

Unified Storage Server



User Guide

Installation & Basic Configuration

References

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

Copyright and Trademarks

Copyright © 2025 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., 735 N Pastoria Avenue, Sunnyvale, CA 94085 TEL: +1 (408) 962-3600, FAX: +1 (408) 962-3661, support@chelsio.com, www.chelsio.com

Copyright ©2025. Chelsio Communications. All Rights Reserved.

CONTENTS

<u>1.</u>	INTRODUCTION		
<u>2.</u>	HARDWARE		11
	2.1	RECOMMENDED HARDWARE	11
	2.2	RESERVED MEMORY MATRIX	13
	2.3	CONFIGURING THE HARDWARE	14
	2.4	INSTALLING THE PRODUCT	15
<u>3.</u>	<u>FIRS</u>	ST BOOT	16
<u>4.</u>	<u>WE</u>	B-BASED MANAGEMENT	17
	4.1	Accessing the Management Interface	17
	4.2	LAYOUT AND NAVIGATION	22
<u>5.</u>	HAF	RDWARE & SOFTWARE FEATURES MATRIX	24
	5.1	Installation and Boot	24
	5.2	NETWORK CONTROLLER SUPPORT	25
	5.3	STORAGE CONTROLLER SUPPORT	26

	5.4	STORAGE MANAGEMENT/VIRTUALIZATION SUPPORT	27
	5.5	System support	28
	5.6	CLIENTS SUPPORTED	29
	5.7	CLIENT PROTOCOLS SUPPORT	30
<u>6.</u>	<u>NET</u>	TWORKING	31
	6.1	Networking overview	31
	6.2	CONFIGURE AND CONTROL NETWORK DEVICES AND SETTINGS	35
	6.3	NETWORK TEAM CONFIGURATION	48
	6.4	Active Connections monitoring	57
	6.5	NETWORK TROUBLESHOOTING TOOLS	61
<u>7.</u>	<u>Sto</u>	DRAGE	71
	7.1	STORAGE OVERVIEW	71
	7.2	STORAGE MANAGEMENT	74
	7.3	STORAGE CONFIGURATION	90
	7.4	OS DISK DEVICES CONFIGURATION	91
	7.5	SOFTWARE RAID ARRAY CONFIGURATION	109

	7.6	ISCSI INITIATOR	. 118
	7.6.1	REMOTE ISCSI TARGETS	. 122
	7.7 N	IVME INITIATOR	. 129
	7.7.1	REMOTE NVME TARGETS	. 131
	7.8	MICROSEMI (MSCC) ADAPTEC SMARTRAID 3154-8I RAID CONTROLLER	. 136
	7.8.1	PHYSICAL DEVICES AND RAID ARRAYS	. 139
	7.9	VOLUME MANAGEMENT	. 142
	7.10	BACKUP AND RESTORE	. 177
<u>8.</u>	ISCSI SAN		<u>209</u>
	8.1	ISCSI OVERVIEW	.209
	8.2	ISCSI DEPLOYMENT	.210
	8.3	CONFIGURING THE CHELSIO ISCSI TARGET	.211
	8.4	ISCSI SUMMARY	.218
	8.5	CREATING A NEW TARGET	.221
	8.6	ISCSI TARGET SUMMARY INTERFACE	.225
	8.7	ISCSI TARGET LUN CONFIGURATION	.233
	8.8	ISCSI TARGET NETWORK PORTALS	.240

	8.9	ISCSI TARGET CLIENTS / INITIATORS	.244
<u>9.</u>	<u>NV</u>	ME-OF	253
	9.1	NVME-oF overview	.253
	9.2 0	REATING A NEW NVME TARGET	. 255
	9.3	NVME TARGET LUN CONFIGURATION	.258
<u>10</u>	FILE	SHARING	265
	10.1	FILE SHARING OVERVIEW	. 265
	10.2	GLOBAL SETTINGS	.270
	10.3	FILE SHARING CONFIGURATION	. 276
	10.4	User Management	. 297
<u>11</u> .	<u>SYS</u>	TEM TOOLS	305
	11.1	System tools usage	. 305
<u>12</u>	APP	ENDIX	337
	12.1	TROUBLESHOOTING	. 337
	12.2	CHELSIO END USER LICENSE AGREEMENT	.346
	12.3	GNU GENERAL PUBLIC LICENSE	.351

1. Introduction

Chelsio's Unified Storage Server is a powerful easy-to-use turnkey solution for creating highperformance 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 the most PCIe Gen4 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.
- NVMe target/initiator services.
- ▶ NFS, CIFS, FTP, HTTP file sharing services.
- Dynamic storage allocation and management.

Snapshots of storage data for zero-downtime backups, and snapshot scheduling.

Extensive hardware support for various network and storage devices.

- Optimized for the PCIe Gen4 System 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.

This is achieved by a combination of Chelsio's high-performance, scalable storage stack, the Linux operating system, and management software. The product fully supports Chelsio T5 and T6-based Unified Wire adapters for full protocol offload for 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 are the recommended hardware platform for typical usage scenarios:

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

CPU: 64-bit, Octa core Intel Xeon CPU	
Memory:	16GB DDR4
PCIe:	3 x PCle Gen3 8-lane
Network:	2 x 10Gbps / 2 x 25Gbps Ethernet (Chelsio T5/T6 recommended)
Storage:	Adaptec SmartRAID 3154-8i RAID controller, with 15000 RPM SAS hard disks

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

CPU:	64-bit, 2 x Octa core Intel Xeon CPU
Memory:	32GB DDR4
PCIe:	2 x PCle Gen3 8-lane, 1 x PCle Gen3 16-lane
Network:	2 x 40Gbps / 2 x 100Gbps Ethernet (Chelsio T5/T6 recommended)
Storage:	2 x NVMe PCIe SSDs

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
Thin Provisioning (TP)	700 MB of additional memory reserved per pool

- Usage examples:
 - If a pool size of 32 TB with 1 GB RAM size is created, then the system requires a minimum of 10GB + 1GB RAM + 1 GB = 12 GB memory.
 - On the above pool, if a 32 TB volume is created and replicated, then the system requires a minimum of 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 to the <u>Storage</u> 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

The Unified Storage Server is provided on IPMI, USB, or CD/DVD ROM ISO media insert the media to your machine and set the drive as *first boot device* in your system's BIOS. The server will then boot accordingly and present an option to install the product onto a hard drive or a USB flash drive. Refer to the *Quick Start guide* for the first-time setup instructions.

3. First boot

Refer to the *Quick Start guide* for configuring your appliance to boot to Unified Storage correctly and configuring it for the first time. When 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. The 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, and 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 the 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

Chelsio	Welcome to Unified Storage Server
	Login to Unified Storage Server on localhost.localdomain
Adn	ninistrative password:
Note list,	e: Please ensure that this website is added to your trusted sites 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 three 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.

Chelsio	Configuration 👎 Help 🔲 Event Logs	Cogout Support
Vertesx1.blr.asicdesigner.	Unified Storage on Vertesx1.blr.asicdesigners.com	
Network Kools	Status: - Scanning System configuration done. Loading data done.	
System time: 10:41:35 System date: 2025 03:19 English v v	 Various management modules are available from the menu on the left. Unified Storage license status: Not licensed Icensing wizard System: Network: 	

Figure 4.2 – Management Interface Layout

5. Hardware & Software features matrix

5.1 Installation and Boot

Features	Version
	V4.0
Bootable CDROM	Y
USB drive	Y
Minimum space required on HDD	10GB
USB flash boot device	Y
SATA flash boot device	Y
IDE disk boot device	Y
NVMe boot device	Y

5.2 Network controller support

Features	Version
	V4.0
T5 adapters support	Y
T5 RDMA support	Y
T6 adapters support	Y
T6 RDMA support	Y
Support for iSCSI/NFS/CIFS/FTP on Chelsio NIC	Y
Bonding support	Y
T5 Offload bonding support	Y
T6 Offload bonding support	Y
Offload bonding modes	Active-Standby, 802.3ad

5.3 Storage controller support

Features	Version	
	V4.0	
Hardware RAID		
Microsemi (MSCC) Adaptec SmartRAID 3154-8i RAID Controller 12Gb/s	Y	
iSCSI Initiator		
Chelsio T5 based adapters	Y	
	Y	
Chelsio T6 based adapters		

5.4 Storage management/virtualization support

Features	Version
	V4.0
Chelsio Thin Provisioning Volume Management (recommended)	Y
Manual snapshots	Y
Instant snapshots (requires Chelsio TP)	Y
Instant restore from snapshot (requires Chelsio TP)	Y
Instant volume clone (requires Chelsio TP)	Y
Snapshot scheduling (requires Chelsio TP)	Y
Max. snapshots/clones per volume Chelsio TP	254
Max. storage pool size (requires Chelsio TP)	Unlimited (2 ⁶⁴ bytes)
Max. volume size	Unlimited (2 ⁶⁴ bytes)
Max. filesystem size - Chelsio TP	Unlimited (2 ⁶⁴ bytes)

Max. open files per client	64k
NAS backing filesystem	XFS

5.5 System support

Motherboard / Processor	PCIe Gen5 system with at least 12 cores
Memory	128 GB
Ethernet NIC	Chelsio T5/T6 adapter
Storage	Any of SAS, SATA, SSD, NVMe
OS Drive	SAS/SATA/SSD/NVMe disk of at least 250GB

5.6 Clients supported

	Version
	V4.0
iSCSI - Microsoft initiator	Y
iSCSI - RHEL 8.x open iSCSI initiator	Y
iSCSI - RHEL 9.x open iSCSI initiator	Y
iSCSI - Chelsio T5 initiator	Y
iSCSI - Chelsio T6 initiator	Y
NFS - RHEL 8.x	Y
NFS - RHEL 9.x	Y
CIFS - Win 2022	Y
CIFS - Win 2025	Y

5.7 Client Protocols support

	Version
	V4.0
	Y
iSCSI Initiator	(Chelsio T5/T6 only)
	Y
NVMe Initiator	(Chelsio T5/T6 only)
	Y
NFS v3 over TCP	(Chelsio T5/T6 only)
	Y
NFS v3 over RDMA	(Chelsio T5/T6 only)
	Y
CIFS over TCP	(Chelsio T5/T6 only)
	Y
FTP	(Chelsio T5/T6 only)



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 the 192.168.1.0/24 network segment and CIFS clients on the 10.1.2.0/8 network segment, the appliance should be directly connected to both network segments through Ethernet switches, 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 Figure 6.1 (a) and Figure 6.1 (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 it is specifically required. 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 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 to the Firewall's documentation to allow iSCSI, NFS, CIFS, FTP, and SSH traffic through the Appliance.



