are five control buttons: 'Start', 'Stop', 'Reverse role', 'Delete replicated volume', and 'Delete configuration'. At the bottom, there are expandable sections for 'Configuration' and 'Synchronization schedule'.

1. Device: pool_1/vol1 5.00 GB primary / source volume	
Status:	Connected
Storage devices:	
Local device:	Up to date / in sync
Peer device:	Up to date / in sync
Local Network Portal: 102.1.1.237:7700	
Remote Network Portal: 102.1.1.62:7700	
+ Configuration:	
+ Synchronization schedule:	

Figure 7.10(d) – Replicated devices with synchronization complete

3.1. Resource Actions

- Start: Local resource connects to the replication link. If the peer is also connected, then replication of data starts (depending on the roles). If peer is not connected, then the local resource waits for connection to be established.
- Stop: Local resource disconnects the replication link, thereby stopping the replication process.
- Reverse role: The Resource role is changed from Primary to Secondary and vice versa.
- Delete replicated volume: The secondary replicated volume is deleted along with the replication configuration and metadata volumes on both peers. This action has no effect on the primary data volume.
- Delete configuration: Only the replication configuration and metadata volumes on both peers will be deleted. Primary and secondary data volumes are preserved.

Note:

- In order to access the replicated data at peer side, please make sure that the resource is made primary by reversing the roles.
- The Reverse role action will work only if the Resource is not in use at the Primary side.

Example: How to use the replicated volume on the peer side

After creating replication and synchronization is complete, follow these steps to access the replicated data on the peer side:

- i. On the Local Machine (where the replicated volume resides), make sure that the volume is free/unassigned.
- ii. On the Peer side, in the **Replication** module, expand the **Replicated devices** section.
- iii. Each replicated volume will be listed with its related status and actions.

- iv. To access a volume, click the **Reverse role** button and click **OK** on the alert box that appears.
- v. After the process of reversing the roles is complete, the status of the replicated volume should change and display *primary / source volume*.
- vi. Now, in the **Volume management** module, select the pool (specified while creating replication on the local side) from the list.

The volume can now be used in any iSCSI,FC and NAS operation.

3.2. Resource Settings

3.2.1. Configuration

Here user can modify local/ peer IP address, port number, synchronization policy, bandwidth limit and enable/disable data compression for a resource. By default, configuration settings selected while creating the resource will be displayed.

Note: Before modifying the IP address, please make sure that the new IP address is already configured.

Configuration:

Local IP address: 104.44.44.145

Local TCP port: 7711 (7700-7956)

Peer IP address: 104.44.44.225

Peer TCP port: 7711 (7700-7956)

Synchronization Policy: Disk synchronous mirroring ▼

Maximum bandwidth allowed: 1 GB/s ▼

Enable data compression:

Save settings

Figure 7.10(e) – Modifying resource configuration

3.2.2. Synchronization Schedule

Here the replication process for a resource can be scheduled. Available options are Continuous and Time slot.

- Continuous: Data gets replicated in real time.
- Time slot: User can specify the exact time at which primary resource will establish connection; therefore start replicating data (provided the peer is connected).

Note: If Wait till Synchronization completes field is checked, replication will continue until completion, thus bypassing the parameters specified in the Time slot option.

Synchronization schedule:

Replication Schedule: Continuous Time slot

Days of week: All Days ▾

Start time: 00 ▾ : 00 ▾ AM ▾

End time: 05 ▾ : 00 ▾ AM ▾

Wait till Synchronization completes:


 Save settings

Figure 7.10(f) – Synchronization schedule settings

7.11 File Systems

7.11.1 File Systems Summary

This section lists the various devices and the corresponding folders on which they are mounted (attached).

The screenshot shows a web-based interface for managing file systems. At the top, it says "Select a folder from the list to view details:". Below this is a list of file systems under the heading "File systems". The list includes: swap, /var/crash (highlighted in blue), / [root], and /var/log. Each item has a checkmark icon to its left. Below the list, the details for the selected folder, /var/crash, are displayed. The details include: Folder: /var/crash, Device: sdh7, File system type: ext3, Status: attached (with a red square icon), and two buttons: "Temporarily detach" (with a red square icon) and "Permanently detach" (with a red X icon). The status section also shows storage information: Free space: 5.479 GB [97%], Used space: 146.108 MB [3%], and Total space: 5.926 GB. The "Attach:" section has a dropdown menu set to "At system startup" and a green checkmark icon. At the bottom, there is an "Edit Quotas" button with a pencil icon.

Figure 7.11.1(a) - File Systems summary for a selected Folder

- **Sections of the interface**

1. Summary

The section on the top displays the list of available Folders while the section on the bottom displays the details of the selected Folder.

1.1. Summary Commands

- **Temporarily/Permanently Detach:** You can choose to detach a device from a folder in two ways: Temporarily or Permanently. Detaching a device permanently results in the specific folder being removed from the list on the top; whereas the specific folder appears in the list but in a disabled state when detached temporarily.

2. Default Quota Settings

In this section, you can set disk usage limits for users/group-users with non-root privileges. Limits can be set as percentage of the total disk size or in KB, MB, GB or TB. Two kinds of limits can be set: Block Hard Limit and Block Soft Limit. Email alert will be sent (if activated), if users cross the lower limit (Block Soft Limit). The Block hard limit is the maximum size of the volume allocated to users. The default values, if not changed, will apply for all the volumes. Once

changed, the new settings will apply only to those volumes which are XFS formatted/re-formatted after the change.

Note:

- The percentage limits will not dynamically change, when the volume (including filesystem) is resized
- When the volume size is less than the block hard/soft limit specified, the hard and soft limit is set to 50% and 25% of volume size respectively by default.
- If Block Soft Limit provided is greater than Hard Limit, then 50% of Hard Limit will be applied as Soft Limit. However this is only possible when Hard Limit is less than the volume size.

Default Quota Settings

Block Hard Limit	<input type="text" value="15"/>	<input data-bbox="894 253 1131 291" type="text" value="% of volume size"/>
Block Soft Limit	<input type="text" value="10"/>	<input data-bbox="894 291 1131 327" type="text" value=""/>
	<input type="button" value="Apply"/>	<ul style="list-style-type: none">% of volume sizeKBMBGBTB

Figure 7.11.1(b) – Default Quota Settings

7.11.2 File system Quotas

You can specify maximum disk space and other privileges for a User/Group on the selected device in this section.

- **Sections of the interface**

- 1. Summary**


This section displays the File system path, Device, Free space, Used space and Total Space and Quotas status.

Summary:

File system path:	/vol1
Device:	pool1/vol1
Free space:	107.317 GB [99%]
Used space:	4.760 MB [1%]
Total space:	107.322 GB
Quotas status:	User accounting enabled User limits enforcement enabled Group accounting enabled Group limits enforcement enabled

Grace times for all users and groups

Space grace time:	<input type="text" value="7"/>	days ▼
Files grace time:	<input type="text" value="7"/>	days ▼

 **Apply**

Note: grace time periods will be rounded off to the nearest value and reported.

Figure 7.11.2(a) - File system summary with quotas status and Grace time

1.1. File Systems Quotas Summary actions

- Grace time for all users and groups: Space/Files grace time is the time for which a User can continue accessing the remaining of allocated disk space after exceeding the Space/Files soft limit. A Global grace time applicable to all Users and Groups can be set here in terms of minutes, hours and days

Note: Setting up a new grace period will not affect Users/Groups which are already in the grace period.

Grace times for all users and groups

Space grace time:

Files grace time:

Figure 7.11.2(b) - Grace time for all users and groups

2. User quotas

2.1. Quota allocation and usage

Here you can view various statistics related to the disk space allocated to different users.

- User: displays a list of current users for the specified Device for whom quota has been set.
- Current space usage: Disk space used by the User.
- Space soft limit: Maximum disk space allocated to a User exceeding which email alerts regarding quota violation are sent. The Space grace period starts at this point.
- Space hard limit: Total disk space allocated to a User.
- Space grace period: Time for which a User can continue accessing the remaining of allocated disk space (only up to Hard limit) after exceeding the Space soft limit. After the Space grace period ends the User is no longer allowed to write on the disk above soft limit.
- Current usage (files) :The number of files created by the user on the specified volume.
- Soft limit (files): Number of files the user is allowed to create after which the File grace period starts.
- Hard limit (files): Maximum number of files the user is allowed to create on the specified volume.

- Files grace period: Time for which a user can continue creating files (only up to Hard limit) after exceeding the Files Soft limit. After the Files grace period ends the User is no longer allowed to create any more files.

User quotas:

Quota allocation and usage:

User	Current space usage size [% of total space]	Space soft limit	Space hard limit	Space grace period	Current usage (files)	Soft limit (files)	Hard limit (files)	Files grace period
aaa	2.439 GB[2%]	2.000 GB	2.500 GB	6days	6	5	8	6days
root	0 KB[0%]	0 (no limit)	0 (no limit)	-	3	no limit	no limit	-

Figure 7.11.2(c) - User quotas allocation and usage table


2.2. Set quota allocation

Quota can be added here for Local user or for a Domain user using NIS/ADS server. Furthermore, Space and Files soft limit, and Space and Files hard limit can be set for the specified User.

Set quota allocation:

User: **Local user**
 ▾

Domain user



Space soft limit: ▾

Space hard limit: ▾

Files soft limit:

Files hard limit:


 **Apply**

Figure 7.11.2(d) - Set quota allocation

3. Groups quotas

3.1. Quota allocation and usage

Here you can view various statistics related to the disk space allocated to different Groups.

- **Group:** displays a list of current Groups for the specified Device for which quota has been set.
- **Current space usage:** Disk space used by the Group.
- **Space soft limit:** Maximum disk space allocated to a Group exceeding which email alerts regarding quota violation are sent. Grace period starts at this point.
- **Space hard limit:** Total disk space allocated to a Group.
- **Space grace period:** Time for which a Group can continue accessing the remaining of allocated disk space (only up to Hard limit) after exceeding the Space soft limit. After the Space grace period ends the Group is no longer allowed to write on the disk above soft limit.
- **Current usage (files):** The number of files created by the Group on the specified volume.
- **Soft limit (files):** Number of files the Group is allowed to create after which the File grace period starts.

- Hard limit (files): Maximum number of files the Group is allowed to create on the specified volume.
- Files grace period: Time for which a Group can continue creating files (only up to Hard limit) after exceeding the Files Soft limit. After the Files grace period ends the Group is no longer allowed to create any more files.

Group quotas:

Quota allocation: Groups

Group	Current space usage size [% of total space]	Space soft limit	Space hard limit	Space grace period	Current usage (files)	Soft limit (files)	Hard limit (files)	Files grace period
ggg	4.471 GB[4%]	3.000 GB	5.000 GB	6days	11	10	15	6days
root	4 KB[0%]	0 (no limit)	0 (no limit)	-	3	no limit	no limit	-
users	2.439 GB[2%]	0 (no limit)	0 (no limit)	-	6	no limit	no limit	-

Figure 7.11.2(e) - Group quotas allocation and usage table

3.2. Set quota allocation

Group quotas can be added here with other options like Space and Files soft limit and Space and Files hard limit.

Set quota allocation: Group

Group: Local group
users ▼

Domain group
[Empty text box] [Search icon]

Space soft limit: 155 GB ▼

Space hard limit: 200 GB ▼

Files soft limit: 1500

Files hard limit: 2800

Figure 7.11.2(f) - Group quotas allocation and usage table

Example: Setting quota allocation for users

Follow the steps mentioned below to set quota allocation for local and domain users for a folder.

- i. In the **File systems** list, select the folder for which you want to set quota allocation and click the **Edit Quotas** button. This will navigate to a new page.
- ii. Expand the **User quotas** section.
- iii. Add a local user or a domain user (using NIS or ADS server). To add a local user, enable the **Local user** radio button and select a user from the drop-down. To select a domain user, enable the **Domain user** and the click on the **Search** icon. Now specify the domain (ADS, NIS) and enter the user name. Click **Apply**.
- iv. Specify the Space soft/hard limit and Files soft/hard limit and click **Apply**.

Note: Similar procedure can also be used to set quota allocation for groups in the **Group quotas** section.

7.12 Lustre

The Lustre file system is a scalable, secure, robust, and highly-available cluster file system that addresses the I/O needs, such as low latency and extreme performance, of large computing clusters.

- **Lustre Clusters**

Lustre clusters contain three kinds of systems:

- File system clients, which can be used to access the file system
- Object storage servers (OSS), which provide file I/O service
- Metadata servers (MDS), which manage the names and directories in the file system

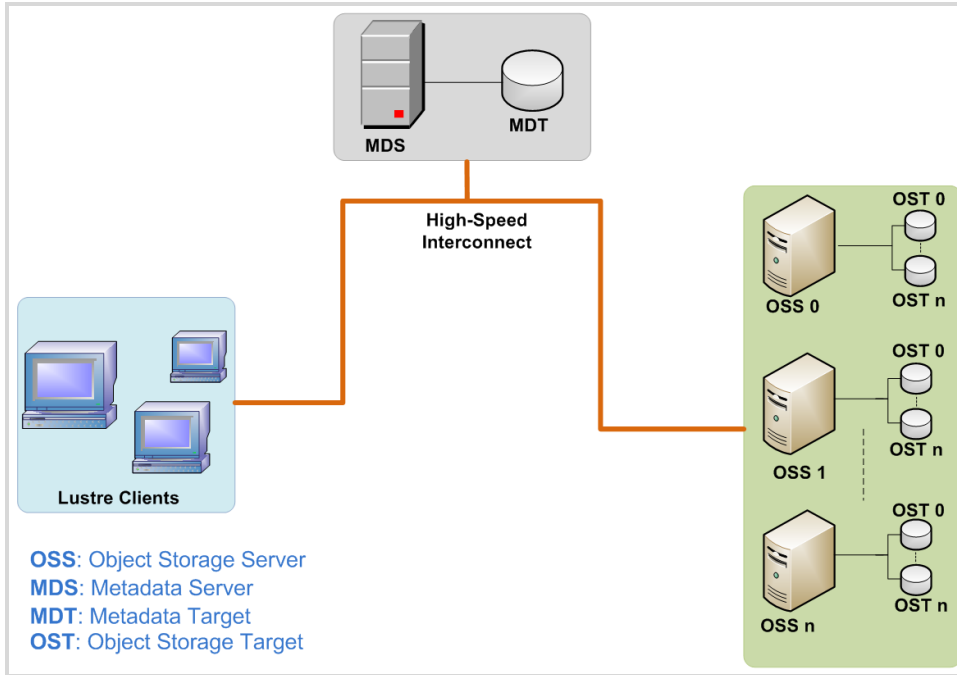


Figure 7.12(a) - Systems in a Lustre cluster

Note: Support to Lustre is provided on two Network types: Infiniband and TCP/IP on Ethernet. To know more about Lustre support, please refer to the **Dependencies** section under **Licensing**.

Example: Creating Lustre files system

- i. In the **Network** module, expand the **Devices** section.
- ii. Select a Chelsio network adapter, and expand the configuration option.
- iii. In the Lustre networking drop-down, select *Lustre TCP* (for T3 adapters) or *Lustre TCP/RDMA* (for T4 adapters) and click **Apply**.
- iv. In the **Volume management** module, select a pool from the Pool list and create a TP volume using the *Allocate new Space/Logical Volume* option. Specify the size and name. You can also use an encrypted volume.
- v. Click **Apply**. The newly created volume will be listed in the **Logical Volumes** section.
- vi. Click on the newly created volume and in the **Actions** drop-down select *Format and use Volume*.
- vii. Use the **Browse** button to locate the folder to mount the volume on.

- viii. For the **Volume Usage** option select **Lustre**.
- ix. Choose the **Lustre target type** as *MGS*.
- x. Click **Apply**
- xi. Create another TP volume by specifying the size and name.
- xii. Follow steps (vi)-(viii)
- xiii. Choose the **Lustre target type** as *MDT*.
- xiv. Specify a name for the file system and the IP Address/hostname of the system where MGS is configured. For HA mode, specify IP Address/hostname of the secondary node where MGS is configured.
- xv. Now select the Chelsio interface on which Lustre networking is to be configured from the drop-down list.
- xvi. Click **Apply**
- xvii. Follow steps (vi)-(viii)

xviii. Choose the **Lustre target type** as *OST*.

xix. Follow steps (xiv)-(xvi)

Lustre file system is configured successfully.

Note: Although, you can create OST on the same server as MGS/MDT, it is recommended that you configure OST on a different USS appliance.

7.13 Backup (non-HA)

Expanding servers, exponential data growth and the need for 24/7 data availability requires an organization to have a reliable backup and recovery solution. The danger of losing mission-critical data due to natural or manmade disasters is of utmost concern for all kinds of businesses- small, medium or large.

USS provides a robust, secure and reliable storage backup and recovery solution for all your enterprise-level data. You can backup important data using the traditional magnetic tape cartridges or virtual tape library. The intuitive wizard helps IT administrators to configure backup, verify and restore **Jobs**. Data can be automatically backed up and verified for integrity on a scheduled basis without user intervention and thus saving administrator's time considerably.

7.13.1 Backup Wizard

- **Add a storage device**

The Backup Wizard helps you in getting started with configuring data backup on magnetic tape devices. You can add a storage device (tape cartridge or create a virtual tape drive), add a backup job or add a tape/volume to the device.

Follow the steps mentioned below to add a tape cartridge as new storage device:

1. Select Device type as "Tape Drive"
2. Select the tape cartridge from the "Tape device" drop-down.
3. You can specify the type of tape media for a physical tape drive. Please ensure that you use an accurate name, such as DDS-5, for the type field.
4. Specify the name of the tape drive and click "Apply"

Follow the steps mentioned below to add a virtual tape drive as new storage device:

1. To add a virtual tape drive, you can either use a logical volume from a storage pool or specify a folder to use directly. To use the first option, select a storage pool from the "Storage Pool for new volume" drop-down. Specify a name for the logical volume or leave it at its suggested default value.
2. Specify the size of the volume that will be created.
3. Specify the folder, to which the newly created volume will be attached/mounted to.
4. Specify the name of the virtual tape drive.
5. Click "Next "

Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as the device type. .

Device type:

Create new virtual tape drive by:

Storage Pool for new volume:

New volume **name**:

New volume **size**:

Folder to use:

Tape drive name:

Figure 7.13.1 (a) - Adding a virtual tape drive.

- **Add a Job**

You can also use the backup Wizard to create a backup job. Follow the steps mentioned below.

1. Select Backup data in the Job type drop-down.
2. Provide a name for the backup job.
3. A default file group template, "System configuration data", pre-configured with folders to backup and other related settings is available. Select it from the drop-down.
4. Similarly, a pre-configured backup schedule, called Schedule template is available. You can choose between two available options: monthly and weekly.
5. Choose the storage device created in the "Add a storage device" section.
6. Choose Default as the media pool.
7. Click "Apply".


Job type:	Backup data ▼
New job name:	backup_job1
Template name:	System configuration data ▼
Schedule :	Monthly cycle ▼
Storage device:	VirtualTapeDrive1 ▼
Media pools:	Default ▼
	

Figure 7.13.1 (b) - Add a backup job.

- **Add a tape**

Next, you can add a Tape / volume to the storage device (virtual or a physical tape drive). Specify the name of the tape/volume and Click Apply.

Device: /virtual_tape

Storage service aware: **yes**

Type: Virtual Tape

Storage device: Pool: pool1, volume: VirtualTapeDrive1

Space: Total: 48.804 GB, used: 4.531 MB, free: 48.800 GB

Name: VirtualTapeDrive1

Actions:

Tape / volume name:

Media pool:

Figure 7.13.1 (c) - Adding a tape/volume to storage device

7.13.2 Backup and Restore




- **Sections of the interface:**

1. Summary

This section shows a brief summary of the backup configuration. User can choose to Start/Stop the Backup manager service. The service can be configured to run automatically by clicking on the Enable Auto Start button.

Summary:

Services status: Backup manager service is **running**
Storage service is **running**
Client service is **running**

Actions:   

Backup jobs configured:	1
Restore jobs Configured:	0
Currently running jobs:	0
Scheduled jobs for next 3 days:	1
Jobs requiring attention:	0
Storage devices:	1
File group templates:	1
Schedule templates:	2
Media pools:	1

Figure 7.13.2 (a) - Summary section of the backup page

2. Jobs

You can view the list of configured jobs here, or modify the job if required.

The screenshot displays a web-based interface for managing backup jobs. At the top left, there is a red minus sign icon followed by the text 'Jobs:'. Below this, the interface is divided into two main sections: 'Jobs:' and 'Details:'.
The 'Jobs:' section contains a list box with the following items: '--Backup jobs--', 'backup1', 'backup2', '--Verify jobs--', 'verify_job1', '--Restore jobs--', and 'restore_job1'. The 'backup1' item is highlighted in blue.
The 'Details:' section provides information for the selected job 'backup1':
- Job name: backup1
- Status: **Scheduling enabled**, not running
- next run scheduled at **23-Apr-13 23:05**
- File group selected: System configuration data
- Backup type: ---Select a Level--- (dropdown menu)
- Schedule: Monthly cycle (dropdown menu)
- Storage device: VirtualTapeDrive1 (dropdown menu)
- Media pool: Default (dropdown menu)
Below the details are three buttons: 'Disable' (with a red X icon), 'Save Changes' (with a green checkmark icon), and 'Save and Run' (with a green play button icon).
The bottom section is titled 'Previous runs of selected job:' and contains a list box with one entry: '2013-04-22 23:05:04'. To the right of this list box are the following details for the selected run:
- Start time: 2013-04-22 23:05:04
- End time: 2013-04-22 23:05:06
- Final Status: **Completed successfully**
- Files backed up: 1,843
- Backup type: Full

Figure 7.13.2 (b) - Configured jobs with details

3. Add a backup/restore Job

Add a backup, restore or verify job here, after you have configured a storage device.

System Configuration data is the default file group template and Monthly and Weekly are the two default Schedule templates available. You can either use these templates with their default settings or edit them in the Templates section. You can also create your own templates. If you want to create a job on a new storage device and media pool, add them in the Storage media page first.

Note: Only one Restore Job can be added and edited from the Jobs section.

Add a backup/restore job:

Job type:	Backup data
New job name:	backup_job1
Template name:	System configuration data
Schedule :	Monthly cycle
Storage device:	VirtualTapeDrive1
Media pools:	Default


 Apply

Figure 7.13.2 (c) - Add a backup job.

Add a backup/restore job:


Job type: Restore data from a backup ▾

New job name: restore_job1

Template name: System configuration data ▾

Storage device: VirtualTapeDrive1 ▾

Media pools: Default ▾

Where to Restore: /restore  Browse..


 Apply

Figure 7.13.2 (d) - Add a restore job

— Add a backup/restore job:

Job type:

New job name:

Template name:

Schedule:

Storage device:

Media pools:

Figure 7.13.2 (e) - Add a verify backup data job

4. Templates

The backup configuration is dependent on templates, for choosing which folders or “file groups” to backup (called file group templates), and the schedule to use for the backup job (called schedule templates).

Default common templates are provided as an example. You can configure existing default templates here for file groups and schedules or create new ones.

1.1. File group templates

1. Choose the template to view and modify, or choose “+Add a File group template” to add a new template.
2. The template name can be modified, and is required when creating a new template.
3. Choose the list of folders that you wish to backup. The “Add” command will display the folder selection pop-up, allowing you to browse the system to find the folders to backup.
4. You can expand the “Settings” link to define additional advanced settings if required.
5. Click “Save template” to save, or you can delete the template if you wish to.

Note: fileset will be disabled if the corresponding job is running.

File group templates: **Details**

+ Add a file group template

- System configuration data template02

Template name:

Folders to backup:

Full path:

+ Settings:

Figure 7.13.2 (f) - Add or modify "File group" templates.

1.1. Schedule templates

1. Choose the template to view and modify, or choose “+Add a Schedule template” to add a template.
2. You can have a list of repetitive schedules, and different types of backup for each schedule, to achieve your backup strategy. There are 2 sample schedules defined, which do different types of backups, based on the day of the week, and the week in the month.
3. Set the schedule settings for an existing schedule by selecting it, or click “Add” to set the settings for a new schedule.
4. The scheduling defaults to every hour of every day of every week of every month in the year. This is not advisable. Please ensure that the time of day to run the backup is set correctly.

Schedule templates:
Details

+ Add a Schedule template

Monthly cycle

Weekly cycle

Template name:

Schedule:

1. Run Full backup
2. Run Differential backup
3. Run Incremental backup

Schedule **Run Full backup at 23:05 on sunday and on first week of the month**

Description: **Run Full backup at 23:05 on sunday and on first week of the month**

Schedule Settings:

Backup type:

Scheduling mode:

Week days:
 Mo
 Tu
 We
 Th
 Fr
 Sa
 Su

Figure 7.13.2 (g) - Add or modify schedule templates



Warnings:

Do not delete a template that is used by a job, because the backup manager will stop functioning due to the invalid configuration.

7.13.3 Storage devices and media

- **Sections of the interface:**

1. Storage devices

This section displays the configured list of backup storage devices, and allows you to add a storage device. Click on a configured storage device to view details for the device. For virtual tape drives, you can create a Tape / volume to be used in the virtual tape drive. For physical tape drives, you can view status of the drive, and attach / detach from the backup manager service, and add physical tape cartridges to be used by the backup service.

Tape drives:	Details
<ul style="list-style-type: none">+ Add a tape driveIBM - DDS Gen5 - DDS-5	<p>Device type: <input type="text" value="Virtual Tape drive"/></p> <p>Create new virtual tape drive by: <input type="text" value="Allocating space from a storage pool"/></p> <p>Storage Pool for new volume: <input type="text" value="---Select a Storage Pool---"/></p> <p>New volume name: <input type="text" value="VirtualTapeDrive1"/></p> <p>New volume size: <input type="text"/> <input type="text" value="GB"/></p> <p>Folder to use: <input type="text"/> <input type="button" value="Browse.."/></p> <p>Tape drive name: <input type="text" value="VirtualTapeDrive1"/></p> <p><input type="button" value="Apply"/></p>

Figure 7.13.3 (a) - Add or modify schedule templates

- To add a new storage device, click “+Add a tape drive”. Now you can choose the type of tape drive to add either a virtual drive or a physical tape drive. For a virtual tape drive, you may choose a Volume management storage pool to allocate disk space, and create

a volume, attach it to a folder, which will be used as the virtual tape drive. Alternately, you can specify a folder to use directly.

- Ensure that you specify the name of the device correctly as required, or leave it at its suggested default value.
- You can specify the type of tape media for a physical tape drive. Please ensure that you use an accurate name, such as DDS-5, for the type field.

1.1. Storage device actions

- **Attach tape drive to backup service**
A new tape drive can be attached to the backup service, where the backup will be taken. Backup service will make use of the next drive only when the current one is full.
- **Rewind/Eject/Eraser the tape**
These actions are used to Rewind/Eject/Eraser the tape in the tape drive. These actions are not available for virtual tape drives.
- **Add new tape to backup service**
This option is used to add a new tape to the backup service when the status of the previous tape is full .

Tape drives: **Details**

+ Add a tape drive
IBM - DDS Gen5 - DDS-5

Device type:

Create new virtual tape drive by:

Storage Pool for new volume:

New volume **name**:

New volume **size**:

Folder to use:

Tape drive name:

Figure 7.13.3 (b) - Add or modify schedule templates

Storage devices:

Tape drives:	Details
<ul style="list-style-type: none">+ Add a tape driveIBM - DDS Gen5 - DDS-5	<p>Device: IBM - DDS Gen5 (Connected)</p> <p>Storage service aware: yes</p> <p>Type: DDS-5</p> <p>Version: COA0</p> <p>Serial: JR00L80</p> <p>Tape media: Loaded</p> <p>TapeAlert monitor: Supported, current status is okay</p> <p>Error correction details:</p> <p>Total errors corrected: reads: 0 writes: 0</p> <p>Data processed (GB): reads: 0.001 writes: 0.001</p> <p>Total uncorrected errors: reads: 0 writes: 0</p> <p>Storage device: /dev/nst0</p> <p>Name: IBM - DDS Gen5</p> <p>Actions: <input type="text" value="---Select an action---"/> <input checked="" type="checkbox"/></p>

Figure 7.13.3 (c) - Add or modify schedule templates

2. Media pools and tapes

User can view the list of configured media pools, and tapes / volumes that are currently configured in a pool. Click on a pool from the Media Pools list to view/edit the corresponding details and tapes added to the pool. To view the status of the tape in that pool, click on the tape name in the 'Tapes cartridges / volumes' list. New media pools can be added from this section.

2.1. Add a media pool

- To add a new media pool, click “+Add a media pool”.
- Ensure that you specify the name of the device correctly as required.
- If the 'Recycle tapes' option is enabled, the tape will be reused after the time specified in the 'Recycle tapes every' option.

Media pools and tapes:

Media Pools:

- + Add a Media Pool
- Default
- pool1

Details:

Pool name:

Recycle tapes: Enabled

Recycle tapes every: days



Figure 7.13.3 (d) - Add or modify schedule templates

Media pools and tapes:

Media Pools:


+ Add a Media Pool
Default
pool1

Details:

Pool name: Default

Recycle tapes: Enabled

Recycle tapes every: 365 days

 Apply

Tapes cartridges / volumes:

tap1

Details:

Tape volume name: tap1

Media type: VirtualTape1

Status: Append

Added on: 2011-03-15 18:20:49

Actions: Erase records from tape

Tape can be erased only when tape status is append/full/used


 Apply

Figure 7.13.3 (e) - Add or modify schedule templates

2.2. Media tape actions

- Erase records from tape

This option is used to erase the contents of the tape.

Note: This option is enabled only when tape status is append/full/used.

8. iSCSI SAN



8.1 iSCSI overview

iSCSI is designed for sharing block storage over TCP/IP networks. This layering allows for deployment of iSCSI storage over long distances or within the data-center over commonly available Ethernet networks. The bandwidth and latency of iSCSI storage depends on the network adapters, system processor and memory capabilities, initiator / client, and target / server used.

Chelsio's iSCSI Target is a stable and high performance product, which enables sharing very large devices (greater than 2 TB) as LUNs / disks to initiators. Multi-path I/O (MPIO) for redundant network topologies and multiple connections per session (MCS) for utilizing high-bandwidth links are supported. It also supports Microsoft Cluster Service Nodes using Microsoft iSCSI initiators, for shared storage, through the SCSI Persistent Reserve or Reserve-Release mechanism. Combined with Chelsio Storage Acceleration hardware, the iSCSI target consumes very low system resources, while providing exceptional performance.

The iSCSI target requires a valid license key, issued by the vendor of the appliance. Once the license key is installed, the target can be configured and started.

8.2 iSCSI deployment

SAN topology design is an essential preparatory step towards a successful deployment. Some of the basic tasks are summarized as follows:

a) Requirements gathering

This includes collecting all current and future expansion / scalability requirements. It also involves estimating the growth of Storage data, and growth of bandwidth usage.

b) Product evaluation and selection

The operating systems and hardware platforms should be already formalized by the first step. This leads to a selection of initiators to use on the particular platform(s) being used. Cost estimation and budgeting is required in case of specialized hardware or services to be used, such as high bandwidth network switches, or a high-capacity WAN link, or a protocol acceleration adapter.

c) Testing and deployment

The final step is to ensure that the products interoperate well, delivering stability, data integrity, and achieving key metrics and SLAs. The test environment can validate the design and product selection, ensuring that the actual production deployment is smooth and successful.

8.3 Configuring the Chelsio iSCSI Target

The Management Interface has an iSCSI section, which is divided into four main sections: Target service summary, Targets, Create a new target and iSNS (Internet Storage Naming Service).

Target service summary displays target service details and allows specifying the acceleration mode of iSCSI, if Chelsio Storage Accelerator hardware is installed in the system. You can also start, stop or restart the service.

The target stack allows for virtualizing iSCSI target services. Multiple targets can be configured based on deployment strategy, or usage model, or for administrative grouping / classification.

There is no limit in the stack, for the number of virtual iSCSI Targets that can be configured or the number of client connections or the number of LUNs shared. It is limited only by the system memory available.

A target is a discrete collection of the following resources on the appliance:

- ▶ The network devices / IP addresses + TCP ports on which iSCSI service is provided
- ▶ Storage devices such as disk drives / logical volumes / RAID arrays which are shared to client initiators
- ▶ Access control rules and authentication parameters.

Each target is identified on the iSCSI SAN by its unique IQN node name. A friendly name or Target Alias is also used. (Please refer to the IETF iSCSI Standard, RFC 3720, for further details, at www.ietf.org.) A minimum of one network device / IP address + TCP port combination is required for a target. Also, one storage device is required to configure a target. The 'Create target' subsection of the Targets section allows for defining a new iSCSI target with the above settings.

If you wish to test the network throughput of the iSCSI connection, choose the built-in Ramdisk LUN. If you use a performance test tool (e.g. IOmeter) that does not care about the data being sent across, use the NULLRW Ramdisk LUN, which discards data, thus avoiding a data-copy and reducing the bottlenecks in the testing process.

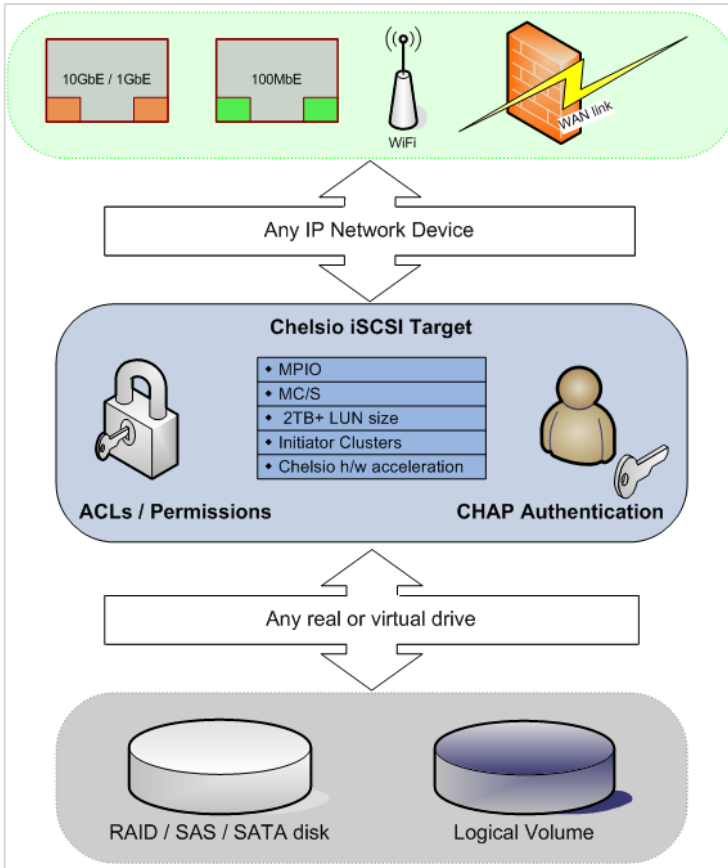


Figure 8.3 (a) – Chelsio iSCSI Target.

This diagram describes the iSCSI Target stack, its relationship with other subsystems in the Appliance, and a brief list of the features that it provides.

Multi-path I/O

This feature allows for an initiator to establish multiple sessions to the target, using multiple network paths. This may be a combination of either using different IP addresses or TCP ports with the same hardware, or using completely redundant switch and network adapter connections and cabling. Thus, if a cable is unplugged, or a hardware fault occurs, data flow between the initiator and target continues uninterrupted. The Chelsio iSCSI Target stack supports MPIO.

Client Operating systems such as Windows and Linux distributions have a built-in MPIO service, which can be configured to use all available paths to an iSCSI drive, once the Initiator establishes the session / connection to the Target using all the available Network paths. Refer the Operating system or Initiator software manual to configure MPIO on the client.

An example configuration would be, to have the Target and Initiator configured for 2 IP addresses, one each on a different Network device, and the Initiator connects to the target on both the IP networks available. These 2 IP networks are physically connected over different Ethernet networks. Refer the following diagram in 8.3 (b) for details.




Figure 8.3 (b) – *Multi-path IO topology.*

This diagram shows an example topology which can achieve hardware and software fault-tolerance at the Network layer.

Depending on the MPIO policy used, the bandwidth of the 2 paths can be aggregated too, to increase throughput.

Multiple connections per session is a mechanism of using multiple TCP connections, within a logical grouping of one session between the initiator and target. Usage of multiple TCP connections allows for properly utilizing high bandwidth 1GbE or 10GbE links. For example, the Microsoft iSCSI initiator supports up to four connections in a session. The Chelsio iSCSI Target stack supports MCS.

Access control rules allow for restricting client initiators to their respective storage devices, and applying fine-grained permissions to each disk drive shared. Each Access Control List (ACL) consists of a set of rules, which can include the initiator's IQN node name, IP address it will use, and the storage devices that it is granted access to. The permissions for each storage device can also be set. The Management Interface has an intuitive ACL configuration subsection.

Challenge Handshake Authentication Protocol (CHAP) Authentication provides security for the SAN, by ensuring that initiators authenticate themselves before being able to access any data on the target. Mutual authentication allows for the target to authenticate itself to the initiator and establish itself as a valid target. CHAP authentication parameters can be configured in the Management Interface.

Dynamic allocation of disk storage can be configured and allocated for iSCSI targets with Volume Management, as detailed in the Storage section of this guide. Logical volumes allocated for iSCSI, can be selected as LUNs / disks for the target to share. If the logical volume is resized, the target dynamically refreshes the size of the device to the initiator, thus allowing the initiator system to use the additional storage capacity immediately.

Caching of I/O for iSCSI is enabled by default. The appliance should be on a UPS setup. If the appliance is not shut down cleanly, all iSCSI devices with caching enabled, will lose any data stored in the cache. Caching can be enabled or disabled in the LUNs / storage exported subsection of the target configuration in the Management Interface.

Internet Storage Naming Service (iSNS) is a method of discovery of iSCSI targets and initiators on the SAN, similar to the Domain Name System (DNS) for networking. iSNS servers provide a repository of iSCSI targets and their networking settings, for initiators to use when finding a suitable target to login to. Targets and initiators need to register with the iSNS server for their data to be available in the repository. The Chelsio iSCSI Target stack provides an iSNS client to register with iSNS servers. iSNS clients can be configured in the Management Interface, by providing the iSNS server details.

For advanced configuration of iSCSI, and fine-tuning, please refer to the iSCSI Target user guide, provided along with this guide.

8.4 iSCSI Summary

- **Sections of the interface:**

1. Target service summary

iSCSI target service details such as Service status, control commands (available only in non-HA mode), number of iSCSI targets configured and running, number of Client Initiator sessions and iSNS details are shown here. You can also select the iSCSI offload mode: TOE or ULP. TOE mode runs iSCSI portal over TCP Offload Engine in terminator ASIC. Whereas, ULP runs iSCSI portal on ULP Hardware acceleration in terminator ASIC.

1.1. Target service actions (non-HA mode):

- **Enable / Disable**

You can choose to start/stop iSCSI targets using this button. Enabling the service also configures it to start automatically on system bootup. The default is to start iSCSI target services automatically.

- **Reload**

Reload all currently running targets configurations.

- Restart

Restarts the target services.

Note: In HA mode, iSCSI Initiators may fail to discover LUN devices on previously connected targets. This primarily happens if the LUN device is full. In such a scenario, go to the **Services** section under **Cluster**, and change the preferred owner to peer (secondary) node and then back again to local (primary) node.







Warnings:

Stopping or restarting all targets will cause data loss to any connected clients.

Target service details:

Service status: **installed, Auto Start Enabled, Target configured and running**

Actions:  **Disable**  **Reload**  **Restart**

Target mode:  **Update**

Targets configured: **1**

Targets running: **1**




Client Initiator sessions: **1**

iSNS details: **No client running**


Figure 8.4 (a) Target service summary and control commands (non-HA mode)

Target service details:

Service status: **installed, Auto Start Enabled, Target configured and running**

Actions:   

Service control is not available when the system is part of a cluster

Target mode: ULP ▾ 

Targets configured: **1**

Targets running: **1**

Client Initiator sessions: **0**

ISNS details: **No client running**

Figure 8.4 (b) Target service summary with disabled control commands (HA mode)

- **Kernel and application installation details**

The support details of Target software is shown here.

Kernel and application support:	
Kernel Module:	Enabled, built as module, module currently loaded
Module Release Version:	5.2.0-001-457.79caa51df2c1.hg+
Control Utility:	installed
iSNS Utility:	installed

Figure 8.4 (c) Target software installation details.

8.5 Creating a new Target

- **New target settings:**

- 1. Target Name:**

This setting is automatically suggested to you. You may use the suggested value, or specify a different one. The IQN name must meet the requirements as specified in the iSCSI IETF standard RFC 3720.

2. Target Alias:

This setting specifies the human friendly name that the Target is identified by. This name is displayed on the navigation menu to the left. Client initiators accessing the Target will report the alias name too.

3. LUN:

Select a single LUN device which can be shared by this iSCSI target. This LUN could be a Volume management device or a disk, or a RAM disk device. You can specify a custom size for a RAM disk device. The permissions for the LUN can also be specified.

Note: For Linux Initiators, after discovering and connecting to a LUN device, format it with XFS or ext4 file system before using it.

Warnings:



- a) The devices shown in gray may already in use by other targets or used for other purposes, such as for storing filesystem data.
- b) Both the RAM disk device options do not preserve data. The data is stored temporarily in memory, while the Target is active, for the regular RAM disk, and is immediately discarded, for the zero-copy option.