Figure 6.1 (a) – *Redundant / Multipathed, 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 <u>Chelsio</u> <u>Support</u> 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 by default at system boot. 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 on 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 the bandwidth status are listed here. A count of the active TCP connections to the Appliance and DNS servers summary is also displayed. IP addresses can be configured for upto three DNS Servers. Hostname can also be configured in this page.
- Summary:	
Ethernet Devices details:	
Ethernet status: IP address: Bandwidth status: Chelsio Storage Accelerator:	8 total, 3 active, 5 with no link. 169.254.0.202, 10.193.185.202, 102.1.1.202 1 x 100 Mb/s 2 x 10000 Mb/s Total 4 ports present, 4 ports offload enabled
Connections details:	
Active TCP connections: Offloaded TCP connections:	2 0
DNS & Hostname:	
DNS status:	2 DNS servers configured
Edit DNS servers:	Server 1: 10.193.184.187 Server 2: 10.193.180.20 Server 3:
Hostname:	Configured Edit

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-pane, with an icon indicative of its current state. A text status, any configured IP address, and the physical MAC address are displayed in the center. Control options for the interface and more information, are available using the buttons on the right. Configuration of the interface can be viewed and changed below, by expanding the configuration link.



	Intel Corpor	ation 1350 Gigabit Network Connec	tion (rev 01)
eth1	Status: IP address	Physical Link down, no IP address	Properties & Statistics
🕂 Confi	MAC Addre	ss: 00:25:90:99:b2:4f	€ Restart

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





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

	Intel Corporation 82576 Gi	gabit Network Connection (rev 01)	
eth0	Status: IP address:	Physical Link up, IP configured 10,193,185,76/22	Properties & Statistics
100 Mb/s	MAC Address: IPMI BMC LAN interface:	00:25:90:10:CD:2A	Stop
	IPMI BMC LAN IP address	10.193.185.77/22	- 😥 Restart
+ Config	uration:		

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

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 two 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 the device specific statistics are grouped accordingly.

Network Devic	e eth1 details
<- Back to Devices list Proper	ties Statistics
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 follows:

 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 that 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.

 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 are 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 recommended 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:	Statically assign 🗸
IP address:	192.168.20.5
Subnet mask:	255.255.0
Gateway:	192.168.20.1
Broadcast address:	
Device activated:	on System startup 🗸
MTU:	1500 Bytes 🗸
	Apply
Add VLAN child dev	ice: 32
VLAN device config	uration: New / blank configuration 🗸
	New / blank configuration
	Migrate eth3 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 allow 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 to 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)	Create new Team	
eth0 [00:07:43:	05:E9:93]	IP Configuration type:	select a method
eth1 [00:07:43:	05:E9:94]	n conngaration type.	Sciectamento
eth2 [E4:1F:13:	:2C:97:98]	IP address:	
✓ eth3 [E4:1F:13:	:2C:97:9A]	Subnet mask:	
		Default Gateway:	
		Broadcast address:	
		Device activated:	on System startup 👻
I		Primary Team member:	eth0 🗸
		Team mode:	Balance-RR -
			Apply

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: Ensure 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 is an example of how to create a Network Team:

1. 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.

- 2. 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.
- 3. If you have chosen **Statically assign**, then provide the IP address, Subnet mask, Gateway, and Broadcast for the team.
- 4. Specify if you want to manually activate the device or automatically on system startup in the **Device activated** field.
- 5. Specify the primary member of the team.

6. 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.



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

Devices: (multi-select)	Create new Team	
eth0 [00:07:43:05:E9:93]	IP Configuration type:	Import from member device 👻
eth1 [00:07:43:05:E9:94]		
eth2 [E4:1F:13:2C:97:98]	Import from member:	eth2
eth3 [E4:1F:13:2C:97:9A]	Imported configuration type:	Statically assign 👻
	IP address:	10.193.184.157
	Subnet mask:	255.255.252.0
	Default Gateway:	10.193.184.1
<u> </u>	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:

Some team modes require a primary member, which can be assigned in this section.

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 an evenly distributed traffic across all the links. It also provides fault-tolerance.

Active Standby:

Only one of the member device 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.

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 at least 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:	yes 💌
(Note: Offload information cannot be shown with na	me resolution enabled)
Show listening services:	yes 💌
Layer 4 protocols:	TCP + UDP 💌
Layer 7 protocols:	All 💌
Decode application layer protocol:	yes 💌
Refresh page every:	30 seconds 💌
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.

Ping a hostname or IP address:	www.ietf.org (IP address or Hostname)
	Ping
Ping result for www.ietf.org:	success. 5 packets transmitted, 5 received, 0% packet loss, time 4024ms 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.



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.

- Advanced settings:	
DNS_Settings	DNS Servers
 DNS search paths: bit asicdesigners.com, asicdesigners.com DNS timeout: DNS retries: DNS load balancing: ARP_Settings ARP filter: disabled 	Server 1: Server 2: Server 3:

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

1. DNS Settings

- DNS servers: You can specify up to three 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 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): The maximum number of servers waiting to open a listening socket at one time.



Warning: Do not change these settings unless you have a clear understanding of their impact. Incorrect values can affect multiple services, network connectivity, and overall system stability.

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.

outing table:						
Network/Host	Subnet Mask:	Gateway	Network Device:	Source IP address	Actions	
default		10.100.104.5	enp3s0		Edit 🔀 Delete	
10.100.104.0	200.200.202.0		enp3s0	100	Edit 🔀 Delete	
			Select V	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 an ARP.

table:				
IP Address:	MAC Address	Network Device:	State	Actions
10.100.104.5	10 25 54 2a ct 17	enp3s0	REACHABLE	🗾 Edit 🐼 Delete
10.100.100.100	to of Teceberth	enp3s0	STALE	🗾 Edit 🐼 Delete
10.100.000.000	101-101-001-241-440-c-1	enp3s0	STALE	🗾 Edit 🐼 Delete
10.100.000.00	70.40330.0010.001.0	enp3s0	STALE	🗾 Edit 🐼 Delete

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.

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. The 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, SATA / SAS, and NMVe.

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, it 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.

d. NVMe: This protocol enables high-performance, low-latency access to NVMe storage devices over a network. It allows NVMe commands to transmit over various network, such as Ethernet and InfiniBand making it possible to share NVMe storage resources across multiple servers in a data center or cloud environment.

7.1.2 Add-on storage controllers

An add-on storage controller is generally a SAS / SATA RAID 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 <u>section 7.2.1</u> under **Storage configuration**.

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 are 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.





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 is being written to both drives simultaneously. This ensures that data will be available if one of the two drives fail, but the write performance is equal to the 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 two 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) x size of the 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 three 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 the 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 four 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 four 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 written actual space. 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 do not 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 at the beginning of the pool. The size of the storage pool that can be created is 2⁶⁴ sectors.
- Zoning: The 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. 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. For more information, refer to the **iSCSI SAN** section.



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 it 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 to 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 to 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 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. That is 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 a 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.

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 the physical disks attached to the system. These may include hardware RAID arrays configured on RAID controllers and LUNs discovered through iSCSI/NVMe 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, and so on.



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.



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 displayed below. Sections in blue are currently configured, whereas sections in green are free / unconfigured.

Storage Server Boot Drive 0	Capacity: 465.76 GB 🛕 SMART
Device properties:	
Serial #: 9VMVPDHM	
Disk cache:	
Read cache: Enabled ~	
Parameters stored on device: No	
r didineters stored on device. No	
Device lavout:	
1 2 3 4 5 6	7

Figure 7.4(c) - OS boot disk

sda Unified Storage Server Boot Drive	Capacity: 111.79 GB 🛕 SMART
Device properties:	
Serial #: 162252404247	
Disk cache:	
Read cache: Enabled V	
Write cache: Enabled V	
Parameters stored on device: No	
Device layout:	
123456780	

Figure 7.4(d) - Partitioned disk

anvme0n1 SAMSUNG MZVLW256HEHP-000L7 Capacity: 238.47 GB				
Device properties:				
Disk cache: Read cache: Write cache: Parameters stored on device Cache segments count:	Enabled V Enabled V ENO 20			
Device layout:				
Physical volume in Volume management Pool Ptest				

Figure 7.4(e) - Disk used in volume management

🖾 sdc MSCC LOGIC	CAL VOLUME Capaci
Device properties:	
Serial #: 8A29F3005E8	
Read cache:	Disabled 🗸
Write cache:	Disabled V
Parameters stored of	n device: No
Device layout:	
Device layout.	

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:

1. Select the disk on which the partition is to be created and click **Device details**.

- 2. In the **Device layout** section, click on the green section.
- 3. From the Actions drop-down, select *Create partition*.
- 4. Enter the partition size in MB, GB, or TB.
- 5. 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 TP by default.

2.1. Chelsio Volume Management (Thin Provisioned, recommended) Using this option, the 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	Actions:	Manage with Volume Management \checkmark
465.76GB	Volume Management usage:	Assign to Pool: Create new pool 🗸
	New Pool name:	
	Volume Management Type:	Chelsio TP (recommended) V
	Pool Size:	8 TB • (>= 32 GB)
	Chunk size:	I28K (recommended) \bigcirc 64K \bigcirc 32K
	Disk for Cache:	Internal V
	RAM Cache :	Disable 🗸
	Allocation Size:	8 🗸 chunks
	Allocation Zones per disk:	16 🗸
	Max volumes in a pool:	32767 (less than or equal to 32767)
	Note: "Maximum volumes ir after creating the pool.	a pool" cannot be changed
		Apply

Figure 7.4(h) - Creating a new TP Pool

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, 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 of the appliance.
- Chunk Size: The chunk size should be a power of two. The allowed chunk sizes are 32, 64, and 128K (default value).
- Allocation size: Number of chunks allocated for every new write.
- Allocation Zones per disk: The 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 the **Storage volume management** page.

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

Devi	ce layou	t:		
1 2		3		
	2 بالد م	Tuno:		Size: 25.00.0D
	SOKJ	Usage:	xfs formatted filesystem, currently detached	Details: unused
		Actions:	Delete Partition	
		Folder Path:	Configure folder to attach Erase Data, Manage with Volume Management Delete Partition	
			Apply	

Figure 7.4(j) - Partition details and actions

- 1. Partition actions
- Erase data and manage with Volume management: This option allows the user to assign the partition to be used in Volume management.
- Delete partition: This option deletes the partition permanently.
- Configure the folder to attach: Using this option, you 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, you 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: 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 reliability indicators in the hope of anticipating failures.