4. Portal Group:

Specify the IP address and TCP port that the Target should provide iSCSI service on. The default iSCSI TCP port is 3260. Change this only if you will configure client initiators to use the different port that you specify. Additional LUNs and Portals can be specified after adding the new Target.

5. Target Redirection (non-HA mode):

By using this feature, you can redirect an initiator to use a different IP address and port instead of the current one to connect to the target. The redirected target portal can either be on the same machine, or a different one. *ShadowMode* allows the redirected portal groups to be on a different USS appliance. Enable this option to use the target redirection feature. Specify the Redirection IP Address (IP of the target to which initiators will be redirected) and the port (using which initiators will connect to the redirected target).

Target Name:	<input type="text" value="iqn.2012-10.V2:3"/>	
	(in iqn.yyyy-mm... format, example: iqn.2004-05.com.chelsio.target1)	
Target Alias:	<input type="text" value="Target-3"/>	
	(A unique short identifier/name for this Target)	
LUN:	LUN Device:	<input type="text" value="MPIO (sdk): 34.18 GB"/>
	RAM Disk Size:	<input type="text" value="32"/> MB
Portal Group:	IP Address:	<input type="text" value="102.50.50.233"/>
		Note: For Cluster mode, it is recommended to use unused Cluster IP addresses to avoid service dependency.
	TCP Port:	<input type="text" value="3260"/> <input checked="" type="checkbox"/> Default
		Valid TCP port can be from 1 - 65535
Target Redirection:	Enable Shadowmode:	<input type="text" value="Yes"/>
	Redirection IP Address:	<input type="text" value="102.50.50.235"/>
	Port:	<input type="text" value="3260"/>

Figure 8.5 - Create new iSCSI Target

8.6 iSCSI target summary interface

- **Sections of the interface:**

- 1. Target summary:**

The IQN name, alias, current status and control commands are available here.

- 1.2. Target control actions:**

- **Start target:** This command allows you to start a target.
Note: Available only in non-HA mode.
- **Stop target:** This command allows you to stop the target if it is running.
Note: Available only in non-HA mode
- **Restart target:** This command will stop and start a running target.
Note: Available only in non-HA mode.
- **Reload configuration:**
The configuration of the target is reloaded, including any changes in disk sizes (for example a change in the size of a Fibre Channel disk or a hardware RAID array disk).
Note: Available only in non-HA mode.

- Delete target: The target is deleted from the configuration file.

IQN Name: [iqn.2010-09.newusr1:1](#)
Alias: Target-1
Mode: Standalone
Status: running

Actions:   

Figure 8.6(a) - iSCSI Target Summary page in Non-HA mode

IQN Name: [iqn.2011-05.SIMHA1:1](#)
Alias: HA-Target-1
Mode: Clustered
Status: running [SIMHA1]


Actions: 

Figure 8.6(b) - iSCSI Target Summary page in HA mode



Warning: Stopping, restarting or deleting a target may cause connected clients to lose data. Ensure that no clients are currently using the target, prior to stopping, restarting or deleting it. Deleting a target cannot be undone. You will need to re-add the target. The data on the LUNs in the target is **not** deleted when deleting the target.

2. Details:

A summary of the LUNs, network portals, access control, CHAP authentication and active Initiator Sessions is shown here.

Details:	
LUNs / Disks configured:	1
LUNs / Disks currently shared:	1
Network Portals configured:	102.66.66.12:65535
Network Portals currently serving:	102.66.66.12:65535,timeout=300000
Access Control:	Not configured (Disabled by default, not recommended)
Authentication:	Not configured (Disabled by default, not recommended)
Initiator Sessions active:	

Figure 8.6(c) - Target configuration details.

3. Parameters:

Advanced settings for the Target are listed here, and a setting can be modified when selected.

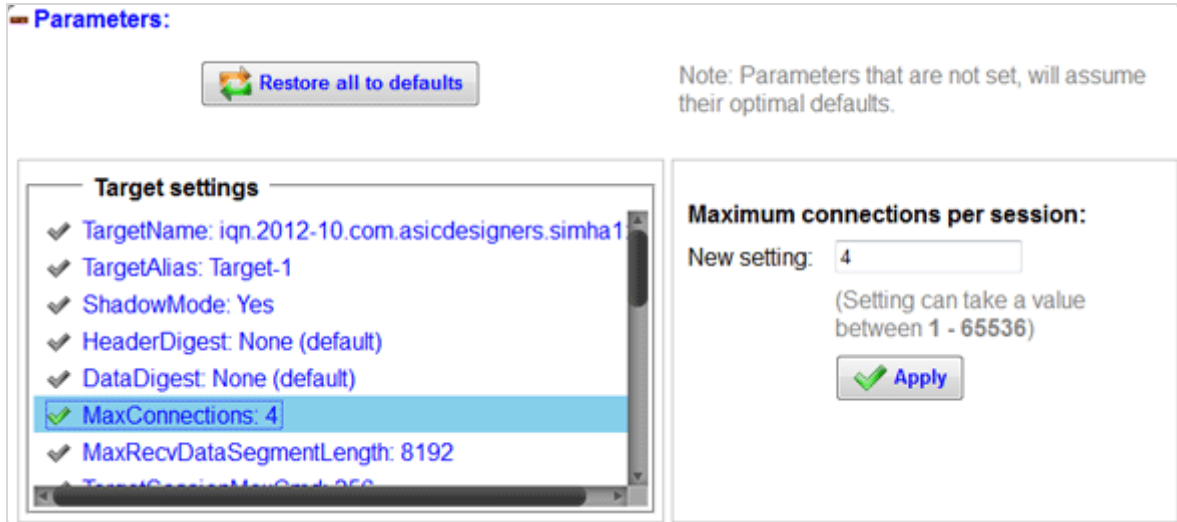


Figure 8.6(d) - Advanced target settings.

3.1. Parameters commands/settings

- Restore all to defaults: Using this command, all the values set for different parameters in the target settings can be reset to their default values.
- Target parameters/advanced settings:

Key	Valid Values	Default Value	Description
TargetName	"<target name>"		A worldwide unique iSCSI target name.
TargetAlias	"<target alias>"		A human-readable name or description of a target. It is not used as an identifier.
ShadowMode	"Yes" "No"	"Yes"	To enable or disable target redirection to external portals.
HeaderDigest	"None"	"None"	To enable or disable

Key	Valid Values	Default Value	Description
	"CRC32C"		iSCSI header Cyclic integrity checksums.
DataDigest	"None" "CRC32C"	"None"	To enable or disable iSCSI data Cyclic integrity checksums.
MaxConnections	1 to 65536	4	Initiator and target negotiate the maximum number of connections requested/acceptable.
MaxRecvDataSegmentLength	512 to 16777215 (224 - 1)	8192	To declare the maximum data segment length in bytes it can receive in an iSCSI PDU.
TargetSessionMaxCmd	1 - 2048	32	To declare the maximum outstanding iSCSI commands per

Key	Valid Values	Default Value	Description
			session.
InitialR2T	“Yes” “No”	“No”	To turn on or off the default use of R2T for unidirectional and the output part of bidirectional commands.
MaxOutstandingR2T	1 to 65535	8	The maximum number of outstanding R2Ts per task.
ImmediateData	“Yes” “No”	“Yes”	To turn on or off the immediate data.
FirstBurstLength	512 to 1677721 5 ($2^{24} - 1$)	65536	The maximum negotiated SCSI data in bytes of unsolicited data that an iSCSI initiator may send to a

Key	Valid Values	Default Value	Description
			target during the execution of a single SCSI command.
MaxBurstLength	512 to 16777215 ($2^{24} - 1$)	262144	The maximum negotiated SCSI data in bytes, of a Data-In or a solicited Data-Out iSCSI sequence between the initiator and target.
DefaultTime2Wait	0 to 3600	2	The minimum time, in seconds, to wait before attempting an explicit / implicit logout or connection reset between initiator and target.
DefaultTime2Retain	0 to	20	The maximum time, in

Key	Valid Values	Default Value	Description
	3600		seconds, after an initial wait.
ErrorRecoveryLevel	0 to 2	0	To negotiate the recovery level supported by the node. <i>Chelsio only supports 0.</i>
DataPDUInOrder	"Yes" "No"	"Yes"	To indicate the data PDUs with sequence must be at continuously increasing order or can be in any order. <i>Chelsio only supports "Yes".</i>
DataSequenceInOrder	"Yes" "No"	"Yes"	To indicate the Data PDU sequences must be transferred in continuously non-

Key	Valid Values	Default Value	Description
			decreasing sequence offsets or can be transferred in any order. <i>Chelsio only supports "Yes".</i>
OFMarker	"Yes" "No"	"No"	To turn on or off the initiator to target markers on the connection. <i>Chelsio only supports "No".</i>
IFMarker	"Yes" "No"	"No"	To turn on or off the target to initiator markers on the connection. <i>Chelsio only supports</i>

Key	Valid Values	Default Value	Description
			"No".

8.7 iSCSI Target LUN configuration

- **Sections of the interface:**

- 1. Current LUN configuration:**

The storage on the system that is shared by this iSCSI Target, is listed here. Options to modify the list are also available (if the iSCSI Target was not auto-mapped from a Fibre Channel Target).

Current LUN configuration:

LUN list (multi-select)

- 0. pool1/vol - 5.00 GB
- 1. pool1/vol2 - 6.00 GB

Edit LUN list:

Edit selected LUN:

Permission: Read-Write Read-Only

Caching: Enabled disabled

(Can only be changed when the target is stopped)

RAM Disk size: MB

LUN Device:

Figure 8.7(a) - LUN listing with options to edit.

1.1. Current LUN configuration options

1.1.1. Edit the LUN list

- **Delete Lun:** This will remove the LUN from the list. It does not affect the data on the storage device, except for RAM disks, where the data is discarded.

Please make sure to click on the **Save LUN configuration** button to save any modifications made.



Warning: Deleting a LUN can cause data loss if a client initiator is accessing the disk, when you save the configuration changes.

1.1.2. Edit the selected LUN:

- **Permission:** You can restrict the access to read-only or allow read + write access. This option is available only when the target (non-HA) or iSCSI service (HA) is stopped.



Warning: All client filesystems do not work well with read-only disks, for example NTFS. You need to verify that the client initiator can access data on a read-only LUN before setting the LUN to read-only.

- **Caching:** It is recommended to not enable this setting unless you have high-speed storage such as Flash / Solid State drives, or hardware RAID controllers, or Fibre Channel controllers, and a high performance system. Also, you should have the system and any attached storage on a UPS / battery backup, since write cache data could be lost in a power outage, causing data loss or corruption of data. This option is available only when the target (non-HA) or iSCSI service is stopped (HA).
- **RAM disk size:** This setting is only applicable to RAM disks. Please ensure that all RAM disk sizes aggregated across all Targets is lesser than half of the system's physical memory.
- **LUN device:** The device being shared can be changed here.

Note: For Linux Initiators, after discovering and connecting to a LUN device, format it with XFS or ext4 file system before using it.

▪ **Warnings:**



- a) The devices shown in gray may already in use by other targets or used for other purposes, such as for storing filesystem data.
- b) Both the RAM disk device options do not preserve data. The data is stored temporarily in memory, while the Target is active, for the regular RAM disk, and is immediately discarded, for the zero-copy option.

2. Add a LUN:

Options to add a LUN to be shared by this iSCSI Target are shown here (if the iSCSI Target was not auto-mapped from a Fibre Channel Target).

- Add a LUN:

Add a LUN by specifying an existing device

Select the Storage device:

RAM Disk size: MB

Add a LUN by allocating space from a storage pool

Storage Pool for the LUN:

New volume name:

New volume size: GB

Add Multiple LUNs

Free LUN list (multi-select)

- 0. pool1/vol4: 66.00 GB
- 1. pool1/vol5: 120.00 GB

Figure 8.7(b) - Add LUN section.

2.1. Add LUN options

2.1.1. Add a LUN by specifying an existing device

A new LUN can be appended to the end of the list, by selecting a device from the menu.

2.1.2. Add a LUN by allocating space from a storage pool

This option allows you to create a logical volume, and assign it as an iSCSI LUN, appending the newly created volume to the LUN list.

- Storage pool: Select the storage pool to create the new logical volume.
- New volume name: You can specify an optional logical volume name (e.g.: iscsi_lun10)

New volume size: You need to specify the new volume size, which will be the size of the new LUN.

2.1.3. Add Multiple LUNs

Using this option you can append multiple storage devices as LUNs to the list.

Example: Adding an iSCSI Target LUN by specifying an existing device

Using this method you can use an existing storage device (logical volume, snapshot device, clone or RAM device) as iSCSI Target LUN. Please note that only free/unassigned devices can be used.

- i. Click the **Add a LUN by specifying an existing device** option to enable it.
- ii. Select from the available storage device (listed in black) to use as iSCSI Target LUN. Devices listed in gray are in use and cannot be selected. To use RAM disk, select *RAM Disk* and enter the size. You can use the *RAM Disk (discard data)* option to test I/O performance.
- iii. Click the **Add LUN** button.

If the device was successfully added, it will appear in the LUN list in the **Current LUN configuration** section.

Example: Adding an iSCSI Target LUN by creating a logical volume.

Using this method, you can create a logical volume and then use it as iSCSI target LUN.

- i. Click the **Add a LUN by allocating space from a storage pool** option to enable it.
- ii. Select the storage pool, to create the logical volume.
- iii. Specify a name for the logical volume and size.
- iv. Click the **Add LUN** button.

If the device was successfully added, it will appear in the LUN list in the **Current LUN configuration** section.

Example: Adding multiple iSCSI Target LUNs.

Using this method, you can add multiple existing storage devices as iSCSI Target LUNs

- i. Click the **Add Multiple LUNs** option to enable it.
- ii. Select multiple devices in the list by clicking on them.
- iii. Click the **Add LUN** button.

If the device was successfully added, it will appear in the LUN list in the **Current LUN configuration** section.

8.8 iSCSI Target network portals

- **Sections of the interface:**

- 1. Current network configuration:**

The IP addresses and TCP ports for this Target are listed in this section. Options to modify the list are available.

- 1.1. Current network portals configuration actions:**

- **Delete Portal:** This command allows you to delete network portals from the list.
- **Edit selected network portal:** You can select a different IP address on the system or specify a different TCP port for the selected entry. You can also add or remove external target portals used for target redirection for the selected entry from the Port redirection list.

Please make sure to click on the **Save network portals configuration** button to save any modifications made.



Warnings:

Deleting a network portal entry can cause client initiators to fail to connect to this Target, if they were configured to use the entry that changed.

Current network configuration:

Portals list (multi-select)

- 1. IP: 102.50.50.113 - Port: 3260
- 2. IP: 102.50.50.116 - Port: 3260
- 3. IP: 102.50.50.235 - Port: 3260

Edit network portals list:

Edit selected network portal:

IP address:

TCP port: Use default iSCSI service TCP port:

Valid TCP port can be from 1 - 65535

Redirect List

- 102.50.50.116:3260
- 102.50.50.235:3260


Portal redirection list:

Add portal to redirect:

Figure 8.8(a) - List of currently configured network portals with options to modify items.

2. Add a network portal:

Network portals can be added here, which is a combination of an IP address, External IP for Redirection and a TCP port, for this target to serve iSCSI client initiators.



The screenshot shows a web form titled "Add a network portal:". It contains the following fields and controls:

- IP address:** A dropdown menu with the text "---Select an IP address---" and a downward arrow.
- External IP for Redirection:** A text input field.
- TCP port:** A text input field containing the value "3260". Below this field is a note: "Valid TCP port can be from 1 - 65535".
- Use default iSCSI service TCP port:** A checkbox that is checked.
- Submit button:** A button with a green checkmark icon and the text "Add network portal".

Figure 8.8(b) - Option to add a network portal.

1.1. Add a network portal options:

- IP address: Select the IP address for the network portal. You may configure the same IP address for multiple portals, if you specify different TCP ports for each entry.
- External IP for Redirection

If Target Redirection is enabled (while creating a target), you can add the IP of the external target to which initiators will be redirected.

- TCP port: Set the TCP port to use. It is pre-filled to the default iSCSI TCP port of 3260. If you wish to use a different port, uncheck the “Use default iSCSI service TCP port” option.

8.9 iSCSI Target clients / initiators

- **Sections of the interface:**

1. Summary

This section lists the ACL and CHAP authentication policies for the target and allows you to modify them.

1.1. Target policy settings

- **Access Control:** This setting turns on or off the access control feature of the iSCSI target service. The iSCSI target is capable of restricting which client initiators can login, as well as what LUNs / disks they are allowed to access. It can also restrict the type of access to the LUN as read-only or allow read-write access.
- **CHAP authentication policy:** This setting specifies the CHAP authentication offered by the target, when a client initiator connects to it.
- **Authentication type:** This setting specifies the type of CHAP authentication done by the target.

- CHAP Mutual authentication Target parameters: Specify the iSCSI target's CHAP username and secret, to provide to client initiators, who wants to authenticate this target as a valid target.

The screenshot displays a configuration summary for an iSCSI target. At the top, a red minus sign icon is followed by the word "Summary:" in blue. Below this, the "Access Control" is set to "Disabled" with a dropdown arrow, and a green checkmark "Apply" button is present. The "CHAP" section is expanded, showing "CHAP authentication policy" set to "Auth preferred, but optional (CHAP,None)" with a dropdown arrow, and "Authentication type" set to "One-way (Client only)" with a dropdown arrow. Under "CHAP Mutual authentication Target parameters:", the "Username" field contains "iqn.2010-12.bolt:1", the "Password" field is empty, and a green checkmark "Apply" button is located below. At the bottom, "Clients connected:" is shown as "1".

Figure 8.9(a) - Summary with policy settings for the iSCSI target.

2. Clients

The clients (iSCSI initiators) currently connected to the system are shown here, along with those that are configured with ACL rules or CHAP authentication parameters.



Figure 8.9(b) - Clients list with settings shown below, and status on the left.

2.1. Clients settings

- Clients list: This area displays the list of currently connected and configured initiators. Click on an initiator to view its configuration information below it.

- **Current status:** If the selected Initiator is successfully connected to a Target (indicated by a green icon), this section provides status information such as number of sessions and TCP connections with port numbers in the Connection list.




Figure 8.9(c) – current status of the iSCSI connection

- **Access control (per client initiator):** On clicking and highlighting an initiator client in the list above, if any access control settings are defined, they are shown below, otherwise the defaults are displayed, with a warning in the status. The access control area has two settings that can be set for fine-grained control. The first is the client initiator's IP address. This specifies the IP address that will be used by the client to connect to this system. The second

is the LUN list that the initiator is allowed to access. LUNs can be masked as required. The list of disks that the client will be able to access, and the order of those disks, is decided by the unmasked LUNs here. You may also restrict the LUN permissions to read-only or allow read + write.

Note: Masking of all the LUNs in the list is not allowed since at least one LUN will have to be available for clients to access.

Access Control:

Status: **Access Control entry is configured. Client IP address restriction configured. Target IP address restriction configured.**

IQN name: iqn.2013-02.com.asicdesigners.Sanga

Client IP address:
Eg: 192.168.1.1, 192.168.1.2

Target IP address:

Portals list (multi-select)

- 0. 102.66.66.113

LUNs access:

All LUNs are accessible with default permissions

Customized LUN access as below:

LUN list (multi-select)

- Target LUN 0: pool1/vol1 - 80.00 GB - LUNNO 0 - RW
- Target LUN 1: pool1/vol2 - 94.00 GB - LUNNO 1 - Masked

Figure 8.9(d) - ACL settings for a client.

- CHAP authentication: The CHAP secret for the initiator can be specified here. The username used is the IQN name of the initiator.




Figure 8.9(e) - Setting up CHAP authentication for Client/ Initiator

Example: Configuring CHAP authentication for an iSCSI target

The procedure mentioned here assumes that you have already created a target and the Initiator has been configured to connect to the target successfully.

- Select the target for which you want to enable CHAP under the iSCSI module. Click on the **Clients/initiators** sub-module.
- In the **Summary** section, select *Enforced* in the *CHAP authentication policy* drop-down.
- The Authentication type will be *one-way* by default.
- Click **Apply**.

- v. Now, in the **Add a Client** section, enter the IQN name of the Initiator. Enter a CHAP secret key for the target. This will enable one-way CHAP authentication. Initiators trying to access the target will have to use the same key to login.
- vi. Click **Save Client settings**.
- vii. The newly added initiator will be listed in the **Clients** section. To change the CHAP secret key, select the client entry in the list and then in the **CHAP authentication** section, enter the new secret key.
- viii. To enable Mutual CHAP authentication, go to the **Summary** section and in the **Authentication type** drop-down, select *Mutual (Client+Target)*. This will enable the password field. Enter a secret key which will be used by the iSCSI initiator to authenticate the target.
- ix. Click **Apply**.