Sections of the interface

1. Summary

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

- Summary	
SMART Details:	
Transport protocol:	SAS
Device:	FUJITSU MAX3073RC Version: D206
SMART support is:	Enabled
Device type:	disk
Serial number:	DQA3P6C05F68
	Disable SMART Support

- Summary SMART Details: see the following Seagate web pages: SATA 2.6, 3.0 Gb/s SATA Version is: ATA Version is: ATA8-ACS T13/1699-D revision 4 Model Family: Seagate Barracuda 7200.12 Serial Number: 9VMVPDHM Firmware Version: CC44 Device Model: ST3500418AS Rotation Rate: 7200 rpm Sector Size: 512 bytes logical/physical SMART support is: Enabled A firmware update for this drive may be available, ==> WARNING: LU WWN Device Id: 5 000c50 02d593332 X Disable SMART Support

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(I) - 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 a hardware manufacturer and can include Power-On Hours, Temperature, Seek Error Rate, and many others. A short test typically completes in under ten minutes, whereas a long test can take several 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

- 1. 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.
- 2. Now, click on the **SMART** button. This will navigate to the **SMART test** page.
- 3. Click on Enable SMART support in the Summary section, if it is disabled.
- 4. In the **Smart Test** section, select the type of test (Short, Long) and view health status or logs using the **Actions** drop-down.
- 5. 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.

Physical Disks (multi-select)		Properties / Actions:
✓ nvme0n1 - 238.47 GB		
✓ sdb - 465.76 GB	ID / Slot:	nvme0n1
🖋 sdc - 465.76 GB	Vendor:	
	Model:	SAMSUNG MZVLW256HEHP-000L7
elect multiple disks in the list above to create a RAID	Actions.	

Figure 7.5(a) - Free physical disks section

1.1. Free / Unassigned Physical device 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 no arrays are 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, and 60. For more information, refer to the RAID levels explanation, at the end of this section.
- **Stripe size**: The RAID algorithm splits incoming data into 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

- 1. 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.
- 2. From the **Actions** drop-down, select *Create a new RAID Array*.
- 3. Select the RAID level (available levels differ depending on the number of disks selected).
- 4. Select the Stripe size (if available).
- 5. 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 hotspare, any available actions for the disk are shown on the right.

— Physical Disks ———		Properties / Actions:
🛷 sda [465.76 GB]		
✓ sdc [465.76 GB]	Array ID:	md0
	Stripe size:	128k
	Status:	active
	Actions:	Part of Volume Management 🗸

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 disk section.

Physical Disks		Properties / Actions:
🖋 sda [931.51 GB]		
✓ sdb [2048.00 GB]	ID / name:	sdb
✓ sdc [931.51 GB]	Model:	ST33000651AS
	Status:	Hotspare Good
	Actions:	Remove Hotspare / Standby -

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 TP by default.

1. 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.

Free	Actions:	Manage with Volume Management 🗸	
465.76GB	Volume Management usage:	Assign to Pool: Create new pool	
	New Pool name:		
	Volume Management Type:	Chelsio TP (recommended) 🗸	
	Pool Size:	8 TB ▼ (>= 32 GB)	
	Chunk size:	○ 128K (recommended) 64K ○ 32K 	
	Disk for Cache:	Internal V	
	RAM Cache :	Disable ~	
	Allocation Size:	16 v chunks	
	Allocation Zones per disk:	16 🗸	
	Max volumes in a pool:	32767 (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 of the appliance.
 - ii. **Chunk Size**: The chunk size has to be in power of two. The allowed chunk sizes are 32, 64, and 128K (default value).
 - iii. Allocation size: Number of chunks allocated for every new write.
 - iv. 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.
 - v. **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 on the **storage volume management** page.

- 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 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.

- Global Settings	
Initiator Properties	
✓ 1. Enable CHAP authentication for a discovery	
✓ 2. Discovery session CHAP password for initiator	Enable CHAP
✓ 3. Discovery session CHAP password for target	authentication for a
✓ 4. Discovery session CHAP username for initiator	uiscovery
✓ 5. Discovery session CHAP username for target	Set this 🗆
✓ 6. Maximum number of data bytes initiator can receive during discovery session	Cotting
✓ 7.	Setting. CHAP V
A ISCSI PDI I Header Direct / Checksum	
Apply	

Figure 7.6 (b) - Initiator Global Settings

3. Chelsio Adapter configuration

This section lists the Chelsio adapter details such as IP address, status, driver, and Mac address.

iSCSI network ports —	Interface: eth2		
🧨 1. eth2	Status: Physical Link	an	
✓ 2. eth3	Iface Configuration:		
	Driver: HW Address:	cxgb4i 00:07:43:4b:98:90	
	IP Address Default: IP Address:	default 102.50.50.192	

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	n.2013-04.dzongri:1 102.20.20.165
Enabled on startup: Yes	
Status: Connected Session ID: 1	Actions: Rescan this target
Connection: 102.20.202.202 -> 102.20.20.165:3260 Iface Name: cxgb4i.00:07:43:10:69:20	

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, provide the target IP address, port details, and the interface you want to bind the target. In the second mode, mention the IP address of the 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.

- Add Target				
Oirectly connect to target				
IP Address:	102.77.77.147			
TCP Port:	3260			
Iface:	eth2 [IP: 102.1.1.62] 👻			
Discover Target				

Figure 7.6.1 (e) - Directly connecting to a target

Query iSNS Server for targets			
IP Address of iSNS Server:	102.88.88.145		
Iface:	default 👻		
Discover Target			

Figure 7.6.1 (g) - Querying iSNS for targets

Example: How to add a remote iSCSI Target manually

Ensure that the iSCSI target to be added must be running before attempting to add.

- 1. From the home page, navigate to **Storage** > **iSCSI Initiator** > **Remote iSCSI Targets**.
- 2. If the target IP address and TCP port number are already known, you can add the target directly. Select the **Directly connect to target** radio button.
- 3. Specify the IP address of the target.
- 4. Specify the TCP port. The port specified here must be the same as provided on the target. If the default option is selected, then the TCP port is 3260.
- 5. Choose a Chelsio interface to bind the target.

- 6. Click **Discover Target.**
- 7. If the target was discovered successfully, it appears in the **Target Details** section. Select the newly discovered target node and select *Login to this target*.
- 8. Click Apply.
- 9. 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 manually. Once the target appears in the **Targets List** in the **Target Details** section, follow these steps:

- 1. From the home page, navigate to **Storage** > **iSCSI Initiator** > **Remote iSCSI Targets**.
- 2. Select and highlight the target node in the Targets List.
- 3. Expand the Target Properties section.
- 4. Scroll down to the *Authentication Type* parameter. Select (one-way) *CHAP* or *Mutual CHAP* from the drop-down.
- 5. 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.
- 6. 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** page under **iSCSI Initiator** module section.

7.7 NVMe Initiator

Sections of the interface

1. Summary

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

Summary	
Initiator NQN N	Name: 2014-08.org.nvmexpress:uuid:00000000-0000-0000-0000-00259099b24e
Service Details:	
Service Status:	Installed, Auto-start Enabled

Figure 7.6 (a) - Initiator summary and control actions

- 1.2. Initiator control actions
- Enable/Disable: You can choose to start/stop NVMe 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 stops and starts the NVMe Initiator service.

7.7.1 Remote NVMe Targets

Sections of the interface

1. Summary

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

– Summary		
Total targets discovered:	1	
Targets enabled on startup:	0	
Total targets Connected:	0	
Total Luns from NVMe-oF targets:	0	

Figure 7.6.1(a) - Remote NVMe Target summary

2. Target Details

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

– Target Details			
Targets List			
1. Portal: 102.50.50.212,4420	Target: nqn.2025-03.com.asicdesigners.blr.t7qa1:1 102.50.50.212		
Enabled on startup:			
Status: Connected	Actions: Logout from this target Actions:		
Session ID:	Apply		
Connection: -> 102.50.50.212,4420			
nace name.			

Figure 7.6.1 (b) - Remote NVMe Target details

1.2. Target Details actions

- **Rescan this target**: Rescan performs 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. Add a Target

Select the target type and provide target IP address, port details, and the interface you want to bind the target. In the iface drop-down, only Chelsio interfaces are listed.

– Add Target	
Directly connect t	a taraat
Directly connect t	Jargel
	TCP 🗸
IP Address:	10.193.187.44
TCP Port:	4420
lface:	~
	Discover Target

Figure 7.6.1 (e) - Directly connecting to a target

Example: How to add a remote NVMe Target manually

Ensure that the NVMe-oF target to be added must be running before attempting to add .

- 1. From the home page, navigate to **Storage** > **NVMe-oF Initiator** > **Remote NVMe Targets**.
- 2. Select the target type from the drop-down menu. The options are TCP and iWARP.
- 3. If the target IP address and TCP port number are already known, you can add the target directly.
- 4. Specify the IP address of the target.
- 5. Specify the TCP/iWARP port. The port specified here must be the same as provided on the target. If the default option is selected, then it is 3260.
- 6. Select a Chelsio interface to bind the target.

- 7. Click Discover Target.
- 8. 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**.
- 9. Click Apply.
- 10. Select the target node to view its properties.

Example: How to access NVMe Target with CHAP authentication enabled

Follow the instructions mentioned above to add an NVMe Target manually. Once the target appears in the **Targets List** in the **Target Details** section, follow these steps:

- 1. From the home page, navigate to **Storage** > **NVMe-oF Initiator** > **Remote NVMe Targets**.
- 2. Select and highlight the target node in the Targets List.
- 3. Expand the Target Properties section.
- 4. Scroll down to the *Authentication Type* parameter. Select (one-way) *CHAP* or *Mutual CHAP* from the drop-down.
- 5. 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, the initiator name of the peer node must also be provided in the *Peer initiator username* field for both one-way and mutual CHAP.
- 6. Click Apply.

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

7.8 Microsemi (MSCC) Adaptec SmartRAID 3154-8i RAID Controller

Sections of the interface

1. Adapter Summary

This section lists the controller details such as Controller Status, Channel description, Controller Model, Controller Serial Number, Physical Slot, Installed memory, Copyback, Background consistency check, Automatic Failover, Defunct disk drive count, Logical devices/Failed/Degraded, BIOS, Firmware, Boot Flash, Driver Version, and Controller Battery Information. Figure 7.8(a) shows the adapter summary page.