Note: For instructions on how to access an iSCSI target with CHAP authentication please refer to **7.6 iSCSI Initiator** section.

Example: Configuring Access control for iSCSI Clients/initiators

Please ensure that no iSCSI service is currently running before you attempt to set access control for a target.

- i. Select the target for which you want to configure access control under the iSCSI module. Click on the **Clients/initiators** sub-module.
- ii. In the **Summary** section, select *Enabled* in the **Access Control** drop-down.
- iii. Click **Apply**.
- iv. Expand the **Clients** section to select the initiator or add a new initiator using the **Add a client** section.
- v. Select the initiator from the list and expand the **Access Control** subsection.
- vi. If not specified while adding the client, enter initiator IP address.
- vii. By default, LUN permissions set in the **LUNs** sub-module for the target will apply. To change them, select the **Customized LUN access as below** option.
- viii. In the LUN list, select a LUN and change the read-write permissions or use the Mask/Exclude button to hide the LUN from the client. Use multi-select to apply similar permissions to different LUNs.
- ix. Click **Save Access Control settings** button.

3. Add a client

A client initiator can be defined here, with its IQN name and an optional IP address that it will use to connect from. Its CHAP authentication parameters can be specified too. If the initiator is part of a cluster environment, please specify the IP addresses of both HA nodes in the *Client IP address* field separated by a comma.

- Add a client:

Add a client ACL and CHAP authentication entry by specifying the parameters below:

Client initiator IQN name:
Eg: iqn.2004-05.com.chelsio.target1

Client IP address:
Eg: 192.168.1.1, 192.168.1.2

Target IP address:

CHAP secret (length of 12 - 16, optional): show password

Figure 8.9(f) - Add client section.

9. USS High Availability

The high-availability feature of the Unified Storage Server depends on correct configuration of hardware and software. iSCSI Targets, NFS and CIFS shares can failover from one USS system to the other, provided all criteria are met.

There is a clustering wizard, which checks and helps in configuring the USS systems in HA mode.

Pre-requisites:

1. Ethernet network:

At least 2 shared Ethernet networks between the USS systems are required. Design the Ethernet connectivity from the USS systems, in such a manner that both the systems can reach a particular network. This allows the systems to communicate with each other over Ethernet, and provides failover Ethernet paths for iSCSI connections. It also allows NAS clients to access shares from either USS system on the same network.

2. SAS:

SAS can be used as the shared storage. SAS HBAs can be used to access an external shared SAS Disk Enclosure. At least one shared disk is required, to allow the USS systems to setup cluster quorum.

3. IPMI:

IPMI (Intelligent Platform Management Interface) capability is required in each USS system. The systems need to be able to reboot each other if they ever reach a critical or hung state. This can be achieved by using IPMI-over-LAN, which is possible when the IPMI enabled LAN ports on each USS system are connected to the same Ethernet network. A static IP address has to be assigned to each USS system's IPMI LAN interface. This static IP address should be reachable from the other node's regular networking interfaces, so it is recommended to use the same IP subnet, to assign addresses for the IPMI LAN interface and the OS networking interface, on both nodes. IPMI configuration is found under "System Tools" in the management GUI.

4. Quorum device:

The clustering functionality requires a shared storage device to use for quorum state information. This device need not be large, a 100MB partition is sufficient. A Shared state partition of 5GB is also required. If a

free disk is selected as Quorum disk, the cluster creation wizard will create the required quorum and shared data partitions.

Once these pre-requisites are met, start the cluster configuration wizard, and specify the cluster configuration details as required. Add at least 2 shared IP addresses on 2 shared networks, and select the disk and partition created in step 4 of the pre-requisites. Specify a password for the IPMI user for both systems.

The cluster should now be created successfully. This allows you to configure clustered iSCSI targets, under the “iSCSI” section, in the “Create new target” option. You can also create clustered NFS and CIFS shares under the File sharing section.

Cluster resources and services can be viewed and modified in the GUI too.

Note: To know about HA dependencies, please refer to the **Dependencies** section under **Licensing**.


9.1 Create Cluster

- Adding peer node

Please ensure that the prerequisites mentioned are met before specifying peer system's address. Enter the peer node's IP address and click Next.

Create a Cluster

Specify peer system's IP address or hostname:

 **Warning: Prerequisites for cluster creation**

Prior to creating a cluster, ensure that you have configured the following correctly on both nodes:

1. Management IP address for both cluster nodes (DNS IP address of node)
2. Fully qualified hostname for both cluster nodes (DNS hostname of node)
3. Minimum of two shared networks should be available.
4. IPMI BMC LAN IP address
5. IPMI BMC user authentication
6. Date, Time, and Timezone should match on both the nodes
7. Network time synchronization (NTP) should be enabled on both nodes
8. Local and peer nodes should not be added to domain.

Figure 9.1 (a) - Create Cluster screen

- Cluster compatibility summary

This section displays the various parameters for which the Local node and Peer node are verified for compatibility. Please note that in order to proceed further all the parameters specified have to match.

Cluster compatibility summary:	
✓ Nodes:	newusr1 newusr2
✓ Date & Time: [mandatory]	Yes 13:42 2011-02-04 Asia/Kolkata NTP service is enabled
✓ Operating System: [mandatory]	Yes Unified Storage Server 2.00.31
✓ Processor: [not mandatory]	Yes Manufacturer: Intel Speed: 3000 MHz Processors: 1 Processor cores: 4
✓ Memory: [not mandatory]	Yes Total memory: 4096 MB 2048 MB modules: 2 Maximum memory supported: 8 GB
✓ Domain name [mandatory]	Yes Both local and peer are not added to any domain.
✓ IPMI BMC: [not mandatory]	Yes BMC: present IPMI version: 2.0 compatible Chassis power policy after failure: always power on BMC LAN IP type: Static IP address BMCs reachable between peers: yes
✓ Storage: [mandatory]	Yes 8 shared disks
✓ Backplane connectivity: [mandatory]	Yes
✓ Network: [mandatory]	Yes 2 shared networks

Figure 9.1(b) - Cluster compatibility summary

Note: For non-Supermicro SBB systems, the *backplane connectivity* parameter will display an incompatibility error (a red cross). If this is the case, you can proceed to the configuration section where you will have to configure the backplane IP addresses manually.

Note: For Supermicro SBB systems please ensure that backplane IP addresses (Backplane cluster connection) in the **Network** section have been configured correctly in the 169.254.x.x network with subnet mask 255.255.0.0

- **Configure the cluster**

This section allows the user to configure various settings like Cluster IP addresses, Management IP addresses, Quorum disk etc. After specifying the required fields click Next to proceed.

Configure the cluster:

Unique Cluster ID: 859B487C

Cluster IP addresses: 102.50.50.13, node: SIMHA2, us: [Remove Cluster IP]

Add IP address: [Add Cluster IP]

Shared network: --Select a network--

IP address: []

Preferred node: SIMHA2

Usage: NAS / File sharing

Backplane IP addresses:

Local node: eth2: 102.25.25.76 (SIMHA1)

Peer node: eth4: 102.25.25.78 (SIMHA2)

Quorum disk:

Shared disk 1: sdd 68.37 GB

Quorum partition: Create a quorum partition

Shared state partition: Create a partition

Shared disk 2: sde 68.37 GB

Quorum partition: Create a quorum partition

Shared state partition: Create a partition

Fencing mechanism: IPMI BMC

IPMI BMC authentication:

Local BMC user: ID: 3 - Name: root

Local BMC password: [] [Show password]

Peer BMC user: ID: 3 - Name: root

Peer BMC password: [] [Show password]

[Next] [Cancel]

Figure 9.1(c) - Configuring the cluster

- **Cluster creation summary**

This section displays the result after proceeding from the Configure the Cluster screen. After the Cluster nodes have been created successfully, the Cluster Summary page displays the Cluster status, services and nodes along with Quorum disks and status.

Note: The Quorum disk status may take some time to reflect the changes.

- Summary:

Cluster status: **Cluster is running**
Local cluster services: **installed, enabled, running**

Actions:

Cluster nodes: **Clus1[local] [online]**
Clus2[peer] [online]

Quorum status: **Cluster is quorate**

Figure 9.1(d) - Cluster Summary

- **Configuration**

Users can use this section to modify IPMI password for both local and peer nodes.

Configuration:

Fencing mechanism:

newusr1 [local]
IPMI IP address: 10.193.185.64
IPMI BMC username: ID: 2 - Name: root ▼
IPMI password: ●●●●●●

newusr2 [peer]
IPMI IP address: 10.193.185.62
IPMI BMC username: ID: 2 - Name: root ▼
IPMI password: ●●●●●●


 **Apply**

Figure 9.1(e) - Configuring IPMI BMC password

Example: How to create a cluster

Follow the steps mentioned below to create a cluster.

- i. In the System summary section, expand the **System** section and click on **Create Cluster** button. You will be directed to the Create Cluster page.
- ii. Before you proceed, it is highly recommended that the requirements mentioned in the Warning box are met. Or else, cluster creation will fail. Now, enter the peer system's IP address or hostname and click **Next**. The local and peer node will now be verified for compatibility.
- iii. After the verification is complete, the **Cluster compatibility summary** will display a summary of various parameters for which the local and peer nodes were verified. All the parameters marked *mandatory* will have to pass, indicated by a green tick and yes, or else the cluster will not be created.
- iv. Now, the cluster will have to be configured. Start by specifying a name for the cluster.
- v. Add a cluster IP by selecting the network and specifying a unique IP address. Also select the preferred node and the usage (iSCSI target/NAS) for that particular IP address. Multiple Cluster IP addresses can be added.
- vi. Specify the IP addresses (used for *heartbeat*) of the local and peer node in the *Backplane IP addresses* section (only for non-Supermicro SBB systems; for Supermicro SBB systems, the IP addresses will be automatically selected).

- vii. Next, you need to specify the quorum disk (cluster's configuration database). From the available disks, select one. You can either use any of the existing partitions on the disk or let the wizard create a new one. You can also configure an additional (optional) quorum disk.
- viii. Choose the fencing mechanism. IPMI BMC is recommended.
- ix. Next Configure IPMI settings like username and password for both nodes and click **Next**.
- x. Cluster with the specified options will now be created.

If the cluster was created successfully, the Cluster Summary page will display the Cluster status, services and nodes along with Quorum disks and status.

9.2 Cluster Resources

- Cluster IP addresses

Users can use this section to add/remove Cluster IP addresses and its usage (iSCSI Targets/ File Sharing).

Cluster IP addresses:

102.50.50.55
104.50.50.57

Select an IP address from the list above for status and actions.

Resource status:
Current node:
Preferred node:
Usage:

Remove IP address:

Add IP address:

Shared network:

IP address:

Preferred node:

Usage:

- iSCSI target
- iSCSI target
- NAS / File sharing

Figure 9.2 (a) - Configuring cluster IP addresses

- **Shared storage devices**

This section displays all the shared LUNs. Use the Rescan shared storage button to discover and add any newly added LUNs to the list.

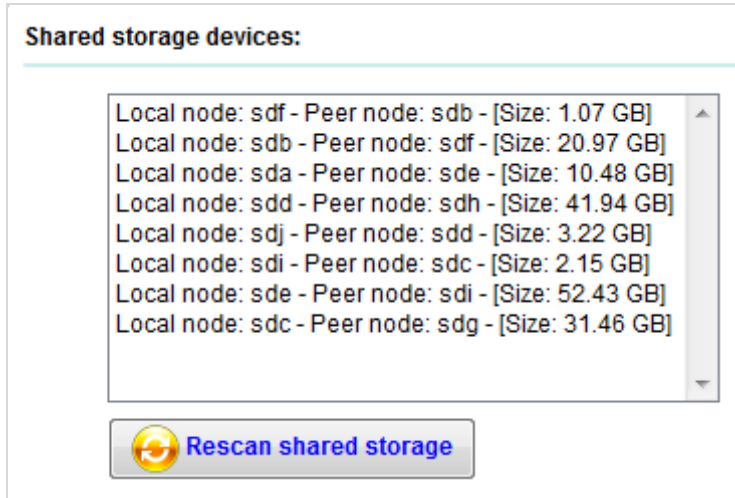


Figure 9.2 (b) - List of shared LUNs

9.3 Cluster Services

Users can use this section to start/stop cluster RAID/pool services and change the preferred owner. All clustered File Shares, File Systems, Storage Pools, Software RAID Arrays and used IP addresses are listed here. Clicking on any of these will display the related resources.

The screenshot displays a web-based interface for managing cluster services. On the left, a 'Services list' table shows a single entry: a green checkmark, a green circle, and the text '[running] md0'. To the right of this list, the service's details are shown: 'Status: Started', 'Actions: Start (green play button) Stop (red stop button)', 'Current owner: cluster5', and 'Preferred owner: cluster5' with a dropdown arrow and a green checkmark button. Below this is the 'Service topology' section, which lists various resources: 'FTP shares: 16', 'HTTP shares: 1', 'Filesystems: 16', 'Storage pools: 2', 'RAID arrays: 1', 'Physical disks', and 'Network interfaces'. At the bottom, there is a 'Resources list:' label.

Figure 9.3 (a) - cluster services

10. File Sharing

10.1 File sharing overview

File sharing protocols and applications are more commonly used by Desktop and Notebook clients. File sharing protocols allow access to data on the file server from almost any operating system and platform, thus providing ease of use for regular users. The supported file sharing protocols are Network File System (NFS), Common Internet File System (CIFS), File Transfer Protocol (FTP) and Hyper Text Transfer Protocol (HTTP).

- ▶ **Network File System (NFS)** is widely used in UNIX / Linux environments. It integrates with UNIX authentication mechanisms, and domain services such as Network Information Service (NIS). IP address or hostname based access control can be used.
- ▶ **Common Internet File System (CIFS)** is the standard file sharing protocol for Microsoft Windows. It integrates with Windows user authentication, as well as Active Directory Domain authentication. Access rules based on hostname or IP address or users are also available.

- ▶ **File Transfer Protocol (FTP)** is a legacy file sharing protocol, widely used over the Internet to transfer files or folders. This protocol does not generally allow mapping a folder hierarchy to a remote client system's folder. Common web-browsers support FTP.
- ▶ **Hyper Text Transfer Protocol (HTTP)** is an application layer network protocol built on top of TCP. HTTP clients (such as Web browsers) and servers communicate via HTTP request and response messages.
- ▶ **Lustre file system** is a scalable, secure, robust, and highly-available cluster file system that addresses the I/O needs, such as low latency and extreme performance, of large computing clusters. Lustre clusters contain three kinds of systems:
 - File system clients, which can be used to access the file system
 - Object storage servers (OSS), which provide file I/O service
 - Metadata servers (MDS), which manage the names and directories in the file system

Note: To know more about Lustre support, please refer to the **Dependencies** section under **Licensing**.




- **Sections of the interface:**

1. CIFS

The current status of the CIFS service is displayed. Options to enable / disable, reload and restart the service are available (only non-HA mode). The number of shares configured and currently active, and the status of any Active Directory domain membership is also listed.

CIFS details:

CIFS services status: **installed, auto-start enabled, running**

Actions:  **Disable**  **Reload**  **Restart**

Configured Shared folders: 2

Currently Shared folders: 2

ADS Domain membership status: **Saved domains: 0**

Current ADS status: **Current domain: none**

ADS service status: **ADS service is stopped.**

Figure 10.1 (a) - CIFS details (non-HA mode)

CIFS details:	
CIFS services status:	installed, auto-start enabled, running
Configured Shared folders:	2
Currently Shared folders:	2
ADS Domain membership status:	Saved domains: 0
Current ADS status:	Current domain: none
ADS service status:	ADS service is stopped

Figure 10.1 (b) - CIFS details (HA mode)




Note: Clients accessing CIFS shares may experience timeout issues, if system load is high or if the back-end (storage) device is slow.

2. NFS

The current status of NFS services such as the NFS daemon, portmapper, locking daemon, etc., is shown here. The number of shares configured and currently active, and the status of any NIS domain membership is also listed.

NFS details:

NFS services status: **installed, auto-start enabled, running**

Actions:  **Disable**  **Reload**  **Restart**

Configured Shared folders: 2

Currently Shared folders: 2

NIS Domain membership status: **Saved domains: 0**

Current NIS status: **Current domain: None**

NIS service status: **NIS service is stopped.**

Figure 10.1 (c) - NFS details (non-HA mode)

NFS details:	
NFS services status:	installed, auto-start enabled, running
Configured Shared folders:	2
Currently Shared folders:	2
NIS Domain membership status:	Saved domains: 0
Current NIS status:	Current domain: None
NIS service status:	NIS service is stopped.



Figure 10.1 (d) - NFS details (HA mode)

3. FTP

The current status of the FTP service, and the shares configured and active are shown here.

FTP details:

FTP service status: **installed, auto-start enabled, running**

Actions:  **Disable**  **Restart**

Configured Shared folders: 2

Currently Shared folders: 2

Figure 10.1 (e) - FTP details (non-HA mode)

FTP details:

FTP service status: **installed, auto-start enabled, running**

Configured Shared folders: 2

Currently Shared folders: 2




Figure 10.1 (f) - FTP details (HA mode)

4. HTTP

The current status of the HTTP service, and the shares configured and active are shown here.

HTTP Details:

HTTP Service Status: **installed, auto-start enabled, running**

Actions:  **Disable**  **Reload**  **Restart**

Configured Shared folders: 2

Currently Shared folders: 2

Figure 10.1 (g) - HTTP details (non-HA mode)

HTTP details

HTTP Service Status: **installed, auto-start enabled, running**

Configured Shared folders: 2

Currently Shared folders: 2

Figure 10.1 (h) - HTTP details (HA mode)

5. Service commands (non-HA mode)

- 5.1. Enable / Disable: This enables or disables the service to start when the appliance boots up (auto-start). Enabling/disabling also starts/stops the service on the appliance.
- 5.2. Reload: This refreshes the service, applying any changes in the configuration file.
- 5.3. Restart: This command stops and starts the service.



Warning: Please be aware that stopping or restarting a service can cause all connected clients to lose connectivity, and possible loss / corruption of any data that the clients were saving.

10.2 Global Settings

- **Sections of the interface:**

1. CIFS settings

Global settings that apply to CIFS service are listed here.

- 1.1. Network devices / IP addresses enabled for CIFS: This is an Access control feature, allowing you to limit the CIFS traffic to certain networks only, if required.
- 1.2. Windows workgroup name: This sets the default workgroup the CIFS service joins. In case you are joining an Active Directory domain, use the “Users” page to configure the domain, this setting will be updated automatically.
- 1.3. Idle time before disconnecting client: Specify a non-zero idle time here, in case there are many idle connections to the system which are not required.
- 1.4. Interval between checks for an inoperative client: Amount of idle time before checking on the client.
- 1.5. Map user to Guest policy: This setting allows you to map client users who fail to authenticate correctly to the guest user on the system automatically. This is not recommended by default, unless you have very specific requirements to do so.

- 1.6. Allow users with blank password: Users without a password can be disallowed from connecting to any CIFS shares. This may affect the guest user, since there is no password assigned to the guest user by default.
- 1.7. NetBIOS support: NetBIOS name resolution is a legacy name resolution mechanism, used by older versions of Windows. Enable it only if it is required for your environment.
- 1.8. Async IO: Enabling this parameter will boost Samba filer server performance. When enabled, server will accept and process other I/O operations, while the requested read or write function continues in the background.

2. NFS settings

Global settings that apply to NFS are listed here. Note that these settings require a restart of the NFS services to be applied.

- 2.1. NFS v2: This setting allows for enabling/disabling the supported NFS protocol version for the NFS service. It is highly recommended to leave it enabled.
- 2.2. MOUNTD, LOCKD, STATD, RQUOTAD service ports: These ports are generally randomly assigned; change them only if you need to configure them to pass through a firewall, where you need a static port.

- 2.3. NFS and portmapping service ports: This is not configurable, since these are registered and well-known ports that all clients will use. These ports are reported here, to allow for any firewall configuration to be updated.
- 2.4. NFS execution threads to spawn: If the service is taking time to respond to new client connections, you can increase this number, to allow for a quicker response to new connections. Note that there may be other factors that could be impacting response to a client that additional threads of the NFS server will not solve.

3. FTP settings

Global settings that apply to FTP are listed here.

- 3.1. Anonymous user login: This setting allows you to enable or disable guest / anonymous access to FTP shares.
- 3.2. Local user logins: This setting allows users configured on the system, or joined to an NIS or ADS domain to login and access FTP shares.



Warning: Local users should be required to use secure FTP, since regular FTP will transmit all passwords via plain text over the network and is insecure by nature.

- 3.3. Secure FTP via SSL, SSL encryption protocols: These settings allow you to enable or disable Secure FTP, and configure encryption for Secure FTP.
- 3.4. Enforce SSL encryption for Local user logins: This setting is highly recommended, if you are allowing local user logins.
- 3.5. Write / Upload to FTP shares: You can completely disable all write access to all shares using this setting.
- 3.6. Anonymous user Write/Upload to FTP shares: Please be careful when enabling this setting. Only enable this if you have a share that may contain possible malicious data, and will never be accessed by local users. It is always better to configure a regular user account for uploading files to a share, and sharing that user credentials as required.

4. HTTP settings

Global settings that apply to HTTP are listed here.

4.1. Server Admin: Here the user can mention the email address, where problems with the server should be e-mailed.

10.3 File sharing configuration

To configure file sharing a path to share, is chosen. This builds the hierarchy of folders appearing in the share, on the client. This 'shared folder' can be attached to an underlying storage device, such as a logical volume, by formatting a logical volume with a file system and attaching it to the folder.

Once the 'shared folder' is attached to the storage device, it has sufficient free space to store data for clients to access. The share can be configured for any or all of the supported protocols (CIFS, FTP, NFS, HTTP), so that the data is simultaneously available on any type of client.

In NFS, options for showing file systems attached to folders inside the parent shared folder can be set. The access level for administrative and non-administrative users can be specified, along with the permissions for the share.

CIFS allows for setting user access lists, and IP address or hostname based access. The shared folder can be given a different share name, which appears to clients.

FTP allows for setting guest or regular user access, and listing a shared folder with a different share name. FTP authentication is transmitted over plaintext, and is inherently insecure. It is not advisable to use FTP with Domain authentication, unless using encrypted Secure FTP.

In addition to listing a shared folder with a different share name, HTTP allows for setting authentication for selected users or all in the Users list.

The Management Interface shows all protocols shares for a particular folder existing on the Appliance, bundled together. This allows for easily viewing all the shares for a particular folder.

- **Sections of the interface:**

1. Shared folders listing

The folders on the system that are shared are listed on this page.

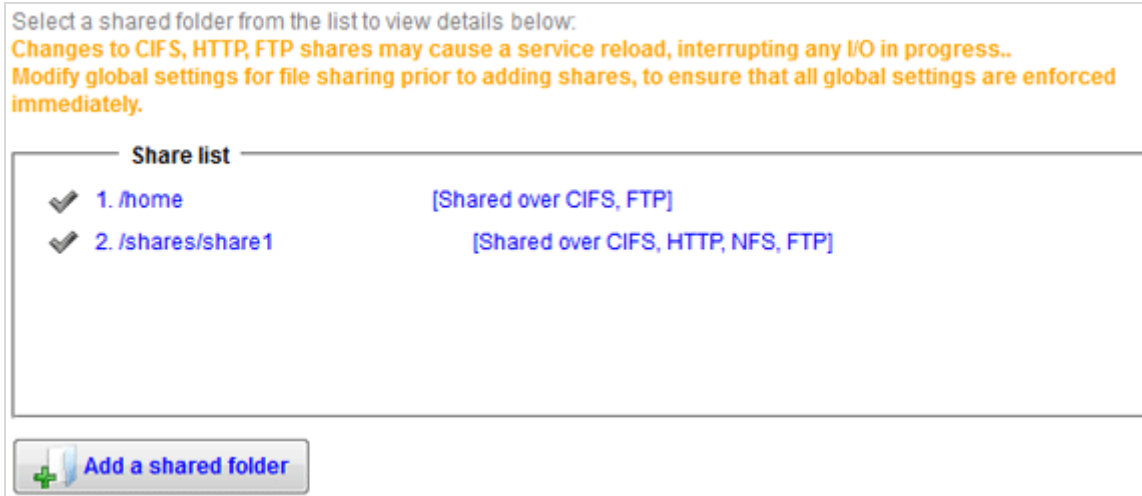


Figure 10.3(a) - list of folders shared on the system.


/shares/share1		 Edit folder ownership and permissions
Status: Folder exists	Storage Device: pool_1/vol1	Currently mounted
Folder size: 4.289 MB	Free space: 5105.711 / 5110.000 MB	Space used: 1 %
Owner: root	Group: root	Specials: None
Access: rwX	Access: rwX	Others: rwX

Figure 10.3(b) - Shared folder list with index number, ownership and permissions details.

1.1. Add a shared folder

Clicking on the **Add a shared folder** button redirects to this page, where the user can share a folder using different protocols, which can be further configured.

Add a new CIFS / NFS / FTP / HTTP Share

Information: Folder and share permissions

Please ensure you check / change ownership and permissions for the folder to the correct user and group after adding the share.
ACLs for allowing or restricting hosts or users can be edited after adding the share. Other settings can also be set by editing the share after addition.

Folder to share:

Folder "/shares" exists.

Share cluster mode: ▼

Share Protocol:

- CIFS
- NFS
- FTP
- HTTP

Figure 10.3(c) - Adding a shared folder in Non-HA mode

Add a new CIFS / NFS / FTP / HTTP Share

i Information: Folder and share permissions

Please ensure you check / change ownership and permissions for the folder to the correct user and group after adding the share.
ACLs for allowing or restricting hosts or users can be edited after adding the share. Other settings can also be set by editing the share after addition.

Folder to share:

Folder "/shares" exists.

Share cluster mode:

Cluster IP addresses:

Cluster IPs (multi-select)

- 1. 125.55.55.5 [node: clus1] [unused]
- 2. 125.55.55.6 [node: clus1] [unused]

Share Protocol:

- CIFS
- NFS
- FTP
- HTTP

Figure 10.3(d) - Adding a shared folder in HA mode

1.1.1. Sections

- Folder to share and share protocols to enable: The folder to be shared can be specified here by typing the full path, or by selecting the folder using the “Browse” button popup menu.



Figure 10.3(e) - Browse for folder popup menu

- Cluster IP addresses (HA mode): Cluster IP using which you want to share the particular folder.
- Add share protocol settings: CIFS, NFS, FTP and HTTP share settings can be specified while adding a shared folder from their respective sections. The section is shown only if the share protocol is enabled.

Note: For HA mode, the folder to be shared should have a volume attached to it before attempting to add it.

- NFS: To add a shared folder using NFS select the NFS checkbox and specify various related settings and permissions like host/subnet/netgroup to which the share is allowed access, Read-Write permissions, Write caching method, etc. Finally click on Apply.
- CIFS: To add a shared folder using CIFS, select the CIFS checkbox and provide the share name. Permissions like Guest/Anonymous access, Read+Write and Read-Only can be set. You can also choose to enable async or sync caching. Finally click on Apply.
- FTP: To add a shared folder using FTP, select the FTP checkbox and provide the share name. You can also enable/disable Guest/Anonymous access for

the particular share. Guest-only access can also be allowed by enabling the "Only Guest access allowed" option. Finally click on Apply.

- HTTP: To add a shared folder using FTP, select the HTTP checkbox and provide the share name. You can choose to enable/disable authentication for users. Selecting All valid Users enables Authentication for all the users in the list, whereas selecting Selected Users enables choosing which users to enable Authentication for. Finally click on Apply.



Warning: Do not share system folders such as the root folder “/” or “/proc” or “/sys” or other system folders.

Example: How to add a new share

Following is an example showing how to add a new share in non-HA mode. For HA mode, first mount a volume to the folder to be shared and then follow the steps mentioned below.

- i. Click on the **Add a shared folder**. This will navigate to a new page.

- ii. Click the **Browse** button to locate the folder to share or enter the path manually if the location is already known.
- iii. If adding a share in HA mode (the *Share cluster mode* is *Clustered*), then select the cluster IP, using which you want to share the folder.
- iv. Enable the adjacent checkboxes for protocol(s) using which the folder has to be shared. The corresponding settings for each protocol enabled, will be displayed below.
- v. Provide a share name (except for NFS) and other related settings and click **Apply**.
- vi. If the share was added successfully, you will be redirected to the **File Shares** page and the newly added folder will appear in the **Share list**.

Note: For more information on how to add a cluster, please refer chapter 9. High Availability.

1.2. Folder actions

1.2.1. Edit folder ownership and permissions: This command allows you to edit the owner of the folder and the permissions for the folder. There are 2 sets of permissions that apply for a client to access a file or folder in a share: The share permissions and the folder permissions (on the filesystem, locally on this system). If either of them is not correctly set, then the folder may not be accessible to clients.

Folder Owner and Group

Folder Owner: Local user: root (Administrator) Domain user:

Folder Group: Local Group root Domain Group

Apply ownership to subfolders and files:

Global folder permissions

Owner access: Read Write Execute

Group access: Read Write Execute

Other users access: Read Write Execute

Specials: Set UID: Set GID: Sticky:

Apply permissions to subfolders and files: No

Advanced folder permissions

None

Edit permissions of user/group in list

Permissions: Read Write Execute

Add new user/group

Local user: Local user: root Local group: Local group: root Domain user: Domain group:

Warning:
Ownership and permissions changes will take effect immediately and any clients connected will be allowed or denied access to files based on them.

Figure 10.3(f) - Edit folder ownership and permissions page.

Example: How to edit folder ownership and permissions for a folder

- i. In the **Share list**, select the folder for which you want to change ownership and permissions.
- ii. Click on the **Edit folder ownership and permissions** button. This will navigate to the folder properties page.
- iii. A local or a domain user can be configured as the folder owner. To make a local user as the owner, enable the **Local user** radio button, and choose from the list of available users. By default, **root** is the owner of the shared folder. To make a domain user as the folder owner, enable the **Domain user** radio button. Then click on the search icon. In the pop-up that appears, select the domain and enter a valid user name. NIS/ADS domain should be configured to make a Domain user as the folder owner. Click **Apply**.
- iv. Similar to users, local or domain groups can be configured as the folder group. When you select a local user as the folder owner, the primary group to which the user belongs becomes the folder group automatically. To make a domain group as the folder group, enable the **Domain Group** radio button. Then click on the search icon. In the pop-up that appears, select the domain and enter a valid group name. NIS/ADS domain should be configured to make a Domain group as the folder group. Click **Apply**.
- v. Configure global permissions for the folder like Read, Write, Execute permission for the folder owner, group and other users (users not belonging to the folder group).
- vi. You can further configure permissions to the folder using the **Advanced folder permissions** section. Please note that the options in this section will be available only if a logical volume has been mounted

on the shared folder. Add/remove a local /domain user or group and configure the permissions for that particular user/group.

vii. Click **Apply changes**.

2. Per shared folder shares list

NFS, CIFS, HTTP and FTP shares of the shared folder are listed below the folder's properties.

The screenshot displays the 'NFS shares' configuration window. At the top, it shows a list of shares with the following details for the first entry:

- 1. /shares/nfs1
- Mode: Standalone

Below the list, there are three buttons: 'Edit share', 'Disable', and 'Delete'. A 'Share details' section is expanded, showing two columns of settings:

Saved settings:	Running settings:
Shared to: 102.50.50.0/24 Type: subnet	Shared to: 102.50.50.0/24 Type: subnet
all: caching: enabled check subtree: no coalesce writes: enabled filesystem id: 9166	caching: enabled check subtree: no coalesce writes: enabled filesystem id: 9166 guest gid:

Figure 10.3(g) – NFS share with details and actions (non-HA mode)

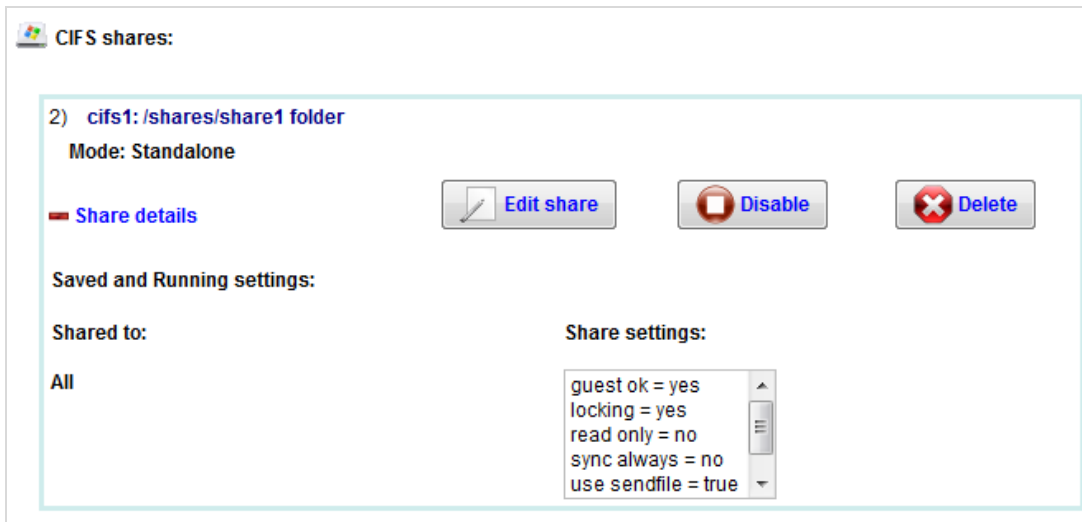


Figure 10.3(h) – CIFS share with details and actions (non-HA mode)





Figure 10.3(i) – HTTP share with details and actions (non-HA mode)



Figure 10.3(j) – FTP share with details and actions (non-HA mode)

 **NFS shares:**

1. **/shares/share1**
Mode: Clustered [clust1]

Share details

Saved settings:		Running settings:	
Shared to: 10.193.184.0/22 Type: subnet	all: caching: disabled check subtree: no coalesce writes: enabled filesystem id: 8353	Shared to: 10.193.184.0/22 Type: subnet	caching: disabled check subtree: no coalesce writes: enabled filesystem id: 8353 guest gid:

Figure 10.3(k) – NFS share with details and actions (HA mode)

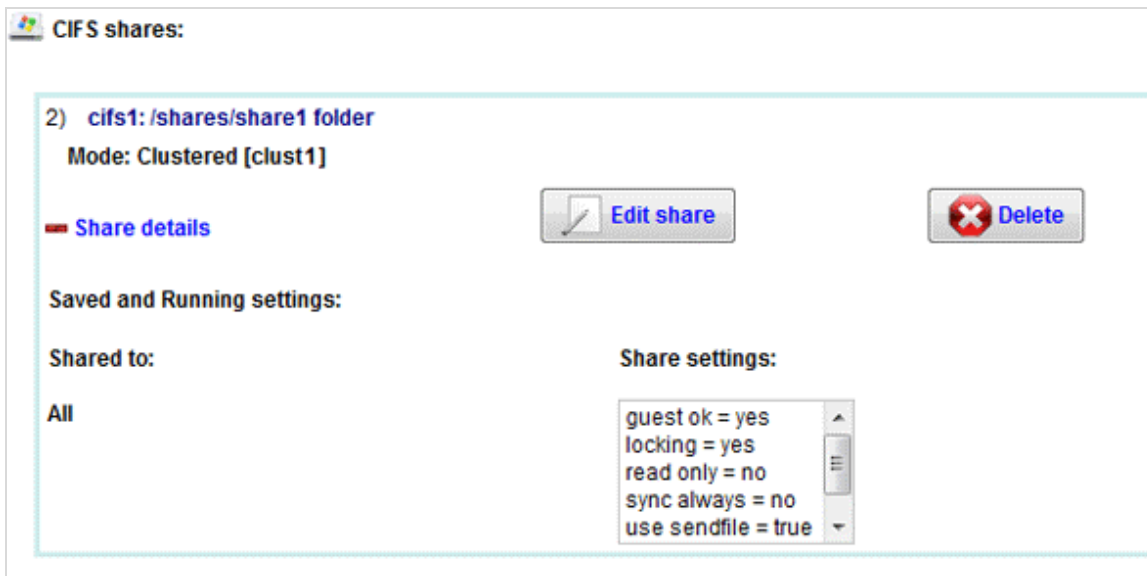


Figure 10.3(l) – CIFS share with details and actions (HA mode)

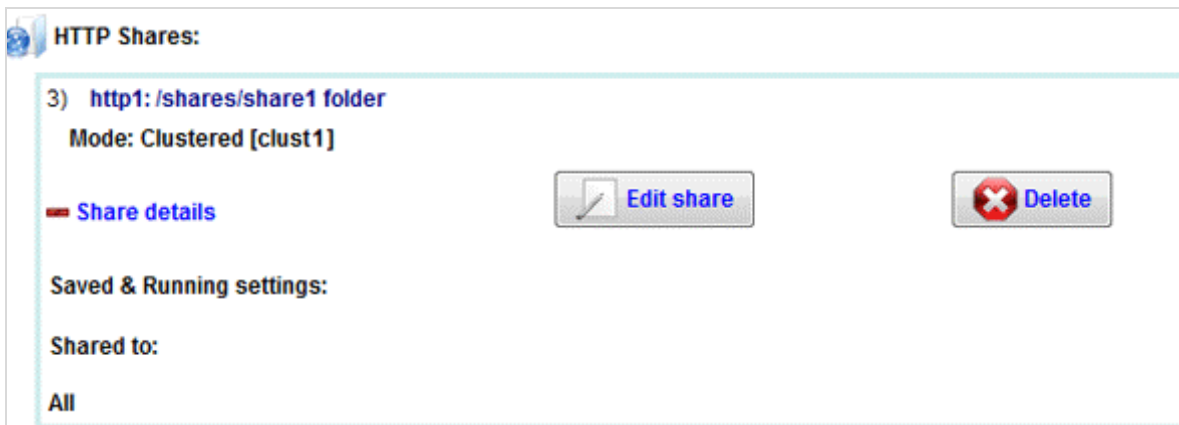


Figure 10.3(m) – HTTP share with details and actions (HA mode)

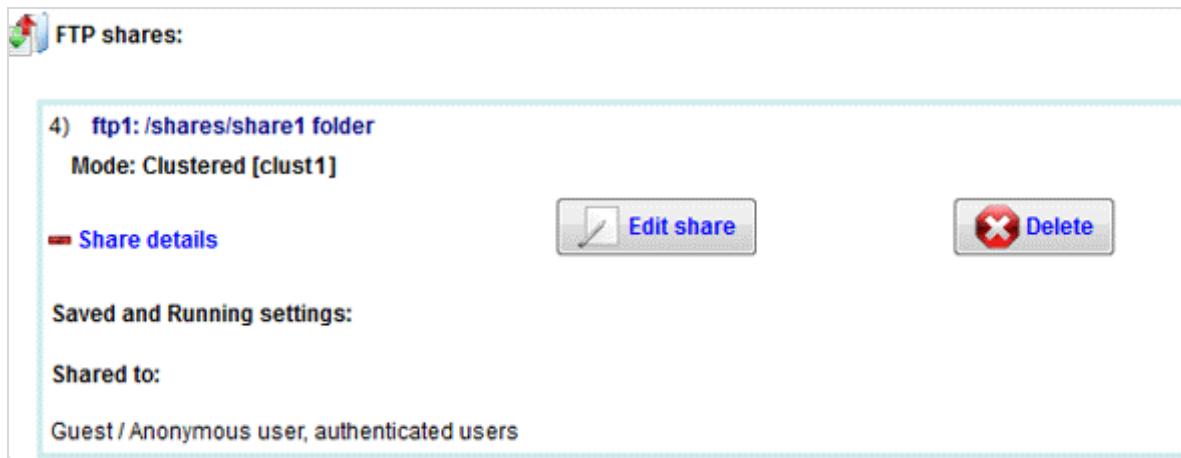


Figure 10.3(n) – FTP share with details and actions (HA mode)

2.1. Per share properties and actions

Current settings and actions for each CIFS, NFS, HTTP and FTP share are shown below the share.

2.1.1. Share Actions

- **Edit share:** This command will display a new page which allows you to edit the folder that is shared, or the share permissions and settings.
- **Enable / Disable share (only for non-HA mode):** This sets the current state of the share. If the share is enabled, it is accessible to clients; otherwise it is not accessible to clients when disabled. CIFS shares retain the enabled / disabled state across service restarts or system reboots. NFS and FTP shares will be re-enabled by default when the service is restarted or the system is rebooted.
Note: This option is not available for HTTP shares
- **Delete:** This will delete the share of that protocol. The data in the folder does not get deleted.



Warning: Disabling or deleting a share may cause clients to lose data. Ensure that no clients are currently using the share, prior to disabling or deleting it. Deleting a share cannot be undone. You will need to re-add the share. The data on the share is **not** deleted when deleting the share.

2.2. Edit share actions

Clicking on **Edit Share** leads to a new page, where user can change various options related to that particular share.

- **NFS Shares:** Here user can change various settings and permissions related to NFS share like specifying the host/subnet/netgroup to which the share is allowed access, changing Read-Write permissions, specifying the Write caching method, etc.

Edit NFS Share:

Shared to:

Change shared to:

Specific settings for 133.33.33.0/24:

Permissions:

Client Root Access level:

Client User Access level:

Write caching:

[+ More settings](#)

Figure 10.3(o) - Edit NFS share settings.