- Adapter Summary		
Controller details:		
Controller Status:	Optimal	
Channel description:	SCSI	
Controller Model:	MSCC SmartRAID 3154-8i	
Controller Serial Number:	8A29F3005E8	
Physical Slot:	3	
Installed memory:		
Copyback:		
Background consistency check:	Idle	
Automatic Failover:		
Defunct disk drive count:	0	
Logical devices/Failed/Degraded:	1/0/0	
BIOS :		
Firmware:	1.32	
Boot Flash:		
Driver Version:	Linux 2.1.24-046	
Controller Battery Information:		

Figure 7.8(a) – *The Adapter Summary page*

2. Settings

These settings are used to configure the controllers. You can rescan for the missing options and restore the default controller configuration. Additionally, you can enable or disable features such as the alarm, copy-back mode, background consistency check mode, automatic failover

mode, and native command queuing. Select the performance mode from the **OLTP/Database** and **Default/Dynamic** options, then click **Apply**. To upload or update firmware, browse for the firmware file and click **Upload**. Figure 7.8(b) shows the settings page.



Figure 7.8(b) – The Settings page

7.8.1 Physical Devices and RAID Arrays

This section displays the physical disks along with their corresponding properties. To modify the disk's properties, select the disk from the left pane. Under **Properties/Action**, select the desired action from the drop-down menu and click **Apply**. Figure 7.8(c), Figure 7.8(d), and Figure 7.8(e) shows the physical devices and RAID arrays.

🚜 Free / unassigned Physical disks			
Physical Disks (multi-select) Device #5(Direct Attached, Slot 5(Connector 1)	Disk ID : Model: Vendor: Drive Type: Status: Actions:	Properties / Actions: 5 ST4000NM025B SEAGATE SAS 12.0 Gb/s Ready Locate / Blink Disk Locate / Blink Disk	Apply
*Select multiple disks in the list above to create a RAID array.** Select a single disk to modify the disk.		Stop Blinking Create Simple Volume Add to existing RAID Array	
		Assign as Dedicated Hotspare	
🚜 Hot Spares		Assign as global HotSpare	
No free disks are available to configure.		Initialize the Disk	
		Clear Disk	

Figure 7.8(c) – Free / unassigned Physical disks

Hot Spares	
No free disks are available to configure.	Properties / Actions:

Figure 7.8(d) – Hot Spares

Physical Disks ———————————————————————————————————		Properties / Actions:
 Device #0(Direct Attached, Slot 0(Connector 0) Device #2(Direct Attached, Slot 2(Connector 0) Device #4(Direct Attached, Slot 4(Connector 1) 	0 1D / Slot: Model: Drive Type: Status: Vendor: Actions:	0 ST4000NM025B SAS 12.0 Gb/s Online SEAGATE Locate / Bilnk Disk ♥

Figure 7.8(e) – Logical devices and RAID Arrays

To modify the properties of RAID arrays, click **Edit** and choose an action from the **Actions** dropdown menu and then click **Apply**. Figure 7.8(f) shows the edit page.

— Physical Disks —————————————————————	7	Properties / Actions:	
Device #0(Direct Attached, Slot 0(Connector 0)()			
Device #2(Direct Attached, Slot 2(Connector 0)()	Array Name:	R5	
Device #4(Direct Attached, Slot 4(Connector 1)()	Status:	Optimal	
	Stripe Size:	1024 KB	
	Write-cache mode:		
	Write-cache setting:		
	Read-cache		
	RAID Level:	5	
	Partitioned:		
	Bootable:	None	
	Power settings:		
	Failed stripes:		
	Actions:	Select an action 🗸	
	New Array name:	Select an action	
		Blink /Identify all drives	pply
		Stop Blinking of all drives in the Array	pply
		Stop Dimining of an unves in the Array	

Figure 7.8(f) – *The Edit page*

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: The size of pools, volumes, and physical devices listed in this section are calculated using the binary convention 1KB=1024bytes, 1MB =1024 KB, and so on.



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. For example, whether the pool is Writeable or not, Chunk Size, Snapshots count, Logical/Allocated Volume count. 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

This is particularly beneficial when USS is upgraded from the previously released version 3.0 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. You can also delete the pool.

🝘 Pool: Memory Type: Legacy Size: 2.00GB Free: 2.00GB 📝 Edit					
Inner Devices:					
Physical devices :			Logical Volume	s :	
\delta sda9					
Details and Actions:					
📷 Pool: Memory					X Cancel
Pool Details:		Usage:			
Writable:	yes	Allocated Volumes Space	e: 0.00MB		
Resizable:	yes	Actual Physical Disk spa	ice: 2.00GB		
Clustered:	no	Actions :		Allocate new Space / Logical Volume	~
Snapshots count:	0	0:		10	
Physical Volume count:	1	Size of volume as % of f	eespace in pool:		
Logical/Allocated Volume co	ount: 0	Size of new Volume:		204 MB 🗸	
		Volume Name (optional)			
					Apply

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

1.3. Pool Actions

1.3.1. 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 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:	Auto Select 🗸
Enable volume encryption:	No 🗸
	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. If 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.

😽 Pool: Ptest				💥 Cancel
Pool Details:		Usage:		
Active: Pool Mode: Status: Maximum limit for pool size:	Yes Async Valid 16777216	Allocated Volumes Space: Actual Physical Disk space:	5.00GB 238.47GB	
Device ID: Chunk Size:	TB 253:0 128K	Space:	33.75MB	Allocate new Space / Logical Volume 🗸
Maximum number of Volumes: Physical Volume count: Logical/Allocated Volume count:	32767 1 1	Size of volume as % of Size of new Volume: Volume Name:	freespace in pool:	10 838349 MB ~
		Device for data: Allocation zone: Enable volume encrypt	ion:	Distributed on all devices Auto Select No
				Apply

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.
 - a. Remove missing physical devices from Pool: Remove any physical device which is not present, in 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: Deletes 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:	Auto Select 🗸
Enable volume encryption:	No 🗸
	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.

Actions :	Allocate new Space / Logical Volume 🗸
Size of volume as % of freespace in pool:	
Size of new Volume:	838349 MB 🗸
Volume Name:	
Device for data:	Distributed on all devices 🗸
Allocation zone:	Auto Select 🗸
Enable volume encryption:	No
	Encrypt using Master Key
	No

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

- Deactivate the Pool: Using this option, the user can disable the selected pool.
- Resize the pool: Using this option, the user can resize the selected pool.

Note: The aggregate size of all the pools configured on the appliance cannot exceed the storage capacity of the appliance.

 Activate/Deactivate: All logical volumes in Pool: Enable/Disable all the 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 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 the 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.

1.4. Physical Devices actions

1.4.1. Logical Volume Management (Legacy)

Click on any one physical device to view the various details and actions.

Details and Actions:		
Physical Volume: nvme0n1 238.4	7GB	💥 Cancel
Volume Details:	Actions:	
Parent device size : 238.47GB	Actions:Select an Action	
Size Allocated to pool: 238.47GB	Select an Action	Apply
Free space: 0.09MB	Extend the physical volume	
	Replace disk	

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 Source Physical
 Volume 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.
- Extend the physical volume: Extend the volume of selected physical device.
 - Replace disk: This option is used to replace the disks by selecting the destination disk and log device.
- 1.4.2. 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 the storage capacity greater than or equal to the current 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 a Log device to store metadata during the live migration process. The selected disk will be free once the process is completed.

- 1. On the home page, navigate to **Storage > Volume management**.
- 2. In the Volume Management module, expand the Storage Pools section.
- 3. In the **Pool list,** select the pool, the physical device of which is to be replaced.
- 4. In the **Physical devices** section, click on the disk to be replaced.
- 5. In the Actions drop-down, select *Replace disk*.
- 6. In the **Physical disks** drop-down, select the disk which will replace the selected disk.
- 7. In the **Log Device** drop-down, select the disk to be used as a log device.
- 8. Click Apply.
- 9. 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.

10. You can cancel the process anytime by selecting either of the disks and selecting *Abort Disk replacement.*

1.5. Logical Volume actions

1.5.1. 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 Figure 7.9(j).

Logical Volume: Memory Pool: Ptest Size: 5.00GB				
Volume Details:	Usage:			
Active: yes	Shared as iSCSI LUN in:	Target-1		
Encrypted: No	Actions:	Select an Action		
Number of Siblings: 0		Select an Action		
Time of Creation: Tue Mar 11 16:55:02 2028	i	Extend/Reduce the Logical Volume		
		Rename Volume		
		Modify data Allocation Ta	arget and Apply	
		Create a clone		
		Schedule Snapshots		
		Create Snapshot		

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 the 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.

olume Details:	Usage:	
Active: yes		Target-1
Read-Write Access: yes	Shared as ISCSI LUN in:	•
Encrypted: No	Actions:	Extend/Reduce the Logical Volume V
Number of Siblings: 0		
Time of Creation: Tue Mar 11 16:55:	02 2025	
	New Volume size:	MB 🗸
	Resize and Refresh running is	SCSI Target

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 the 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 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 the 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 back up the data at the point of creating the snapshot.

Logical Volume: Memory Pool: Ptest Size: 5.00GB				
Volume Details:	Usage:			
Active: yes	Unused			
Read-Write Access: yes	Actions:	Create Snapshot		
Encrypted: No				
Number of Siblings: 0	Snapshot Volume Name:			
Time of Creation: Tue Mar 11 16:55:02 2025				
	This creates a point-in-time images to this point-in-time copy. Warning: No IO should be running snapshot of the volume.	ge of the Volume Ptest_Memory and allows access ng to the Volume, to get a consistent, good		
		Apply		

Figure 7.9(m) - Create Snapshot

1.5.2. 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. Similarly, Logical Volume size can also be decreased. **Note**: Logical volume size cannot be decreased if the logical volume is part of a target and the target is running currently.

/olume Details:	Usage:	
Active: yes	Unused	
Read-Write Access: yes	Actions:	Extend/Reduce the Logical Volume 🗸
Encrypted: No		
Number of Siblings: 0		
Time of Creation: Tue Mar 11 16:55:02 2025	New Volume size:	MB ~

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.

• 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, the user can schedule snapshot creation for a volume on an hourly, daily, or weekly basis.

Logical Volume: Memory Pool: Ptest Size: 5.00GB				
Volume Details:	Usage:			
Active: yes	Unused			
Read-Write Access: yes	Actions:	Schedule Snapshots		
Encrypted: No				
Number of Siblings: 0	Hourly Snapshots:	Disabled		
Time of Creation: Tue Mar 11 16:55:02 2025	Daily Snapshots:	Disabled		
	Weekly Snapshots:	Disabled		
		Apply		

Figure 7.9(p) - Scheduling a daily snapshot

Example: Scheduling Snapshots

- 1. Select the unencrypted logical volume for which you want to schedule snapshots. This will navigate to the logical volume properties/actions page.
- 2. From the Actions drop-down, select Schedule Snapshots.
- 3. Select if you want the snapshot to be created on an hourly, daily, or weekly basis by selecting the checkbox for each.
- 4. 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.
- 5. For weekly snapshots, you can select the days on which snapshots will be created.
- 6. 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*.
- 7. Click Apply. You can view the details of the scheduled snapshot in the Settings section.

- 8. To change any settings, click on the volume and select **Schedule Snapshots** from the **Actions** dropdown.
 - Create Snapshot: With this option, the 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.

Logical Volume: Memory Pool: Ptest Size: 5.00GB				
Volume Details:	Usage:			
Active: yes	Unused			
Read-Write Access: yes	Actions:	Create Snapshot		
Encrypted: No				
Number of Siblings: 0	Snapshot Volume Name:	Snapshot1		
Time of Creation: Tue Mar 11 16:55:02 2025				
	This creates a point-in-time image to this point-in-time copy. Warning: No IO should be running snapshot of the volume.	of the Volume Ptest_Memory and allows access to the Volume, to get a consistent, good		
		Apply		

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.

1.6. Snapshot Actions

1.6.1. Logical Volume Management (Legacy)

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

Snapshot list	Snapshot Details:	Usage:
✓ 1. Name: VirtualTapeDrive1-Snapshot1	Active: Yes	Space used in snapshot:
	Encrypted: No Time of Creation: Fri Dec 27 19:13:57 2024	Used: 0.00%, Free: 100.00%
		Not attached to a folder, not assigned as iSCSI LUN.
		Actions:Select an Action 🗸
		Select an Action
		Rename the Snapshot Volume
		Delete the Snapshot
		Roll back to this snapshot
		Deactivate the Snapshot

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 the 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.

1.6.2. 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 deletes 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 that 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.
- $\circ~$ 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:

- $\circ\,$ 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

The following are the steps to roll back a logical volume to a particular snapshot:

- Before a logical volume is rolled back, please ensure that the volume and the snapshot are free/ unused. For example, the volume and snapshot should not 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, ensure that they also are free/unused.
- 2. In the **Snapshot-list**, select the snapshot to which you want the logical volume rolled back.
- 3. If the snapshot is not active, select **Activate the Snapshot** from the **Actions** drop-down and click **Apply**.
- 4. 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. The user can also unlock any newly added encrypted volumes.

Encryp	tion Set	tings:		Select an Action	-						
Note: C	orrespo	onding	recovery k	ceys a <mark>Select an Action</mark> Reset Master Pass Phrase Reset volume pass phrase		ons	ettings	5.			
Snapst	ot Sche	eduling	Settings	Upload old master key Upload master key for new vols							
	Hourly	Snapst	nots time:		10	•	minut	tes	pastt	the h	our
	Daily S	napsho	ots time:		00	•	: 05	•	АМ	•	
	Weekly	Snaps	hots time	c .	00	•	: 15	•	AM	•	
Schedu	led sna	pshots	:			∕∕Ap	oply				
	Туре	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 a minimum of six characters.

lew Pass phrase:	•••••	
Confirm Pass hrase:	•••••	
Please paste the inencrypted recovery key here:	BEGIN 853 MIIBOWIBAAJBAG z4jGqmMHAct9r: NwNVQDOJLIFd2: YQIAAPFWLGHHD JJIEAvHDQYw1aJ 1H0YwDxvM8fWvg tVZr2umDhrTkn2	PRIVATE KEY DfDLS8ujGSj1gLtrtn50AhRBwMaMWsvRogNjFQH0oU8bXNV63Hz f8857buGGZuJTRWUaGG3UZMCAwEAAQJATm+7xxAV66mXHH6IgrUN YXCTEQpd3WLHybasJTqHIz5/yt5oqnu7q/XIToAn/Kx1s304PQ9 gky+q71HVc1CGHqerIEtcDABeedDxVzKvV3AiEA72V0qMUymN8z KAVU20e9EcVBk7Gy6UCIQCFVSrIrk23pv+KvCRuz70Uygs8XqBR QIge8bM05rCH3KAX08znzgbS2/R2fM+w4TPFWRPRZSjipUCIQDh NNGX85FkFIIMtte7R8JFjsKUMkFCQ== DVUATE_EXX

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:

- \circ $\;$ The new pass phrase has to be a minimum of six 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:	u1	•
Encrypted Volume:	vol3	•
Please paste the unencrypted recovery key here:	BEGIN B MIIBOWIBAAD /o7sho4mNG03 ItbNRAhZ3Rm lgECIQDODH2- vtIScQPJH/jC /ve4it97Y4nx REFIAC1UZJB9 END B3R	<pre>SA PRIVATE KEY BANSIRhbF4RYkgQhdNHVEH3ASJZDCK8KQ381SOKVDHjtoxMyNQIIu tieQniLAY72.+d2kuFBNgOOCAwEAAQJBAIKNEE2rpBxm2xaQuJJ1 kkqA2qptINIVy8LCh02tHrOCXY8dSxFKUIw8J/kMaVOCPCF3xVEU c2Rorzddo/1WXQXuJpu6iJX0Y2H/hXJsqT3LBQIhAOXCM08M6+Wo odWvC2v6s8c6BaJPc9LJALASuxo1Fm6+cOhGK21DeaR14Q220WDg mQIhAMB0/ypvJ3bz+Av+w9VI9zqiVXz+hkucYWFIqQWqC1pJALBq DDQNNEIMQ31sQSmurkrVGv6yT9Z5dw== } PRIVATE KEY</pre>
	Note: Make sure	you give the valid recovery key.
e Go Back		Change Pass Phrase

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. For example, after the reinstallation of the operating system. 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 u	nlocks all the old encrypted volumes, and re-encrypt with current master key.
Current Master pass phrase:	•••••
Please paste the unencrypted recovery key here:	BEGIN 85% PRIVATE KEY MIIBOgIBAAJBALmmBI8hAllpCPr9mcnhfeC6gnw5L6mxykNwjWsmHTdkfXV02DTC wSsfuYzXHNP4VJHWJHGOWSMY1sK/EIH1sCAwEAQJACAcObHxLfWzK9:jjAB9 yPuKgwv58+hFEAKKVTAbV8jpVvZR3czFHdTeXINvgVHob8W+6o+tZtt2ZtvzR/c AQIhAOTJgqAIpMuf1f4btqicIH3j7ZHnC9nzRmISH6GEZNtbAiEAz7rzZm4uL83T dyjCLzAZTxYpDg461Ru/SNtXPegSOQECICGK1WP84GwZOVuXpIfDm9/Y8pNZVBt3 lvG5xn1f4dnRAiBeOK/6F/CMHObRf5LEng6Uvrp0cJP16K5Lgem58PWuAQIhAMWW iCeIYN7JcUCjIdVhFd1ypBMwt/1H91KGmIDbgR3f END 85% 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 and 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 Backup and Restore

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 enterpriselevel 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.

7.10.1 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.

Creating a Virtual Storage device

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

- 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.
- 4. Specify the name of the virtual tape drive.
- 5. Click Next.

Tape drives:	Details
+ Add a Tape drive FileChgr1-Dev1 - File1 FileChgr1 Dev2 - File1	Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as the device type.
FileChgr2-Dev1 - File2	_ Device type: Virtual Tape Drive ✔
FileChgr2-Dev2 - File2	Create new
	VIITUAI TAPE (Specifying a folder path VIITUA
	Use folder path: Browse
	Tape drive

Figure 7.10.1 (a) - Adding a virtual tape drive

6. After creating the virtual storage device, proceed to the **Creating a Media pool** section to continue with the backup process.

Creating a Physical Storage device

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

1. Navigate to Storage > Backup > Storage media > Storage devices.

- 2. Click +Add a Tape drive under Tape drives section.
- 3. Select **Device type** as **Tape Drive**.
- 4. Select the tape device.
- 5. Enter the **Media type** for the device.
- 6. Enter the tape drive name.
- 7. Click Apply.

Tape drives:	Details
+ Add a Tape drive FileChgr1-Dev1 - File1 FileChgr1-Dev2 - File1 FileChgr2-Dev1 - File2 FileChgr2-Dev2 - File2	Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as the device type Device type: Tape Drive Tape device:Status:
	Media type for device: (Eg.: DDS-5 / DLT / DAT) Tape drive name:

Figure 7.10.1 (b) - Adding a physical tape drive

8. After creating the physical storage device, proceed to the **Creating a Media pool** section to continue with the backup process.
Creating a Media pool

To create a media pool, follow these steps:

- 1. Navigate to **Storage > Backup > Storage media > Media pools and tapes**.
- 2. Select +Add a Media Pool under Media Pools section.
- 3. Enter the pool name.
- 4. Enter the appropriate value for recycle and click **Add Pool**.

Media Pools:	Details:
+ Add a Media Pool	Pool name: Pool1
File	Recycle tapes: 🗹 Enabled
Scratch	Recycle tapes every: 90 days 🗸
	Add Pool

Figure 7.10.1 (c) - Create a Media pool

Add a new tape to backup service

Note: Perform this step only once for each tape.

To add a new tape to backup service, follow these steps:

- 1. Navigate to **Storage > Backup > Storage media > Storage devices**.
- 2. From Tape drives, select Physical tape or Virtual tape to proceed with Add to backup service.
- 3. Select Add new tape to backup service from the Actions.
- 4. Enter the Tape / volume name.
- 5. Select the **Media pool**.
- 6. Click Apply.

pe drives:	Details
eChgr1-Dev1 - File1 👞 Ň	ote: If you have a Tape drive, you may add it here by selecting 'Tape
eChgr1-Dev2 - File1 dr eChgr2-Dev1 - File2	ive' as the device type
eChgr2-Dev2 - File2	Device: IBM - DDS Gen5 (Connected)
M - DDS Gen5 - DDS-5 S	torage service aware: yes
	Type: DDS-5
	Version:
	Serial: JR00L80
	Tape media: Loaded
	TapeAlert monitor: Supported, current status is okay
E	rror correction details:
	Total errors corrected: reads: 0 writes: 0
I	Data processed (GB): reads: 0.001 writes: 0.000
	Total uncorrected errors: 0 writes: 0
	Storage device: /dev/nst0
	Name: IBM - DDS Gen5
	Actions: Add new tape to backup service
	Madia paoli (Deeld and

Figure 7.10.1 (e) - Adding a new tape to backup service using the physical tape drive

Tape drives:	Details
FileChgr1-Dev2 - File1 FileChgr2-Dev1 - File2 FileChgr2-Dev2 - File2	Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as the device type
IBM - DDS Gen5 - DDS5	Device: /virtual_tape
Virtual lapeDrive'i - Virtual lape	Storage service aware: yes
	Type: Virtual Tape
	Storage device: Device not mounted, currently attached to /dev/nvme0n1p5
	Space: Total: 344.151 GB, used: 7.411 GB, free: 336.740 GB
	Name: VirtualTapeDrive1
	Actions: Add new tape to backup service 🗸
	Tape / volume name: VirtualTape1
	Modia pools (Peeld, etc.)