More settings

Set Anonymous user as:

Set Anonymous group as:

Hide inner mounted filesystems:

Hide filesystem if accessed from parent tree:

Secure client TCP port < 1024 required:

Lock request Authentication required:

Subtree checking:

Write coalescing:

Share folder only if mounted:

Filesystem ID (+ve number) to use:

Figure 10.3(p) - Edit NFS share additional settings.

- CIFS Shares
 - a. Share settings: Using this option, user can edit various share settings like Share name and also include comments. Share permissions like Read+Write and Read-Only can be set. User can choose to enable async or sync caching.

Share settings

CIFS Share name:
(for clients to access as \\< hostname >\< share name >)

Share comment:
(Displayed to clients when accessing the share)

Show in shares / browse list:

Permissions:

Write caching:
Note: Caching is only available if the underlying filesystem allows caching

Changes to CIFS shares will cause a service reload, interrupting any I/O in progress.

Figure 10.3(q) - CIFS Shares

- b. Access Control: Here, Administrators can allow/deny particular Client/Hosts and Users/Groups, access to the share. Enabling Guest/Anonymous access grants the guest Read and Write permissions to the share.

Access control

Guest / Anonymous access:

Only Guest / Anonymous access (Enabling this will disable regular user access and all User ACLs) allowed:

+ Clients/Hosts ACL:

+ Users / Groups ACL:

Changes to CIFS shares will cause a service reload, interrupting any I/O in progress.

Figure 10.3(r) - CIFS Access control settings

Example: Adding/Removing Access Control for CIFS shares

Here is an example on how to configure access control for a client/host.

- i. Expand the **Clients/Hosts ACL** section.
- ii. By default, any client/host with the required permission is allowed access to the CIFS share. To change this, in the **Allowed Clients** section, select the client/host from the drop-down and click on **Add allowed client**. You can also specify any other client/host not present in the drop-down by selecting *Custom value* and specifying the client/host IP or network.
- iii. Similarly, you can restrict a particular client/host from accessing the CIFS share. To do this, add the client/host from the drop-down, in the **Denied Clients** section, or use the *Custom value* option to add a client/host not present in the drop-down.

To allow a user/group access to a CIFS share, follow the steps mentioned below:

- i. Expand the **Users/Groups ACL** section.
- ii. To add a local user, enable the *Local user* radio button in the **Allowed Users/Groups** section, and select a user from the drop-down. To add an ADS user, enable the *ADS user* radio button and click on the adjacent search icon button. In the domain search pop-up, select the ADS domain and enter a valid user name. Click Apply.
Note: If ADS domain doesn't appear in the pop-up, please verify that the domain has been added correctly and that the active directory client service is running in the **Users** sub-module.

- iii. Similarly, you can also add a local or ADS group.
- iv. Click **Add allowed user**. The user/group added will appear in **Allowed Users/Groups** list.
- v. Repeat steps (ii), (iii) and (iv) to add more users/groups.
- vi. Select the user/group in the list and change read-write permissions if required.

Users/Groups can be denied access to shares in a similar way by adding them in the **Denied Users/Groups** section.

- FTP Shares: In addition to editing share name, Administrators can grant the guest Read and Write permissions to the share by Enabling Guest/Anonymous access.

FTP Share name:	<input type="text" value="home"/> (for clients to access as ftp://< hostname >/< share name >)
Guest / Anonymous access:	<input type="checkbox"/>
Only Guest access allowed:	<input type="checkbox"/>
Note: This will disable access for regular users, depending on other FTP settings.	
<input type="button" value="✔ Apply changes"/> <input type="button" value="✘ Discard changes"/>	

Figure 10.3(s) - Edit FTP share settings

- HTTP Shares: In addition to editing share name, Administrators can choose to enable/disable authentication for users. Selecting All valid Users enables Authentication for all the users in the list, whereas selecting Selected Users enables choosing which users to enable Authentication for.

Edit HTTP Share shares

HTTP Share name	<input type="text" value="shares"/> <small>(for clients to access as http:\\< hostname >\< share name >)</small>
Enable Authentication	<input checked="" type="checkbox"/>
Authentication available for:	All valid Users <input type="button" value="v"/>
Modify Users	<div style="border: 1px solid black; padding: 5px;"><p style="text-align: center;">User list (multi-select)</p><p>1. root</p></div>

Figure 10.3(t) - Edit HTTP share settings.

10.4 User Management

- Sections of the interface:

1. Local Users and Groups

The current list of local users and groups on the system is displayed here, and options to add users and groups are provided. Users or Groups can be also be modified.

Local Users and Groups:

Local Users:

The screenshot displays the 'Local Users and Groups' management window. On the left, a list of local users is shown with checkmarks next to each name. The 'root (administrator)' user is highlighted in blue. The right-hand side of the window shows the configuration details for the selected user, 'root'. The fields are as follows:

User name:	root
Full name:	<input type="text" value="root"/>
UID:	0
Primary group:	<input type="text" value="root"/>
CIFS mapped:	<input checked="" type="checkbox"/>
User enabled:	<input checked="" type="checkbox"/>
Login shell enabled:	<input checked="" type="checkbox"/>
Change password:	<input type="checkbox"/>
Set password:	<input type="text"/>
Confirm password:	<input type="text"/>

At the bottom of the configuration panel, there are two buttons: 'Delete User' (with a red X icon) and 'Save changes' (with a green checkmark icon).

Figure 10.4(a) - Local users management

1.1. Local Users and Groups settings

1.1.1. Local users list: The list of local users is given here. Clicking and highlighting a user displays the user's details, settings and actions. CIFS mapping is required for the user to login and access CIFS shares. Enabling the login shell will allow a user to access the local terminal prompt, which is not recommended. The user's password can be assigned here. The user can also be disabled or deleted if required.

1.1.2. Add local user: Specify the user's name, group, and password. The login shell is disabled by default, but it can be optionally enabled.

Note: If your USS appliance is part of an NIS domain, adding a local user with the same name as NIS user will result in an error.

1.1.3. Local groups list: The list of local groups is given here. Clicking and highlighting a group displays the group's details, settings and actions. Users can be added or removed from the group. NIS or ADS domain users can also be added to a local group, to allow domain users to access a local file or directory with the same level of access as a local user. A user cannot be removed from a group, if that group is the user's primary group.

1.1.4. Add local group: Specify the group's name and add the group.

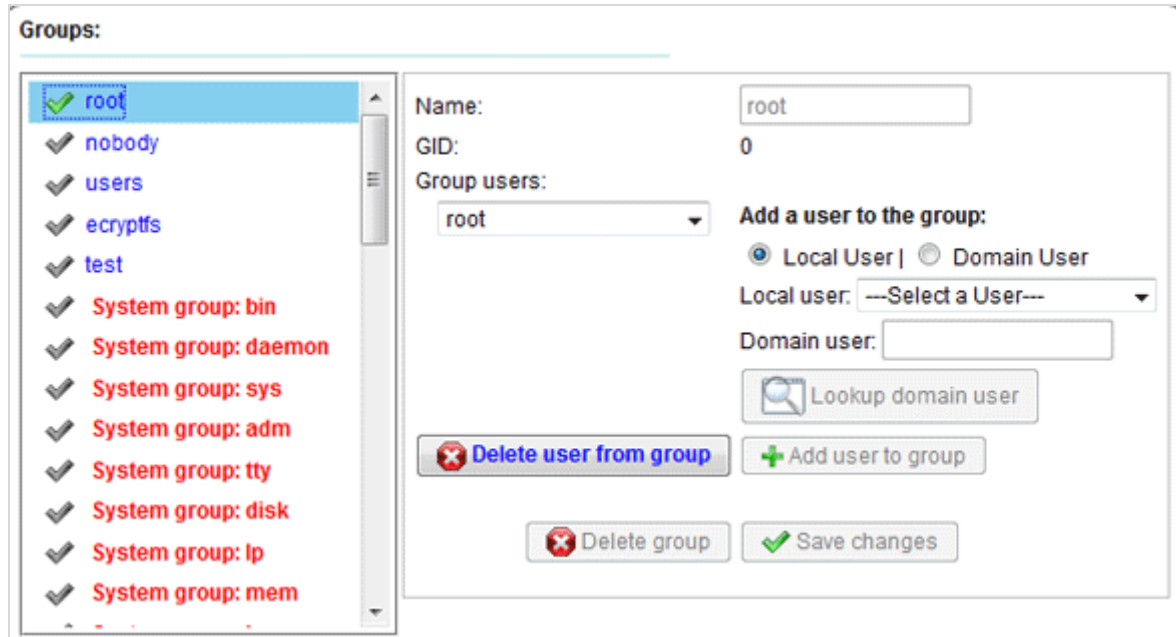


Figure 10.4(b)- Managing groups

Example: Adding a Local user

- i. In the **Users** section, under **File sharing**, expand the **Local Users and Groups**.
- ii. Two users, *root* and *nobody*, are created by default. You can select them and change properties like password and enable CIFS mapping (required for the user to login and access CIFS shares).
- iii. In the **Add a local user** section, provide the user name.
- iv. Other parameters are optional. If not specified, **users** will be assigned as the primary group for the newly created user. Enabling the login shell will allow a user to access the local terminal prompt, which is not recommended. If you want to enable authentication for the user, specify a password in the **Set password** and **Confirm password** fields.
- v. Click **Add User** to add the user. The newly added user will appear in the **Local Users** list.
- vi. You can edit the above configuration and more options by selecting the user in the list.

Note: In HA mode, adding a local user will automatically replicate that user at the peer node.

2. NIS domain membership

Network Information Service (NIS) domain authentication allows for the appliance to be part of an NIS domain, and access the NIS user database, and authenticate users based on their NIS credentials. This allows for centralized user management. NIS credentials are transmitted over plaintext and are inherently insecure.

The current status and configuration settings of the NIS domain client is shown here. There are options to control the client service and to configure domain settings. Use the “join NIS domain” menu option to configure an NIS domain.

2.1. NIS domain client commands and settings

2.1.1. Service control: Enable / Disable, Restart: These commands will affect the currently running service and allow you to decide if it should start on system startup. If you are using NIS for authenticating users on your UNIX / Linux network, you should enable the service to start automatically.

2.1.2. Current NIS domain status and settings: The currently joined NIS domain (if any), will be shown here. You can change the settings for the domain, such as the domain server to use. The domain can also be removed from the configuration.



Note:

- NIS domain details will also be displayed in the File sharing page under NIS details.

- In HA mode, after adding NIS domain, it will automatically replicate at peer node.

- NIS domain membership:

Status: Installed, enabled, domain configured, running

Actions:  **Disable**  **Restart**

Current NIS settings: Domain: **nis.uss.com**
Server: **nalanda**

Saved NIS Domain: Domain: nis.uss.com
NIS Servers: 10.193.184.188
Default Domain: Yes

Domain:

Manual Server configuration:



 **Apply changes**  **Delete Domain**

Figure 10.4(c) - NIS domain client, with a NIS domain configured.

3. ADS domain membership

Active Directory Service (ADS) domain authentication allows the Appliance to join an ADS domain, and authenticate users based on their ADS credentials.

The current status and configuration of the Active directory domain is shown here. There are options to control the Active directory client service and to remove domain membership. Use the “join ADS domain” menu option to configure the ADS domain.

3.1. ADS domain client commands and settings

3.1.1. Service control: Enable / Disable, Restart: These commands will affect the currently running service and allow you to decide if it should start on system startup. If you are using ADS for authenticating users on your Windows network, you should enable the service to start automatically.

3.1.2. Current ADS domain status and settings: The currently joined ADS domain (if any), will be shown here. You can change the settings for the domain or leave the joined domain, after which you can delete the domain.

Note:



- ADS domain details will also be displayed in the File sharing page under CIFS details.
- In HA mode, after adding ADS domain, it will automatically replicate at peer node.



Warning: Do not stop or disable domain client services if they are currently configured and running. Doing so will affect all file sharing clients.

— ADS domain membership:

Status: Installed, enabled, domain configured and joined, running

Actions:  

Current ADS settings

Domain:	ACT.COM
Server:	veera.act.com

Saved ADS Domain:

Domain:	ACT.COM
Domain Controller:	veera.act.com
Default Domain:	Yes

Domain:

Domain Controller:

Leave Domain now:




Figure 10.4(d) - ADS domain client, with ADS domain configured.

Example: How to delete ADS domain

Follow the steps mentioned below to delete ADS domain.

- i. In the **ADS domain membership** section, click on the *Leave Domain now* checkbox located at the bottom right.
- ii. Enter the Administrative username and password.
- iii. Click **Apply changes**. The Administrative credentials will now be verified and the page will refresh automatically. The status box will display errors if any.
- iv. Expand the **ADS domain membership** section again. If no error was reported, you should now only see the *Domain* and *Domain Controller* name at the bottom right.
- v. Now, click the **Delete Domain** button.

The ADS domain is now deleted. You can verify this by accessing the **ADS domain membership** section. You should get the page to join an ADS domain.

Note: In HA mode, deleting ADS domain on local node will also delete it from the peer node.

11. System Tools

11.1 System tools usage


The System Tools section of the Management Interface provides maintenance commands such as a shutdown or restarting the appliance. The date, time, and time zone can be set, and network time synchronization using the Network Time Protocol (NTP) can be enabled. NTP is recommended if there are any file-sharing services running on the system. It guarantees the accuracy of file timestamps. The modules present in this section are as follows:

1. Logs

System Logs for iSCSI or NAS services can be viewed or downloaded in the Log viewer. There are options to view three different log files and the ability to filter through them by using a standard set of criteria or by using a custom text search.

To use, choose a system log file to view and optionally use the filter criteria before selecting the

View Log button.



Choose a log to view:

Log Filter:

Filter logs using custom text:

Number of lines to display per page:

Figure 11.1(a) - Log selection criteria

1.1. Logs Options settings

- Choose a log to view: There are three primary log files that can be viewed. The drop down box controls which one:
 - i. Current kernel log – equivalent to Linux’s dmesg
 - ii. /var/log/messages
 - iii. IPMI BMC System Event Log

The two action buttons to the right of View Log button operate on the selected log file:

- i. Download Log File – Downloads the entire log file to the viewing pane.
- ii. Clear Log File – Deletes the selected log file.

The actions on the log files are not limited to these two. Log filtering can be used on the selected log file with the next three controls.

- Log Filter: To help find relevant portions of the log file, pre-selected filter criteria can be selected from the drop down dialog box. Choose the relevant filter that fits.

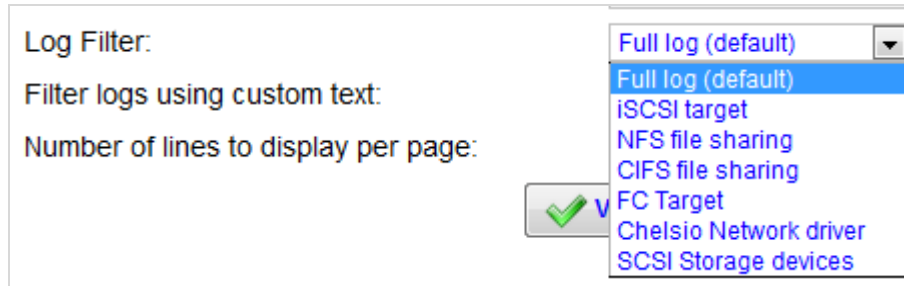


Figure 11.1(b) - Filter criteria for log files

The filter criteria choices details are:

- i. Full log (default) – No filter applied
 - ii. iSCSI target – Shows only the iSCSI target log entries
 - iii. NFS file sharing – Shows only the NFS file sharing logs
 - iv. CIFS file sharing – Shows only the CIFS (Windows) file sharing logs
 - v. FC Target – Shows only FC Target log entries
 - vi. Chelsio Network driver – Shows only the Chelsio driver logs
 - vii. SCSI Storage devices – Shows only the storage devices logs
-
- Filter logs using custom text: The text entered in the text box will be used to display only lines of the log file that contain an exact match.
 - Number of lines to display per page: The number entered on this line will limit the displayed number of lines of the log file.

2. Performance monitoring

Performance monitoring is possible, with monitors for Processor Usage, Memory usage, Network and Disk Throughput, Processor List and Cache Statistics.

To use, choose the desired category to be monitored, refine the choice with available options, and then Click the Add button. The performance data in a line-graph format will be displayed below. Furthermore, the user can chose to view performance data from multiple categories simultaneously. To do that, choose another category and options (if any) and click the Add button. The relative performance data will get added in the bottom in the display area. A particular performance data can also be removed by clicking on the close icon located on the right side of each graph.

In case of Cache Statistics, performance data of only one storage pool can be viewed at a time. To view data for a different pool, close the graph that is already running and select the required pool from the **Options** drop-down and click **Add**.

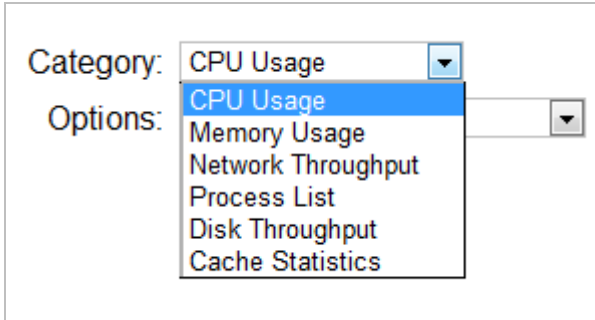


Figure 11.1(c) - Performance Monitoring selection criteria

2.1. Category: This allows for the selection of the parameter to be monitored. Choices here include:

- CPU Usage
- Memory Usage
- Network Throughput
- Process List
- Disk Throughput
- Cache Statistics

2.2. Options: To refine the parameter to be monitored, options can be used, which are set through drop down dialog boxes. The options content for each category is different depending on the parameter chosen. A table that describes each is as follows

Category	Options	Description
CPU Usage	Average CPU Usage	Average percentage usage of all the CPUs present in the system.
	CPU and all processors usage	Percentage usage of individual CPUs present in the system along with the average usage.
Memory Usage	N/A	Percentage of memory used.
Network Throughput	Network device	Rx, Tx and Bi-directional Throughput of the selected Network interface.
Process List	N/A	List of all the processes running in the system.
Disk Throughput	Disk device	Read and Write Throughput of the selected Physical Disk.
Cache Statistics	RAM/SSD cache enabled pool	Percentage of Cache hit, miss and usage for the selected pool




Figure 11.1(d) – Graph displaying Average Processor usage and Memory usage


3. System configuration data

The configuration of the appliance can be backed up from here, and it is recommended to backup the system configuration, and download it and store it in a secure location, once the appliance is fully configured. The Reset configuration to installation defaults option will reset all configuration data to the defaults from the installation of the software.

3.1. Restore System Configuration: To restore a previously saved backup configuration data file, enter the path and file name in the text box. Alternatively, navigate to the file using the browser pop-up box from pressing the Browser button. Then press the Upload button. The Reset configuration to installation defaults option will reset all configuration data to the defaults from the installation of the software. It will not cause any loss of data stored in any volumes or physical disks, but it will remove any configurations that were made till now, in any part of the software. A system restart is required to restore all of the system settings.

Backup System configuration:

Download the backup of System configuration data as of now (Tue Feb 15 16:22:46 2011):

 **Download**

Restore System configuration:

Upload a previous backup of System configuration data:

Reset configuration to installation defaults:

Note: Restoring configuration from a backup or to installation defaults will not affect any data on any data volumes, but it will stop all services and restart the system.

Note: A system restart is required to apply all restored settings correctly!

Figure 11.1(e) - System backup and restore of configuration data

4. System software update

The System update subsection allows for updating the system software (Unified Storage itself) with any upgrade packages or latest version of USS ISO image itself, provided by the vendor of the appliance. Once the system update file/ISO is obtained, navigate to the file using the browser pop-up box from clicking the **Browse...** button. Then click the “Upload System Update file and apply update” button. The update status can be seen on the screen. System services will not be disrupted during the ISO upgrade process. Once the update is completed a system restart is required.

In case of HA systems, follow the procedure mentioned below to upgrade:

- i. Fail-over all services to node2.
- ii. Upgrade node1 using the ISO provided.
- iii. Reboot node1.
- iv. Fail-back all services to node1.
- v. Upgrade node2 using the ISO provided.
- vi. Reboot node2.