Figure 7.10.1 (f) - Adding a new tape to backup service using the virtual tape drive

Detaching the tape drive

To detach the tape drive from the storage, follow these steps:

- 1. Navigate to Storage > Backup > Storage media > Storage devices.
- 2. Select the tape drive under Tape drives.
- 3. Select the Detach tape drive from backup service option from Actions.



Figure 7.10.1 (g) - Detaching the tape drive

Attaching the tape drive to backup service

To attach tape drive to backup service, follow these steps:

- 1. Navigate to Storage > Backup > Storage media > Storage devices.
- 2. Select the appropriate tape drive under Tape drives.
- 3. Select the Attach tape drive to backup service option from Actions.

4. Click Apply.

Note: The Attach tape drive option must be used after Adding a tape drive.

Storage devices:	
Tape drives:	Details
Tape drives: FileChgr1-Dev1 - File1 FileChgr1-Dev2 - File1 FileChgr2-Dev1 - File2 FileChgr2-Dev2 - File2 IBM - DDS Gen5 - DDS5	Details Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as the device type Device: IBM - DDS Gen5 (Connected) Storage service aware: yes Type: DDS5 Version: Serial: JR00L60 Tape Alert monitor: Supported, current status is okay Error correction details: Table processing of the status of the
	Total errors corrected: reads: 0 writes: 0
	Total uncorrected reads: 0 writes: 0 errors:
	Storage device: /dev/nst0 Name: IBM - DDS Gen5 Actions: (Attach tape drive to backup service V

Figure 7.10.1 (h) - Attaching the tape drive

7.10.2 Backup

Sections of the interface

1. Summary

This section shows a brief summary of the backup configuration. The user can choose to start or 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: Disable Auto Start	Stop Restart
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 7.10.2 (a) - Summary section of the backup page

2. Jobs

Using this option, you can view the list of configured jobs and modify them if required.

Jobs:	
Jobs:	Details:
Backup jobs	Job name: backup1
backup1 backup2 Verify jobs	Status: Scheduling enabled, not running next run scheduled at 23-Apr-13 23:05
verify_job1	File group selected: System configuration data
restore_job1	Backup type:Select a Level 🔻
	Schedule : Monthly cycle -
*	Storage device: VirtualTapeDrive1 -
	Media pool: Default
	X Disable Save Changes Save and Run
Previous runs of selected job:	Details:
2013-04-22 23:05:04	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 7.10.2 (b) - Configured jobs with details

3. Creating a Backup Job

Using this option, you can add a backup job, after the storage device configuration.

The **System Configuration** is the default file group template. 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. To create a job on a new storage device and media pool, add them to the **Storage media** page.

To create a backup job, follow these steps:

- 1. Navigate to **Storage > Backup > Add a backup/restore job**.
- 2. Select the job type from the drop-down menu and enter the job name.
- 3. Select the **Template name, Schedule, Storage device**, and **Media pools**.
- 4. Click Apply.

Job type:	Backup data 🗸
New job name:	
Template name:	Full Set 🗸
Schedule :	WeeklyCycle ~
Storage device:	Select a Device 🗸
Media pools:	Default 🗸

Figure 7.10.2 (c) - Creating a backup job

4. Editing a Backup Job

To edit a backup job, follow these steps:

- 1. Navigate to **Storage > Backup > Jobs**.
- 2. Expand the **Jobs** option and select the **Backup job** under the **Jobs** section.
- 3. Select the required backup option from the **Backup type** drop-down menu. The available options are:
 - Full backup
 - Differential backup
 - Incremental backup

4. Click Save Changes.

Jobs:	Details:
Backup jobs	Job name: job1
job1 Verify jobs Restore jobs	Status: Scheduling enabled, not running next run scheduled at 20-May-25 23:0
RestoreFiles	File group selected: Full Set
	Backup type:Select a Level 🗸
	Schedule : -Select a Level-
	Storage device: Differential backup
•	Media pool: Incremental backup

Figure 7.10.2 (d) - Editing a backup job

5. Running a Backup Job

To run a job, follow these steps:

- 1. Navigate to **Storage > Backup > Jobs**.
- 2. Select the required job.
- 3. Click Save and Run.

Jobs:	Details:
Backup jobs	Job name: Job1
Job1	Status: Scheduling enabled, not running
Restore jobs	next run scheduled at 27-May-25 23:0
RestoreFiles	File group selected: Full Set
	Backup type: Full backup
	Schedule : WeeklyCycle 🗸
	Storage device: IBM - DDS Gen5 🗸
•	Media pool: Pool1

Figure 7.10.2 (e) - Running a backup job

7.10.3 Restore

To restore data, the steps are the same for both virtual and physical tape drives. Follow the steps below to restore the backed-up data:

- 1. Navigate to **Storage > Backup > Jobs**.
- 2. Select the Restore Files option under Jobs.
- 3. From the drop-down menu, select the job to be restored.

- 4. Browse and select the location to restore the data.
- 5. Click Save and Run.

Jobs:	Details:
Backup jobs	Job name: RestoreFiles
Restore jobs	Status: not running
RestoreFiles	Job to Restore: BackupClient1 - 2025-04-22 23:05:03 ▼
	Where to //tmp/bacula-restores Browse
	Save and Run

Figure 7.10.3 (a) - Restore Job

7.10.4 Verify

Creating a Verify backup job

To verify data, the steps are the same for both virtual and physical tape drives. Follow the steps below to verify the backup job:

- 1. Navigate to Storage > Backup > Add a backup/restore job.
- 2. Select Verify Backup data from Job type.
- 3. Enter the New job name.
- 4. Select the Template name, Schedule, Storage device, and Media pools.
- 5. Click Apply.

Job type:	Verify Backup data
New job name:	
Template name:	Select a File group 🗸
Schedule :	Select a Schedule V
Storage device:	Select a Device 🗸
Media pools:	Select a Media Pool 🗸

Figure 7.10.4 (a) - Creating a Verify backup job

Running a Verify backup job

To run a verify backup, follow the steps below:

- 1. Navigate to Storage > Backup > Add a backup/restore job.
- 2. Select Verify Job under Jobs.
- 3. Click Save and Run.

:	
Jobs:	Details:
Backup jobs	Job name: VerifyJob
job1 Verify jobs	Status: Scheduling enabled, not running
VerifyJob	File group selected: Full Set
Restore jobs	Schedule : WeeklyCycle 🗸
RestoreFiles	Storage device: IBM - DDS Gen5 🗸
	Media pool: Pool1
-	🗙 Disable 🛷 Save Changes 🚺 🔊 Save and Run

Figure 7.10.4 (b) - Running a Verify backup job

7.10.5 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). The 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 it 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 the template, or you can delete the template if required.

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

Add a file group template	Template name: System Configuration]
Catalog	Folders to backup: //afs	
System Configuration	Add	
	- Remove	-
	Full path:	
	+ Settings:	
		-

Figure 7.10.5 (a) - 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 two 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 a Schedule template 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 recommended. Ensure that the time of day to run the backup is set correctly.

Schedule templates:			Deta	ails				
Add a Schedule template	Template name:	WeeklyCycle						
veeкiyCycleAnerbackup	Schedule:	1. Run Full backup 2. Run Differential b 3. Run Incremental	ackup backup					
	Remove			•				
		A Save T	omplate	C Delete	Template	1		
		Sale I	empiate	O				
	Schedule F Description: r	Run Full backup	p at 23:05	on sunda	y and o	n first w	veek of th	ne
	Schedule F Description: r	Run Full backup nonth	p at 23:05	on sunda	y and o	n first w	eek of th	ne
	Schedule F Description: r Schedule Setting Backup ty	Run Full backup nonth gs: rpe: Full backup	o at 23:05	on sunda	y and o] n first w	eek of th	ne
	Schedule F Description: r Schedule Setting Backup ty Scheduling mo	Run Full backup month gs: rpe: Full backup ode: Week days	• at 23:05	i on sunda	y and o	j	veek of th	ne
	Schedule F Description: r Schedule Setting Backup ty Scheduling mo Week da	Run Full backup month gs: rpe: Full backup ode: Week days ays: Mo Tu We	p at 23:05	i on sunda	y and o	n first w	reek of th	ne
	Schedule F Description: r Schedule Setting Backup ty Scheduling mc Week da Weeks in the moi	Run Full backup month gs: ype: Full backup ode: Week days ays: Mo Tu We nth: Specific weel	p at 23:05	i on sunda	y and o	n first w	reek of th	ne
	Schedule F Description: r Schedule Setting Backup ty Scheduling mc Week da Weeks in the mod	Run Full backup month gs: (pe: Full backup yde: Week days ays: Mo Tu We nth: Specific week Week	p at 23:05	i on sunda a Su nonth ∽ /e Th Fr Sa	y and o a Su	n first w	reek of th	ne
	Schedule Description: F Schedule Setting Backup ty Scheduling mc Week da Weeks in the more	Run Full backup month gs: rpe: Full backup ode: Week days ays: Mo Tu We specific weet Week Week First	p at 23:05	i on sunda a Su nonth ~ /e Th Fr Sa 3 4 5 6	y and o a Su 7	n first w	reek of th	ne
	Schedule Description: F Schedule Setting Backup ty Scheduling mc Week da Weeks in the moi Specific Week(s	Run Full backup month gs: Fpe: Full backup ode: Week days ays: Mo Tu We Neek Week Week First So of Second	• • • •	i on sunda a Su nonth ✓ /e Th Fr Sa 3 4 5 6 0 11 12 13	y and o a Su 7 3 14	n first w	reek of th	ne
	Schedule F Description: F Schedule Setting Backup ty Scheduling mc Week da Weeks in the mou Specific Week(s the mou	Run Full backup month gs: Full backup vek days wo Tu We Note: Veek veek veek veek of Second nth: Third	 p at 23:05 p at 23:05 Th Fr Sak(s) of the m Mo Tu W 1 2 3 8 9 10 15 16 17 	a Su a Su nonth ✓ Fe Th Fr Sc 3 4 5 6 0 11 12 12 7 18 19 20	a Su 7 3 14) 21	n first w	reek of th	ne

Figure 7.10.5 (b) - 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.