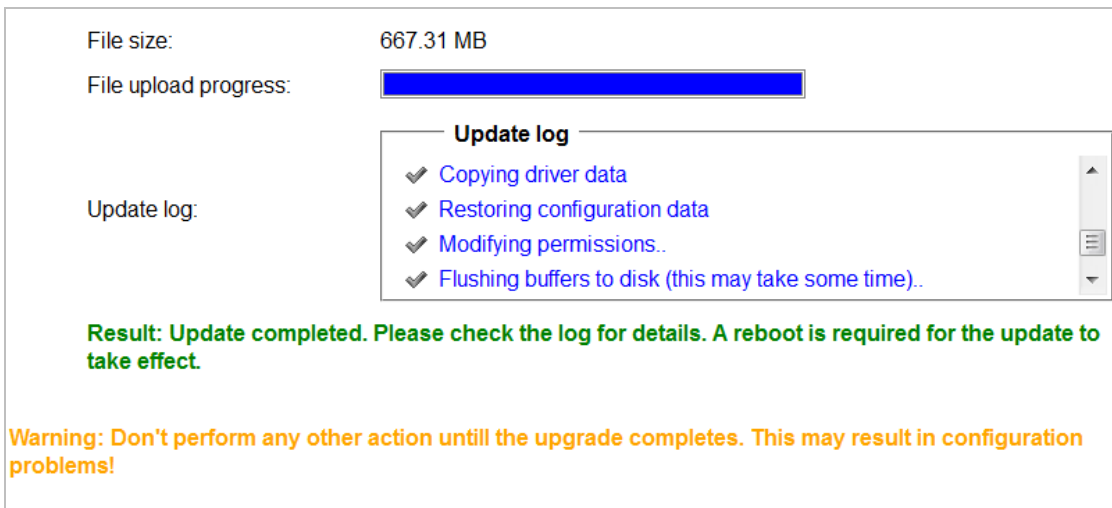


Figure 11.1(f) - System software update

Note:

- This feature is supported in USS version 2.2.0 and above. For appliances running with USS lower than 2.2.0, please boot your system using the ISO image CD/DVD and use the *update install* option or run a fresh installation. Refer the **Quick Start Guide** for more information.

- A minimum 2GB of free system memory is required for a successful update.



Warning: Do not navigate away from the page while the upgrade is in progress. This will abort the process. Also, performing any other action during upgrade might cause configuration issues. Hence, the appliance should be left idle till you see the successful upgrade confirmation message.

5. Virus scanning

USS uses the ClamAV antivirus to detect and remove file based viruses.

5.1. Scanning Folders/Directories for Viruses

- The Browse button can be used to locate the folder/directory to be scanned.
- The Add Folder to Scan button adds the selected folder to the list of folders to be scanned. Step 1 can be used to add more folders. The Scan button scans the listed folders for file based viruses.

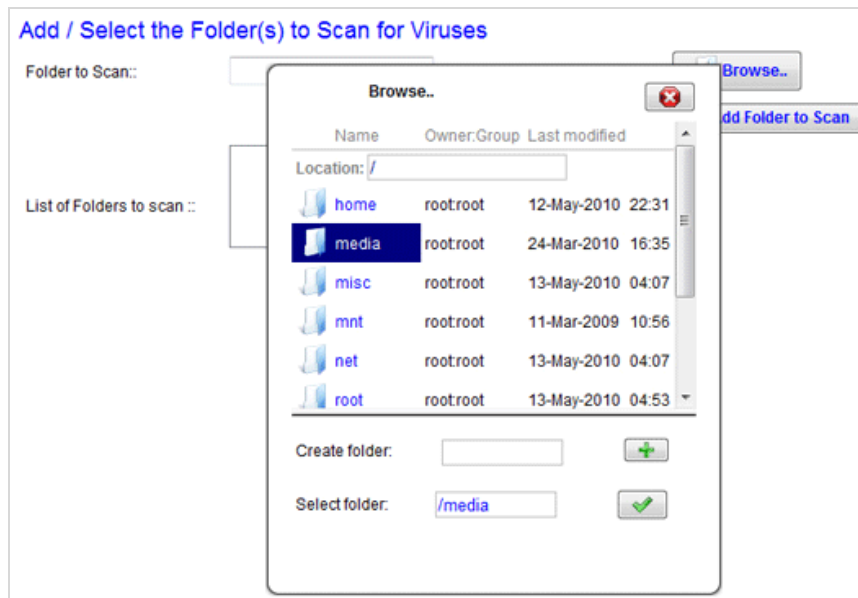


Figure 11.1(g) - Adding folders to virus scan

5.2. Scan Result: This section displays the summary of the scan including Scanned paths, Number of directories and files scanned.

5.2.1. Scan Result Commands

- Clear Results: This command resets all the values including the Scan Result and the infected files found during the scan.
- Delete selected files: Infected files detected during the scan can be removed by using this command. Users can choose the files to be deleted by selecting the check box adjacent to the file name.




Figure 11.1(h) - Deleting viruses

Example: Scanning your system for viruses

Follow the steps mentioned below to scan you system for viruses and infections.

- i. Use the **Browse** button to locate the folder/directory to be scanned.
- ii. Click the **Add Folder to Scan** button to add the selected folder to the list of folders to be scanned.
- iii. Repeat steps (i) and (ii) to add more folders.
- iv. Now click the **Scan** button. USS will now scan the listed folders for file based viruses. You can stop the scan anytime using the **Stop Scan** button.
- v. The result of the scan will be displayed including scanned paths, number of directories and files scanned. If any infected files were detected during the scan, they will be listed here.
- vi. Choose the infected files to be removed by clicking on the corresponding check box and click Delete selected files.

Clear the scan results and reset the values using the **Clear Results** button.

5.3. Update Configuration

You can use the Update Configuration section to update virus definitions manually or configure automatic updates on a daily or weekly basis.

For manual update, download the virus databases (*main.cvd* and *daily.cvd*) from ClamAV official website and upload it. To schedule automatic update, click on *Enable Automatic Update* and select either *daily* or *weekly*. If you opted for daily updates, you can configure the time at which the virus database will be updated. For weekly update, choose the days on which you want the database to be updated and also set the time for update. If your USS appliance uses a proxy server to connect to the internet, you will need to specify the settings here.

Update Configuration

Clam Anti Virus Version: 0.97.5
 Virus Database Version: 16448

Manually update the virus database

Upload daily.cvd file:

Upload main.cvd file:

Configure automatic update of database

Enable Automatic Update:

Daily Weekly

Scheduled days: Su Mo Tu We Th Fr
 Sa

Schedule Time: 04 : 00 AM

Enable Proxy Settings:

HTTP Proxy Server:

HTTP Proxy Port:

HTTP Proxy User:

HTTP Proxy Password:

Figure 11.1 (i) - Configuring manual & automatic virus updates

6. Hardware Profile Information

The Hardware Profile Information page displays information regarding the USS appliance's hardware like processor, memory, BIOS, and USS software installed. You can download the information and also verify the hardware platform.

1.1. Download Complete Hardware Profile Info

Use this button to download the complete hardware profile information for troubleshooting purposes. The downloaded file can be viewed using any text editor.

1.2. Verify Hardware Platform

You can check if your USS appliance's hardware was verified by Chelsio USS QA team using this button. If not, you can contact Chelsio support team with the hardware profile info file at support@chelsio.com to do so.

1.3. Download in XML Format

This button generates an XML file containing information about the system. This file needs to be sent along with the support file to the Chelsio Support team for trouble shooting any issues.




System Information:		BIOS Information:	
Vendor:	IBM	Vendor:	IBM Corp.
Model:	System x3650 M2 -[7947IDS]-	Version:	-[D6E145CUS-1.06]- 0.0
Serial Number:	99G0259	Release date:	03/01/2010
Processor:		Chassis:	
Vendor:	Intel(R) Corporation	Vendor:	IBM
Model:	Intel(R) Xeon(R) CPU E5506 @ 2.13GHz	Asset Tag:	none
Count:	2	Thermal status:	Other
Cache:	Internal Cache Level 1: 32 kB, Internal Cache Level 2: 256 kB, Internal Cache Level 3: 4096 kB	Power supply status:	Safe
Current speed:	2130 MHz	Type:	Other
Memory:		Software:	
Type:	None	Operating system:	Unified Storage Server 3.0.0-134
Count:	2048 MB: 2 No Module Installed: 14	Kernel:	Linux 2.6.18.US_v2.0 #12 SMP Sun Mar 24 22:01:34 EDT 2013
			
			

Figure 11.1(j) – Hardware Profile Information

7. IPMI BMC

The IPMI BMC subsection allows you to monitor the health of the system, including multiple sensors such as temperature and fan speed. It also allows remote management of the system, including power-cycling the system remotely, and some BMCs may allow remote KVM access. The BMC has a user database, and the users can be configured here, to allow remote management.

- **Sections of the interface**

1. Summary

1.1. Summary Commands

- **Reset BMC:** You can reset the BMC Firmware using two options: Warm reset and Cold reset (Power off, on).

Summary:

BMC details:

IPMI version:	2.0
BMC firmware:	1.33
Reset BMC Firmware:	<input type="button" value="Warm reset"/> <input type="button" value="Warm reset"/> <input type="button" value="Cold Reset (Power Off, On)"/> <input type="button" value="Reset BMC"/>

(Only required if the BMC is not responsive)

IPMI drivers status:

ipmi_devintf	loaded
ipmi_si	loaded

BMC sensors status:

0_75_VTT_PG	0x0 discrete
1_05V_PG	0x0 discrete
1_2_AUX_PG	0x0 discrete
1_8V_PG	0x0 discrete
1_8_PLL_PG	0x0 discrete
3_3V_PG	0x0 discrete
Ambient_Temp	20.000 degrees C
CMOS_Battery	0x0 discrete
CPU_Bus_PERR	na discrete
CPU_Init_Err	na discrete
CPU_Machine_Chk	na discrete
CPU_Protocol_Err	na discrete

Figure 11.1(k) - IPMI BMC Summary

2. BMC Configuration

2.1. Chassis

- Power restore policy determines how the system or chassis behaves when AC power returns after an AC power loss. Available options are always-off, always-on and previous.

Chassis:

Power restore policy:
(after power loss)

previous
always-off
always-on (recommended)
previous


A screenshot of a BMC configuration interface. At the top, the word "Chassis:" is displayed in bold. Below it is a horizontal line. The main configuration area shows the label "Power restore policy:" followed by "(after power loss)" in a smaller font. To the right of this label is a dropdown menu. The dropdown menu is currently open, showing four options: "previous", "always-off", "always-on (recommended)", and "previous". The "always-on (recommended)" option is highlighted with a blue background. The dropdown menu has a small downward-pointing arrow on its right side.

Figure 11.1(l) - Power restore policy

2.2. IPMI Over LAN

In this section you can set up IPMI over LAN either by DHCP or Static IP

IPMI over LAN:

MAC address: **00:25:90:14:2e:ba**

IP address type:

IP address:

IP netmask:

IP gateway:

VLAN: Enabled: ID: (1 - 4094)

Figure 11.1 (m) - Power restore policy

2.3. IPMI Users

Users with various Privilege levels can be setup here. A user can be setup using custom settings or using default values. You can also edit various user specific settings using the "Edit user" section.

IPMI users:

Maximum users: 10

Users:

- ID: 1 - User: not configured - Level: not configured
- ID: 2 - User: root - Level: administrator
- ID: 3 - User: uss - Level: no access
- ID: 4 - User: not configured - Level: not configured**
- ID: 5 - User: not configured - Level: not configured

User ID: 4

User name: user_288u

Set password:

Password: Show password

Privilege level: User

Callin / Callback: callback

Link authentication: Enabled

IPMI messaging: Disabled

User status: Disabled

Figure 11.1(n) - IPMI Users settings

8. UPS Management Controller

Using this utility, you can monitor and administer your UPS hardware.

- **Sections of the interface**

1. Summary

The summary section displays information regarding UPS hardware like model number, battery charge status, input and output voltage, serial port device id, etc.

UPS Details:	
Device model:	PW5115
Status :	OL
Battery charge:	95
Output voltage:	115
Input voltage:	115
Device serial:	GF236A0348
Port:	/dev/ttyS0
Beeper status:	enabled

Figure 11.1(o) - UPS hardware summary



2. UPS Configuration

In this section, you can view and configure UPS management service status. You can start, stop or restart the service. The **Enable** button will start the service as well as configure it to start during boot up (auto-start). The **Disable** button will stop the service and disable auto-start.

To change the serial port, plug the serial cable into the new port on the appliance and then select the corresponding port device id from the **Change Serial Port** drop-down list. Next, Click **Apply**. You will have to restart the appliance for changes to take effect.

UPS Configuration:

Service status: **installed, auto start enabled, Running**

Action :  **Disable**  **Restart**


Change Serial Port :  **Apply**

Figure 11.1(p) - UPS Configuration

9. Date – time

The system time, date, and time zone can be configured from this system utility tool. Additionally, system time can be synchronized with a network time server.

To change the date, time, time zone, and the network time synchronization, configure the settings in the user interface. The details on each are described further below. Once all of the settings are configured, press the Apply button for them to take effect.

Description	Current setting	Change Setting
Date:	08 April 2013	New Date: <input type="text" value="08 April 2013"/> <input type="button" value="Edit"/>
Time:	20:11:48	New Time: <input type="text" value="20"/> : <input type="text" value="11"/> : <input type="text" value="48"/> (hh:mm:ss)
Timezone:	Asia - Kolkata +0530 IST	(UTC/GMT offset Code Continent/City) <div style="border: 1px solid black; padding: 5px;"> <p style="text-align: center;">Timezone</p> <ul style="list-style-type: none"> ✓ -1100 WST Pacific/Apia ✓ -1100 SST Pacific/Pago_Pago ✓ -1100 SST Pacific/Midway ✓ -1100 NUT Pacific/Niue ✓ -1000 TKT Pacific/Fakaofu ✓ -1000 TAHT Pacific/Tahiti ✓ -1000 HST Pacific/Johnston ✓ -1000 PST Pacific/Honolulu </div>
Network Time synchronization:	<input type="button" value="Enabled"/> Currently running Sync with 10.193.182.23	NTP servers: <input type="text" value="10.193.182.23"/> <input type="button" value="Remove server"/> Add custom NTP server hostname / IP : <input type="text"/> <input type="button" value="Add"/>

Figure 11.1 (q) - System time / date control

2.4. Date

The date can be manually changed through the pop-up calendar. To invoke it, press the Edit button. Then navigate to the current month/day and press Apply. This will fill in the New Date field under the Change Setting column.

Note: Option to change date is available only if *Network Time synchronization* is disabled.

Description	Current setting	Change Setting
Date:	08 April 2013	New Date: 08 April 2013 <input type="button" value="Edit"/>
Time:	20:13:17	New Time: 20 : 13 : 17
Timezone:	Asia - Kolkata +0530 IST	(UTC/GMT offset Code Continent) Timezone <ul style="list-style-type: none">✓ -1100 WST Pacific/Apia✓ -1100 SST Pacific/Pago_P✓ -1100 SST Pacific/Midway✓ -1100 NUT Pacific/Niue✓ -1000 TKT Pacific/Fakaofu✓ -1000 TAHT Pacific/Tahiti✓ -1000 HST Pacific/Johnston✓ -1000 HST Pacific/Honolulu

Select the Date

April

Su	Mo	Tu	We	Th	Fr	Sa
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30				

2013

Figure 11.1(r) - Calendar pop-up

2.5. Time

The time can be changed by setting the dropdowns for the hour, minute, and second (hh:mm:ss) under the Change Setting column.

Note: Option to change time is available only if *Network Time synchronization* is disabled.

2.6. Timezone

The timezone can be changed by navigating to and selecting the Code or Continent/City in the selection list.

2.7. Network Time synchronization: Using this option, user can synchronize their system time by connecting to a NTP server. More than one custom NTP servers can be added.

10. Alerting

Email alerting is provided for different types of failures or events that may occur on the system. Set the email recipient list and the SMTP server address, and choose the events that should be alerted, in the Email alerts page. If you have an SNMP monitoring station, provide the IP address of the SNMP monitoring station and the port, and SNMP traps will be sent to that monitoring station with the alert information.

- **Sections of the interface:**

- 1. Summary**

- 1.1. Service details

User can view Alert service details here and Restart the Service. Autos Start is enabled by default, hence the Alerting Service will run automatically when the system starts.

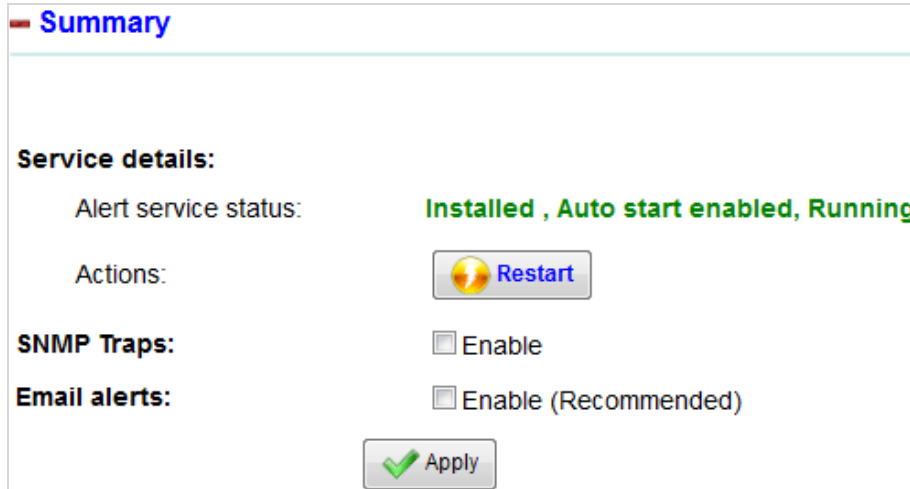


Figure 11.1(s) - Summary section with service details, SNMP and Email alerts settings

1.2. SNMP Traps

SNMP traps enable an agent to notify the management station of significant events by way of an unsolicited SNMP message. User can enable this feature by specifying the hostname /IP address and port id of management station and the events for which alerts are to be received in the Settings section.

SNMP Traps: Enable

Trap Sink Hostname / IP Address:

Trap Sink Community Port:

Trap Sink List:

Sink List	
<input checked="" type="checkbox"/>	102.55.55.55 162
<input checked="" type="checkbox"/>	102.55.55.56 162

Figure 11.1(t) - Setting up SNMP alerts

Example: Setting up SNMP traps

- i. Enable SNMP traps by clicking on the check box adjacent to “SNMP traps”
- ii. Specify the hostname /IP address and port id of the management station and click **Add**. Use this step to add more devices.
- iii. Click **Apply**.
- iv. Now under the “Settings” section, select the events for which alerts should be sent.
- v. Click **Apply**.

1.3. Email Alerts

Users can also choose to receive email alerts. More than one email id can be specified separated by commas. SMTP authentication can be enabled in which case a user name and password should be provided.

Email alerts:	<input checked="" type="checkbox"/> Enable
Email Recipients:	<input type="text" value="youremail@chelsio.com"/>
Sender Email Address:	<input type="text" value="mail@mailserver.com"/>
Mail Server:	<input type="text" value="10.192.180.15"/>
SMTP Port:	<input type="text" value="25"/>
SMTP Authentication:	<input checked="" type="checkbox"/> Enable
User Name:	<input type="text" value="user_23"/>
Password:	<input type="password" value="....."/>
Confirm Password:	<input type="password" value="....."/>
<input type="button" value="✔ Apply"/>	

Figure 11.1(u) - setting up Email alerts

Example: Setting up Email alerts

- i. Enable email alerts by clicking on the check box adjacent to “Email alerts”
- ii. Enter the email id of the recipients separated by commas.
- iii. Enter the email address from which you would like to receive the alerts in “Sender Email Address” in the format of<name>@<domain>.<extension>.
- iv. Enter the IP or host name of the mail server and SMTP port number.
- v. If your SMTP server requires authentication, then enable SMTP authentication here and provide the username and password.
- vi. Click **Apply**.
- vii. Now under the **Settings** section, select the events for which alerts should be sent.
- viii. Click **Apply**.

2. Settings

Using the Settings section, user can customize various options, based on which SNMP traps/Email alerts will be send.

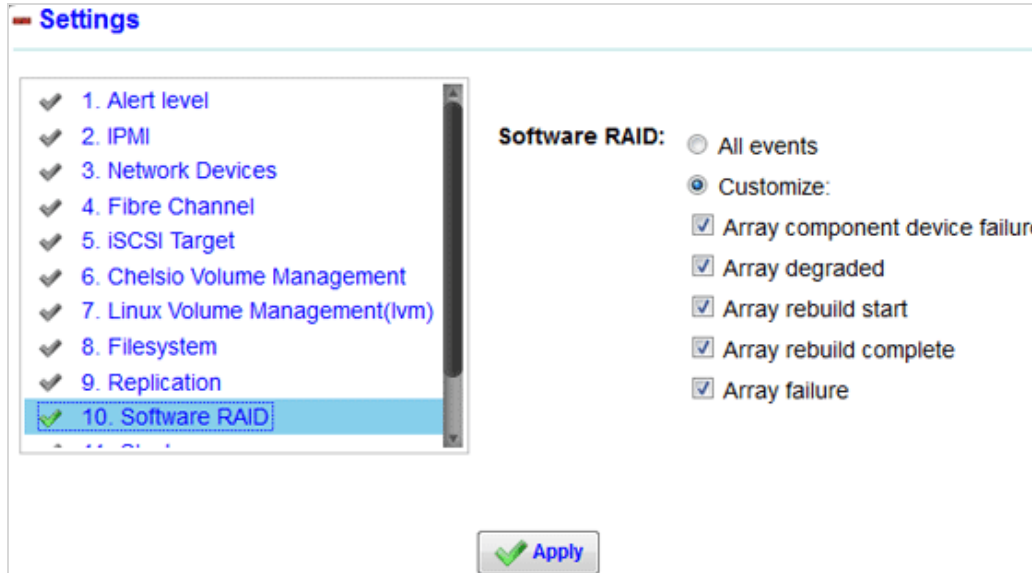


Figure 11.1 (v) - configuring alert settings

The options available are dependent on the hardware and the features licensed.

- i. *Alert level* will enable users to receive three kinds of alerts: Information (all notifications), Warnings (warnings and errors) and Errors. You can decide the time interval for synchronous events (filesystem size, COW snapshots, filesystem quota, etc) at which the alerts will be gathered and sent by specifying the poll time.
- ii. *IPMI* provides alerts on system health.
- iii. *Network Devices* alerts will alert the user on a device losing or acquiring an Ethernet link.
- iv. *Fibre Channel* alerts will alert the user on a device losing or acquiring the link on the FC port.
- v. *iSCSI target* alerts for an initiator login or logout, and security events such as an initiator ACL deny or CHAP authentication failure.
- vi. *Chelsio Volume Management* alerts will alert the administrator when a thin provisioned pool is running low on physical disk space, where the amount of allocated space is greater than the physical disks capacity.
- vii. *Linux Volume management* alerts are for snapshots space running low, which will cause the snapshot to become invalid. Using Linux Volume management is not recommended.

- viii. *Filesystem* alerts provide details of any filesystems that are running low on space, and if any users are crossing their assigned quotas for a filesystem.
- ix. *Replication* provides multiple alerts for different replication states and error conditions.
- x. *Software or Hardware RAID (SAS Raid)* alerting is provided for RAID array degrade / failure and rebuild events.
- xi. *Cluster* provides alerts when cluster membership changes.
- xii. *SMART disk monitoring* provides alerts on the health of any physical disks that support SMART technology in the disk firmware, and for which monitoring has been enabled.

11. Administrator password

The administrative password for the root account can be changed. To do this, simply enter the old and new password in the respective text boxes and repeat the new password in the *confirm password* text box and then press the “Change password” button.



The image shows a web form for changing the administrator password. It consists of three text input fields stacked vertically, each with a label to its left: "Old password:", "New password:", and "Confirm password:". Below these fields is a button with a green checkmark icon and the text "Change password". The entire form is enclosed in a light gray border.

Figure 11.1 (w) - Administrator’s account password change

12. Shutdown, restart

Halting and restarting the system immediately, or after a certain delay is also available in the interface. Either action can occur immediately if desired or scheduled up to 3 hours in advance. To execute on either a shutdown or reboot, dial in the amount of minutes in the drop down dialog box (00 minutes being immediate), and press the appropriate button for either Halt / Shutdown or Reboot / Restart. If it's not an immediate action (set to 00 minutes), then either action can be canceled by pressing the "Cancel scheduled Shutdown" button. In HA mode, command to shutdown or reboot peer node from the local node is available.



Figure 11.1 (x) - System Control for shutdown or reboot



Figure 11.1 (y) - System Control for shutdown or reboot peer node in HA mode

The granularity of the amount of minutes that the shutdown or reboot can be scheduled is as follows:




Figure 11.1 (z)- Shutdown / Reboot scheduling time

12. Appendix

12.1 Troubleshooting

If your USS appliance is not functioning as expected, you are requested to go through this section which addresses USS related issues and their solutions, before contacting the support team. You are also advised to check README and Release Notes which contain distribution specific problems included in this release and possible workaround.

1. **While adding a local user in the Local Users and Groups section, I get the following error “*Adding user..failed! Details: useradd: user user1 exists*” even though a local user with that name doesn’t exist.**

This error occurs if your USS appliance is part of an NIS domain, and the name of the local user you’re trying to add already belongs to an NIS user. To resolve this, access your USS appliance’s CLI and execute the following command:

```
[root@host]# ypcat passwd
```

The above command will list all the users in NIS domain database. Now try creating a local user again but with a name which doesn’t appear on the list.

2. **While running I/O operations, USS appliance was rebooted abruptly (e.g. due to power loss). Now the drive on which operations were running is not listed in the *Disk devices* section anymore.**

This is a kernel behavior and may not occur every time. If it does, follow these steps to resolve it: reseal the drive on which I/O operations were running. Next, access the **Disk devices** section under **Storage**, and use the **Rescan all storage devices** button. The drive should appear in the list again.

Note: It is highly recommended that the appliance is shutdown or restarted using the power options available in USS (System Tools->Shutdown, restart). Abrupt shutdown or reboot may cause I/O errors and data loss.

- 3. I have configured network interfaces to be used as cluster backplane correctly. However, on trying to create a cluster, the *backplane connectivity* parameter displays error (red Cross) during Cluster Compatibility Check.**

This is an expected behavior when using non-Supermicro SBB systems. If all other mandatory parameters were matched and passed (indicated by a green tick), then you can proceed to the configuration section where you will have to provide interfaces for backplane manually.

- 4. The System section displays “*This platform is not validated for clustering*”. What does it mean? Can I still create cluster?**

The Chelsio QA team ensures that USS is tested with a large number of different hardware configurations. It is possible that your server’s hardware was not verified for cluster and hence the message. However, you can proceed with the cluster creation process. To create cluster successfully, you will have to ensure that the prerequisites (mentioned in the Create Cluster page) are met.

5. The *Create Cluster* button is disabled/grayed out.

The **Create Cluster** button will be disabled if your USS appliance is not licensed with **HA** and **Chelsio Thin Provisioning (TP)** features. An appliance with a non-TP license provides only LVM (Logical Volume Management) for SAN management and HA (cluster) is only supported with TP. You should contact Chelsio support team and apply for the appropriate license.

6. I forgot the administrative (root) password. How do I reset it?

- In case of standalone (non-HA) setup, you can use the procedure mentioned below to reset the root password:
 - i. Reboot/boot-up USS machine.
 - ii. During boot, press any key when prompted to enter GRUB menu.
 - iii. Select the third option, i.e. *Chelsio Unified Storage vX.X.X (recovery)*.
 - iv. Now, at the BASH prompt, run the following command:

```
bash-3.2# mount -o remount,rw /sysroot
```

- v. Run the following command. This will change the root directory to */sysroot*

```
bash-3.2# chroot /sysroot
```

- vi. Finally, change the password:

```
sh-3.2# /usr/bin/passwd
```

- vii. Enter the new password:

```
Changing password for user root
New UNIX password:
```

- viii. Confirm the new password:

```
Retype new UNIX password:
```

- ix. Reboot for changes to take effect.

- For HA setup, the procedure of resetting the root password is slightly different. Follow the steps below:
 - i. Shutdown the secondary (peer) node.
 - ii. Reboot the primary (local) node.

- iii. Reset the password on the primary node by following the steps ii-viii mentioned for non-HA setup.
- iv. Shutdown the primary node.
- v. Boot up the secondary node.
- vi. Again, reset the password on the secondary node by following the steps ii-viii mentioned for non-HA setup. Please ensure that the password set for primary node is also set for the secondary node.
- vii. Reboot the secondary node.
- viii. Boot up the primary node.

7. I had configured my cluster as iSCSI Initiator and logged on to a remote target successfully. After some time previously discovered LUN is not seen even after trying the *Rescan this target* option.

This error occurs if the LUN is full. To resolve this, you will have to relocate the cluster service to the primary (local) node. To do this, go to the **Services** section under **Cluster**. In the **Preferred owner** drop-down list, select peer (secondary) node and click on the button with green tick. After the service relocates successfully, repeat the same step but select the primary node in the **Preferred owner** drop-down list this time.

8. I forgot master/volume pass phrase (key). How do I reset it?

You will need the corresponding recovery key to reset master or volume pass phrase. In other words, to reset master pass phrase you will have to provide the recovery key generated when the **Volume Management** module was accessed for the first time. To reset volume pass phrase, you will have to provide the recovery key generated while encrypting the volume.

To reset master pass phrase, follow these steps:

- i. Access the **Volume Management** module
- ii. Next, expand the **Settings** section.
- iii. Select the appropriate option in the **Encryption Settings** drop-down list depending on the type of pass phrase you want to reset. Click **Apply**.
- iv. Enter the new pass phrase consisting of minimum 6 characters in the *New Pass Phrase* field.
- v. Re-enter the same pass phrase again in the *Confirm Pass Phrase* field.
- vi. If you selected *Reset volume pass phrase* in step (iii), then select the pool on which the volume was created in the **Pool** drop-down.
- vii. Select the volume for which you want to reset the pass phrase.
- viii. Depending on the type of pass phrase being reset, copy and paste the correct recovery key.

- ix. Click **Change Pass Phrase**.

The pass phrase is now reset.

9. While copying data on FC Target LUN using Windows client, USS becomes unresponsive and LUN becomes inaccessible.

This error can occur if the storage pool is over-provisioned and pool hosting the FC target LUN runs out of space. To resolve this, reboot USS and client and then add more disks to the pool. To do this, go to the **Disk devices**, and select a free disk and click **Disk details**. Next, under **Device layout**, select the green section and choose *Manage with Volume Management* in the **Actions** drop-down list. Next, in the **Volume Management usage** drop-down list, select the pool on which the LUN was created. Finally, click **Apply**. Once successfully added, try copying data on client again.

You can avoid this error by keeping a watch on the disk space usage using the email-alerts feature which notifies you when pool is running low on disk space. To know more on how to configure email-alerts refer **Alerting** section under **System Tools** ([Click here](#)).

10. I created a storage pool using Micron SSDs as a cache device in a cluster. Local node fails due to abrupt shutdown but cluster service doesn't *failback* when the node is powered on.

After powering on the local node, Micron SSD disks will undergo rebuild process. Cluster service cannot failback unless this process is complete and will continue to run on the peer node in the meantime.

Once the rebuild process is complete, you can transfer the service to local node in the **Services** module available under **Cluster**.

To check the status of process, please view the dmesg entries in the Logs module available under System Tools.

Note: It is recommended that battery backup unit (like **UPS**) should be present to avoid data loss in case of system failure.

11. Service fail-over/fail-back in a cluster takes long time causing I/O errors on the client side.

When a large number of cluster services (e.g. pool, software RAID) are running, the time taken to failover or failback can exceed the default client setting of 60 seconds causing IO errors. This issue

can be resolved by increasing the disk I/O timeout value and maximum request hold time on the client side.

- **Increasing I/O timeout on Windows client**

- i. Click the **Start** button.
- ii. Depending on the Windows version, either type *regedit.exe* in the **Search box** (e.g. Windows 7) or click **Run** and then type *regedit.exe* (e.g. Windows XP). This will open the **Registry Editor**.
- iii. Locate the *HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk* registry key.
- iv. Locate the *TimeOutValue* entry. Right-click and select **Modify**.
- v. Select **Decimal** and set the *Value data* field to a value greater than 60. Click **OK**.
- vi. Reboot machine for changes to take effect.

- **Increasing I/O timeout on Linux client**

Run the following command:

```
[root@host]# echo <timeout value> > /sys/block/sdX/device/timeout
```

- **Increasing maximum request hold time on Windows client**

- i. Click the **Start** button.
- ii. Depending on the Windows version, either type *regedit.exe* in the **Search box** (e.g. Windows 7) or click **Run** and then type *regedit.exe* (e.g. Windows XP). This will open the **Registry Editor**.
- iii. Click **Computer**.
- iv. On the **Edit** menu click **Find** or use the keyboard shortcut *Ctrl+F*.
- v. In the **Find** box, type the text *MaxRequestHoldTime* and click **Find Next**.

The *MaxRequestHoldTime* entry should be located in a similar path:

```
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\0000\Parameters
```

- vi. Right-click on the entry and select **Modify**.

- vii. Select the **Decimal** option and set the *Value data* field to 180. Click **OK**.
- viii. Reboot machine for changes to take effect.

- **Increasing maximum request hold time on Linux client**

- i. Edit the iSCSI configuration file, *iscsid.conf*

```
[root@host]# vim /etc/iscsi/iscsid.conf
```

- ii. Set the value of *node.session.timeo.replacement_timeout* to 180.

```
node.session.timeo.replacement_timeout = 180
```

- iii. Save changes and exit.

iv. Restart iSCSI daemon

```
[root@host]# /etc/init.d/open-iscsi restart
```

12.2 Chelsio End User License Agreement

IMPORTANT: PLEASE READ THIS SOFTWARE LICENSE CAREFULLY BEFORE DOWNLOADING OR OTHERWISE USING THE SOFTWARE OR ANY ASSOCIATED DOCUMENTATION OR OTHER MATERIALS (COLLECTIVELY, THE "SOFTWARE"). BY CLICKING ON THE "OK" OR "ACCEPT" BUTTON YOU AGREE TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, CLICK THE "DECLINE" BUTTON TO TERMINATE THE INSTALLATION PROCESS.

1. License. Chelsio Communications, Inc. ("Chelsio") hereby grants you, the Licensee, and you hereby accept, a limited, non-exclusive, non-transferable license to install and use the Software with one or more Chelsio network adapters on a single server computer so as to function as a storage device that may be accessed by one or more other computers over a network. You may also make one copy of the Software in machine readable form solely for back-up purposes, provided you reproduce Chelsio's copyright notice and any proprietary legends included with the Software or as otherwise required by Chelsio.

2. Restrictions. This license granted hereunder does not constitute a sale of the Software or any copy thereof. Except as expressly permitted under this Agreement, you may not:

(i) reproduce, modify, adapt, translate, rent, lease, loan, resell, distribute, or create derivative works of or based upon, the Software or any part thereof; or

(ii) make available the Software, or any portion thereof, in any form, on the Internet. The Software contains trade secrets and, in order to protect them, you may not decompile, reverse engineer, disassemble, or otherwise reduce the Software to a human-perceivable form. You assume full responsibility for the use of the Software and agree to use the Software legally and responsibly.

3. Ownership of Software. As Licensee, you own only the media upon which the Software is recorded or fixed, but Chelsio retains all right, title and interest in and to the Software and all subsequent copies of the Software, regardless of the form or media in or on which the Software may be embedded.

4. Confidentiality. You agree to maintain the Software in confidence and not to disclose the Software, or any information or materials related thereto, to any third party without the express written consent of Chelsio. You further agree to take all reasonable precautions to limit access of the Software only to those of your employees who reasonably require such access to perform their employment obligations and who are bound by confidentiality agreements with you.

5. Term. This license is effective in perpetuity, unless terminated earlier. You may terminate the license at any time by destroying the Software (including the related documentation), together with all copies or modifications in any form. Chelsio may terminate this license, and this license shall be deemed to have automatically terminated, if you fail to comply with any term or condition of this Agreement. Upon any termination, including termination by you, you must destroy the Software (including the related documentation), together with all copies or modifications in any form.

6. Limited Warranty. If Chelsio furnishes the Software to you on media, Chelsio warrants only that the media upon which the Software is furnished will be free from defects in material or workmanship under normal use and service for a period of thirty (30) days from the date of delivery to you.

CHELSIO DOES NOT AND CANNOT WARRANT THE PERFORMANCE OR RESULTS YOU MAY OBTAIN BY USING THE SOFTWARE OR ANY PART THEREOF. EXCEPT FOR THE FOREGOING LIMITED WARRANTY, CHELSIO MAKES NO OTHER WARRANTIES, EXPRESS OR IMPLIED, AND HEREBY DISCLAIMS ALL OTHER WARRANTIES, INCLUDING, BUT NOT LIMITED TO, NON-INFRINGEMENT OF THIRD PARTY RIGHTS, MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow the exclusion of implied warranties or limitations on how long an implied warranty may last, so the above limitations may not apply to you. This warranty gives you specific legal rights and you may also have other rights which vary from state to state.

7. Remedy for Breach of Warranty. The sole and exclusive liability of Chelsio and its distributors, and your sole and exclusive remedy, for a breach of the above warranty, shall be the replacement of any media furnished by Chelsio not meeting the above limited warranty and which is returned to Chelsio. If Chelsio or its distributor is unable to deliver replacement media which is free from defects in materials or workmanship, you may terminate this Agreement by returning the Software.

8. Limitation of Liability. IN NO EVENT SHALL CHELSIO HAVE ANY LIABILITY TO YOU OR ANY THIRD PARTY FOR ANY INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR PUNITIVE DAMAGES, HOWEVER CAUSED, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR RELATED TO THE LICENSE OR USE OF THE SOFTWARE, INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR LOSS OF ANTICIPATED PROFITS, EVEN IF CHELSIO HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL CHELSIO'S LIABILITY ARISING OUT OF OR RELATED TO THE LICENSE OR USE OF THE SOFTWARE EXCEED THE AMOUNTS PAID BY YOU FOR THE LICENSE GRANTED HEREUNDER. THESE LIMITATIONS SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

9. High Risk Activities. The Software is not fault-tolerant and is not designed, manufactured or intended for use or resale as online equipment control equipment in hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life support machines, or weapons systems, in which the failure of the Software could lead directly to death, personal injury, or severe physical or environmental damage.

Chelsio specifically disclaims any express or implied warranty of fitness for any high risk uses listed above.

10. Export. You acknowledge that the Software is of U.S. origin and subject to U.S. export jurisdiction. You acknowledge that the laws and regulations of the United States and other countries may restrict the export and re-export of the Software. You agree that you will not export or re-export the Software or documentation in any form in violation of applicable United States and foreign law. You agree to comply with all applicable international and national laws that apply

to the Software, including the U.S. Export Administration Regulations, as well as end-user, end-use, and destination restrictions issued by U.S. and other governments.

11. Government Restricted Rights. The Software is subject to restricted rights as follows. If the Software is acquired under the terms of a GSA contract: use, reproduction or disclosure is subject to the restrictions set forth in the applicable ADP Schedule contract. If the Software is acquired under the terms of a DoD or civilian agency contract, use, duplication or disclosure by the Government is subject to the restrictions of this Agreement in accordance with 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors and 49 C.F.R. 227.7202-1 of the DoD FAR Supplement and its successors.

12. General. You acknowledge that you have read this Agreement, understand it, and that by using the Software you agree to be bound by its terms and conditions. You further agree that it is the complete and exclusive statement of the agreement between Chelsio and you, and supersedes any proposal or prior agreement, oral or written, and any other communication between Chelsio and you relating to the subject matter of this Agreement. No additional or any different terms will be enforceable against Chelsio unless Chelsio gives its express consent, including an express waiver of the terms of this Agreement, in writing signed by an officer of Chelsio. This Agreement shall be governed by California law, except as to copyright matters, which are covered by Federal law. You hereby irrevocably submit to the personal jurisdiction of, and irrevocably waive objection to the laying of venue (including a waiver of any argument of forum non conveniens or other principles of like effect) in, the state and federal courts located in Santa Clara County, California, for the purposes of any litigation undertaken in connection with this Agreement. Should any provision of this Agreement be declared unenforceable in any jurisdiction, then such provision shall be deemed severable from this Agreement and shall not affect the remainder hereof. All rights in the Software not specifically granted in this Agreement are reserved by Chelsio. You may not assign or transfer this Agreement (by merger, operation of law or in any other manner) without the prior written consent of Chelsio and any attempt to do so without such consent shall be void and shall constitute a material breach of this Agreement.

Should you have any questions concerning this Agreement, you may contact Chelsio by writing to:

Chelsio Communications, Inc.

370 San Aleso Ave.

Sunnyvale, CA 94085

12.3 GNU General Public License

GNU General Public License version 2 for the Linux Kernel, Operating System and associated software

GNU GENERAL PUBLIC LICENSE
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program"

means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program).

Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place

counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution

of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE

OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc.,
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program
`Gnomovision' (which makes passes at compilers) written by James Hacker.
```

```
<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.

12.4 Other licensing information

Certain software included in the Operating System may be licensed under older or newer versions of the GNU GPL, or various other licensing terms such as the BSD license, or proprietary licenses, etc., as decided by the Authors / Owners of that particular software. All such licensing terms apply for usage of the particular software.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH CHELSIO PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN CHELSIO'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, CHELSIO ASSUMES NO LIABILITY WHATSOEVER, AND CHELSIO DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF CHELSIO PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. CHELSIO PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS. CHELSIO MAY MAKE CHANGES TO SPECIFICATIONS AND PRODUCT DESCRIPTIONS AT ANY TIME, WITHOUT NOTICE.