7.10.6 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 the 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 the status of the drive, and attach or detach from the backup manager service, and add physical tape cartridges to be used by the backup service.

Tape drives:	Details
+ Add a Tape drive FileChgr1-Dev1 - File1 FileChgr1-Dev2 - File1	Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as the device type
FileChgr2-Dev1 - File2	Device type: Virtual Tape Drive 🗸
FlieChgiz-Devz - Fliez *	Create new virtual Allocating space from a storage pool
	Storage Pool for new volume:Select a Storage Pool V
	New volume name: VirtualTapeDrive1
	New volume size: GB ~
	Folder to use: Browse.
	Tape drive VirtualTapeDrive1
	Apply

Figure 7.10.6 (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.
 The backup service will make use of the next drive only if the current drive is full.
 - Rewind/Eject/Erase the tape: These actions are used to Rewind/Eject/Erase 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 FileChgr1-Dev1 - File1 FileChgr1-Dev2 - File1	Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as the device type
FileChgr2-Dev1 - File2	Device type: Virtual Tape Drive 🗸
FileChgr2-Dev2 - File2	Create new virtual Allocating space from a storage pool
	Storage Pool for new volume:Select a Storage Pool V
	New volume name: VirtualTapeDrive1
	New volume size: GB 🗸
	Folder to use: Browse.
	Tape drive name: VirtualTapeDrive1
	Apply

Figure 7.10.6 (b) - Add or modify schedule templates

Tape drives:	Details
+ Add a Tape drive FileChgr1-Dev1 - File1 FileChgr1-Dev2 - File1	Note: If you have a Tape drive, you may add it here by selecting 'Tape drive' as th device type
FileChgr2-Dev1 - File2	Device: /tmp
FileChgr2-Dev2 - File2	Storage service aware: No (Storage service needs to be restarted)
	Type: File1
	Storage device: Device not mounted, currently attached to /dev/sda6
	Space: Total: 69.938 GB, used: 5.228 GB, free: 64.710 GB
	Name: FileChgr1-Dev1
	Actions: Add new tape to backup service 🗸
	Tape / volume name: Volume1
	Media pool: Default 🗸

Figure 7.10.6 (c) - Add or modify schedule templates

2. Media pools and tapes

The 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:	Details:
+ Add a Media Pool	Pool name: pool2
pool1	Recycle tapes: 📝 Enabled
-	Recycle tapes every: 90 days

Figure 7.10.6 (d) - Add or modify schedule templates

Media pools and tape	S:
Media Pools:	Details:
+ Add a Media Pool	Pool name: Default Recycle tapes: V Enabled Recycle tapes every: 365 days -
Tapes cartridges / volumes:	Details:
tap1	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.

Figure 7.10.6 (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 the 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 depend 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.

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 the 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. (Refer to the IETF iSCSI Standard, RFC 3720, for further details, at <u>www.ietf.org</u>.) 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 (for example, 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 figure describes the iSCSI Target stack, its relationship with other subsystems in the Appliance, and a brief list of the features that it provides.

Copyright $\ensuremath{\mathbb{C}}$ 2025. Chelsio Communications. All Rights Reserved.

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 to 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 two IP addresses, one each on a different Network device, and the Initiator connects to the target on both the IP networks available. These two IP networks are physically connected over different Ethernet networks. Refer to 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 two 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.

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, 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:

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 target 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 to the local (primary) node.



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

Targe	t service details:	
	Service status:	installed, Auto Start Enabled, Target configured and running
	Actions:	Disable Reload
	Target mode:	ULP 👻 Vpdate
	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

Kernel and application installation details

The support details of Target software are shown here.

abled, built as module, module currently loaded
).0
stalled
t 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 other settings. 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. 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 grey are 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. 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:

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:	[iqn.2025-03.com.testing.vertesx1:2 (in iqn.yyyy-mm format, example: iqn.2004-05.com.chelsio.target1)				
Target Alias:	Target-2 (A unique short identifier/name for this Target)				
LUN:	LUN Device: RAM Disk Size:	Ptest/Memory: 5.00 GB 32 MB			
Portal Group:	IP Address: TCP Port:	102.50.50.192 ✓ 3260 ✓ Default Valid TCP port can be from 1 - 65535			
Target Redirection:	Enable Shadowmode: Redirection IP Address: Port:	Yes 102.50.50.235 3260			
Create iSCSI Target					

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.



Figure 8.6(a) - *iSCSI Target Summary page*

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 are 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:

Кеу	Valid Values	Default Value	Description
TargetName	" <target name="">"</target>		A worldwide unique iSCSI target name.
TargetAlias	" <target alias="">"</target>		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" "CRC32C"	None	To enable or disable iSCSI header Cyclic integrity checksums.
DataDigest	"None" "CRC32C"	None	To enable or disable iSCSI data Cyclic integrity checksums.
MaxConnection s	1 to 65536	4	Initiator and target negotiate the maximum number of connections requested/acceptable.
MaxRecvDataSe	512 to	8192	To declare the maximum data segment

Кеу	Valid Values	Default Value	Description
gmentLength	16777215 (224 - 1)		length in bytes it can receive in an iSCSI PDU.
TargetSessionM axCmd	1 - 2048	32	To declare the maximum outstanding iSCSI commands per session.
InitialR2T	"Yes" "No"	No	To turn on or off the default use of R2T for unidirectional and the output part of bidirectional commands.
MaxOutstandin gR2T	1 to 65535	8	The maximum number of outstanding R2Ts per task.
ImmediateData	"Yes" "No"	Yes	To turn on or off the immediate data.
FirstBurstLengt h	512 to 16777215 (2 ²⁴ - 1)	65536	The maximum negotiated SCSI data in bytes of unsolicited data that an iSCSI initiator may send to a target during the execution of a single SCSI command.
MaxBurstLengt	512 to	262144	The maximum negotiated SCSI data in

Кеу	Valid Values	Default Value	Description
h	16777215 (2 ²⁴ - 1)		bytes, of a Data-In or a solicited Data-Out iSCSI sequence between the initiator and target.
DefaultTime2W ait	0 to 3600	2	The minimum time, in seconds, to wait before attempting an explicit / implicit logout or connection reset between initiator and target.
DefaultTime2Re tain	0 to 3600	20	The maximum time, in seconds, after an initial wait.
ErrorRecoveryL evel	0 to 2	0	To negotiate the recovery level supported by the node. <i>Chelsio only supports 0</i> .
DataPDUInOrde r	"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> .

Кеу	Valid Values	Default Value	Description
DataSequencel nOrder	"Yes" "No"	Yes	To indicate the Data PDU sequences must be transferred in continuously non- 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 "No".</i>

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.

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.

• 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 grey 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.

Add a LUN by specifying an existing device	
Select the Storage device:	Select a device 🗸
RAM Disk size:	32 MB
Add a LUN by allocating space from a storage pool	
Storage Pool for the LUN:	Select a Storage Pool 🗸
New volume name:	
New volume size:	GB 🛩
Add Multiple LUNs	
Free LUN list (multi-select)	
✓ 0. Ptest/Memory: 5.00 GB	
✓ 1. sdc: 7451.98 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 (for example, 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.

- 1. Click the Add a LUN by specifying an existing device option to enable it.
- 2. Select from the available storage device (listed in black) to use as iSCSI Target LUN. Devices listed in grey 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.
- 3. 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.

- 1. Click the Add a LUN by allocating space from a storage pool option to enable it.
- 2. Select the storage pool, to create the logical volume.
- 3. Specify a name for the logical volume and size.
- 4. 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

- 1. Click the Add Multiple LUNs option to enable it.
- 2. Select multiple devices in the list by clicking on them.
- 3. 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)	Edit network portals	list:		
✓ 1. IP: 102.50.50.192 - Port: 3260	S Delete Portal			
	Edit selected network portal:			
	IP address:	102.50.50.192		
	TCP port:	Use default iSCSI service TCP port:		
		Valid TCP port can be from 1 - 65535		
		Redirect List		
	Portal redirection list:			
	Add portal to redirect	Add Remove		
Save network portals configuration				

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.

- Add a network portal:				
IP address:	Select an IP address 🗸			
External IP for Redirection:				
TCP port:	3260 Valid TCP port can be from 1 - 65535			
Use default iSCSI service TCP port:				
	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 require to authenticate this target as a valid target.

ummary:	
Access Control:	Disabled
CHAP	
CHAP authentication policy:	Disabled V
Authentication type:	One-way (Client only)
CHAP Mutual authentication	Target parameters:
Username:	
Password:	show password
	Apply
Clients connected:	None

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.

lients:			
Clie	nts / Initiators		
🛛 🛷 🛑 [Offline	e] iqn.2004-05.com.chelsio.target1		
– 🏮 Current	status		
Current Current status:	Disconnected		
Current Current status: Sessions:	status Disconnected 0		

Figure 8.9(b) - Clients list with settings shown below, and status on the left

2.1. Client settings

- Clients list: This area displays the list of currently connected and configured initiators. Click on an initiator to view its configuration information below.
- 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 Contro	
Status:	Access Control entry is configured. Client IP address restriction configured. Target IP address restriction not configured.
IQN name:	iqn.2004-05.com.chelsio.target1
Client IP address:	10.193.187.56 Eg: 192.168.1.1, 192.168.1.2
Target IP address:	Any IP 🗸
LUNs access:	All LUNs are accessible with default permissions All CONSTRUCT Second seco
Save Access Control	I settings Delete ACL

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.

— 🔒 CHAP authentication:	
CHAP secret:	is set
CHAP secret (length of 12 - 16, optional):	Show password
	Apply Selete

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.

- 1. Select the target for which you want to enable CHAP under the iSCSI module. Click on the **Clients/initiators** sub-module.
- 2. Under the Summary section, select Enforced in the CHAP authentication policy drop-down.
- 3. The Authentication type will be **one-way** by default.
- 4. Click **Apply**.
- 5. 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.
- 6. Click Save Client settings.
- 7. 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.
- 8. 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.
- 9. Click Apply.

Note: For instructions on how to access an iSCSI target with CHAP authentication, refer to the **iSCSI Initiator** section.

Example: Configuring Access control for iSCSI Clients/initiators

Ensure that no iSCSI service is currently running before you attempt to set access control for a target.

- 1. Select the target for which you want to configure access control under the iSCSI module. Click on the **Clients/initiators** sub-module.
- 2. In the **Summary** section, select **Enabled** from the **Access Control** drop-down.
- 3. Click **Apply**.
- 4. Expand the **Clients** section to select the initiator or add a new initiator using the **Add a client** section.
- 5. Select the initiator from the list and expand the **Access Control** subsection.
- 6. If not specified while adding the client, enter initiator IP address.
- 7. 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.
- 8. 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.
- 9. 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 also be specified. If the initiator is part of a cluster environment, 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:	iqn.2004-05.com.chelsio.target1 Eg: iqn.2004-05.com.chelsio.target1			
Client IP address:	10.193.187.56 Eg: 192.168.1.1, 192.168.1.2			
Target IP address:	Any IP 🗸			
CHAP secret (length of 12 - 16, optional):	Show password			
	Save Client Settings			

Figure 8.9(f) - Add client section
9. NVMe-oF

9.1 NVMe-oF overview

Non-Volatile Memory Express (NVMe-oF) protocol enables high-performance, low-latency access to NVMe storage devices over a network. It allows NVMe commands to transmit over various network, such as Ethernet and InfiniBand making it possible to share NVMe storage resources across multiple servers in a data center or cloud environment. Figure 10.1 shows the NVMe-oF page.

Sections of the interface

1. Target service details

Service status displays the status of the NVMe service. Using the **Actions** button, you can disable, reload, or restart the NVMe service. **Targets configured** displays the number of configured targets. **Targets running** displays the number of running targets.

2. Kernel and application support

Kernel Module displays the status and descriptions of the kernel module. Click **Load Module** to load the kernel module. **Test Load** displays the status and descriptions of the kernel module. Click **Load Module** to load the kernel module. **Test Load** displays the status of the load, **Module Release Version** displays the version of the module, and **Control Utility** indicates whether the utility is installed.

et service details:	
Service status:	installed, Auto Start Enabled, target not configured
Actions:	X Disable Colorad Reload
Targets configured:	0 (choose 'Create New Target' under Targets in the navigation menu)
Targets running:	0
el and application support:	
Kernel Module: Module Release Version:	Enabled, built as module, module currently loaded
Control Utility:	installed

Figure 10.1 – The NVMe-oF page

9.2 Creating a new NVME Target

New target settings

1. Target Name:

This setting is automatically suggested. You may use the suggested value or specify other settings. The IQN name must meet the requirements as specified in the NVMe-oF IETF standard RFC 3720.

2. Target Alias:

This setting specifies the human-friendly name when the Target is identified. This name is displayed on the navigation menu on the left. Client initiators accessing the Target will also report the alias name.

3. LUN:

Select a single LUN device which can be shared by this NVMe-oF target. This LUN could be a Volume management device or a disk, or a RAM disk device. Specify the custom size for a RAM disk device.

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

Warnings

- a) The devices shown in grey are 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:

Select the Protocol from the drop-down. The supported protocols are **TCP** and **iWARP**. Select the IP address and enter the TCP port that the target should provide NVMe-oF service. The default NVMe-oF TCP port is 4420. Change this port if you configure the client initiators to use the different port that you specify. Additional LUNs and portals can be specified after adding the new target.

Target Name:	nqn.2025-03.com.testing.vertesx1:1 (in iqn.yyyy-mm format, example: nqn.2004-05.com.chelsio.target1)				
Target Alias:	Target-1 (A unique short identifier/name for this Target)				
LUN:	LUN Device: RAM Disk Size:	Ptest/Memory: 5.00 GB			
Portal Group:	Protocol: IP Address: TCP Port:	TCP ✓ 102.50.50.192 ✓ 4420 ✓ Default Valid TCP port can be from 1 - 65535			
		Create NVME Target			

Figure 10.2 - Create a new NVMe-oF Target

9.3 NVMe Target LUN configuration

Sections of the interface

1. Current LUN configuration:

The storage on the system that is shared by this NVMe Target, is listed here. Options to modify the list are also available.

0. Ptest/Memory - 5.00 GB - LUNNO: 0 Edit selected LUN: Permission: Read-Write Read Caching: Enabled	NO: 0
Edit selected LUN: Permission: Read-Write Read Caching: Enabled disal	
Permission: Read-Write Caching: Permission: Permissio	Edit selected LUN:
	Permission: Read-Write Read-Onl
(Can only be changed when the target is s	(Can only be changed when the target is stopp
RAM Disk size: MB	RAM Disk size: MB
LUN Device:Select a device	LUN Device:Select a 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, where the data is discarded.

Ensure 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 or NVMe-oF service 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.

• 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 grey 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 NVMe Target are shown here.

Select the Storage device: Select a device	Add a LUN by specifying an exis	ting device		
RAM Disk size: 32 MB Add a LUN by allocating space from a storage pool Storage Pool for the LUN:Select a Storage Pool V New volume name: New volume size: GB V Add Multiple LUNs Free LUN list (multi-select) V 0. Ptest/Memory: 5.00 GB 1. sdb: 465.76 GB V 2. sdc: 7451.98 GB		Select the Storage device:	Select a device	\sim
Add a LUN by allocating space from a storage pool Storage Pool for the LUN:Select a Storage Pool V New volume name: New volume size: GB V Add Multiple LUNS Free LUN list (multi-select) V 0. Ptest/Memory: 5.00 GB V 1. sdb: 465.76 GB V 2. sdc: 7451.98 GB		RAM Disk size:	32 MB	
Storage Pool for the LUN:Select a Storage Pool New volume name: New volume size: Add Multiple LUNs Free LUN list (multi-select) V 0. Ptest/Memory: 5.00 GB V 1. sdb: 465.76 GB V 2. sdc: 7451.98 GB	Add a LUN by allocating space f	rom a storage pool		
New volume name: New volume size: GB ✓ Add Multiple LUNs Free LUN list (multi-select) ✓ 0. Ptest/Memory: 5.00 GB ✓ 1. sdb: 465.76 GB ✓ 2. sdc: 7451.98 GB		Storage Pool for the LUN:	Select a Storage Pool-	V
New volume size: Constraints (multi-select) Constraints (multi-select) Constraint		New volume name:		
Add Multiple LUNs Free LUN list (multi-select) ✓ 0. Ptest/Memory: 5.00 GB ✓ 1. sdb: 465.76 GB ✓ 2. sdc: 7451.98 GB		New volume size:	GB 🛩	
Free LUN list (multi-select) ✓ 0. Ptest/Memory: 5.00 GB ✓ 1. sdb: 465.76 GB ✓ 2. sdc: 7451.98 GB	Add Multiple LUNs			
 ✓ 0. Ptest/Memory: 5.00 GB ✓ 1. sdb: 465.76 GB ✓ 2. sdc: 7451.98 GB 	Free LUN list (multi-select)			
 ✓ 1. sdb: 465.76 GB ✓ 2. sdc: 7451.98 GB 	✓ 0. Ptest/Memory: 5.00 GB			
✓ 2. sdc: 7451.98 GB	✓ 1. sdb: 465.76 GB			
	✓ 2. sdc: 7451.98 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 NVMe-oF 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 (for example, 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 NVMe 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 NVMe Target LUN. Please note that only free/unassigned devices can be used.

- 1. Click the Add a LUN by specifying an existing device option to enable it.
- 2. Select from the available storage device (listed in black) to use as NVMe Target LUN. Devices listed in grey 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.
- 3. 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 NVMe Target LUN by creating a logical volume.

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

- 1. Click the Add a LUN by allocating space from a storage pool option to enable it.
- 2. Select the storage pool, to create the logical volume.
- 3. Specify a name for the logical volume and size.
- 4. 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 NVMe Target LUNs.

Using this method, you can add multiple existing storage devices as NVMe Target LUNs

- 1. Click the Add Multiple LUNs option to enable it.
- 2. Select multiple devices in the list by clicking on them.
- 3. Click the **Add LUN** button.

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

10. File Sharing

10.1 File sharing overview

The file sharing protocols and applications are more commonly used by Desktop and Notebook clients. The 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), and File Transfer Protocol (FTP).

- 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.

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 are also listed.



Figure 10.1 (a) - CIFS details

Note: Clients accessing CIFS shares may experience timeout issues, if the 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., are shown here. The number of shares configured and currently active, and the status of any NIS domain membership are also listed.



Figure 10.1 (c) - NFS details

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: X Disable 6 Res	tart
Configured Shared folders: Currently Shared folders:	0 0

Figure 10.1 (e) - FTP details

4. Service commands

- 4.1. 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.
- 4.2. Reload: This refreshes the service, applying any changes in the configuration file.
- 4.3. 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, if 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.
- 1.6. Allow users with blank passwords: 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, the server will accept and process other I/O operations, while the requested read or write function continues in the background.



Figure 10.2 (a) – CIFS settings

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 enabled: 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 listening 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 listening 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.



Figure 10.2 (b) – NFS settings

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.



Figure 10.2 (c) – FTP settings

10.3 File sharing configuration

To configure file sharing, select the path to share option. 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, and NFS), 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 to set guest or regular user access and lists the shared folder with a different share name. FTP authentication is transmitted over plaintext, and it is inherently insecure. It is not recommended 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.

——— Share list ————		
🖋 1. /home	[Shared over CIFS]	
2. /var/lib/samba/drivers	[Shared over CIFS]	
3. /var/tmp	[Shared over CIFS]	

Figure 10.3(a) - List of folders shared on the system

/home	Redit folder ownership and permissions	
Status: Folder exists	Storage Device used:	Configured, but currently not mounted!
Owner: root	Group: root	Specials: None
Access: rwx	Access: rwx	Others: rwx

Figure 10.3(b) - Shared folder list with index number, ownership, and permission 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
0	Information:	Folder and share permissions
	Please ensure adding the sha ACLs for allow by editing the s	you check / change ownership and permissions for the folder to the correct user and group after re. ing or restricting hosts or users can be edited after adding the share. Other settings can also be set share after addition.
Folde	er to share:	/ Browse
Share	e cluster mode:	Standalone 🗸
Share	e Protocol:	 CIFS NFS FTP HTTP
		Apply Cancel

Figure 10.3(c) - Adding a shared folder

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 pop-up menu.

В	rowse	
Name	Owner:Group	Last modified
Location:	1	
🗾 afs	root:root	16-May-2022 1
🤳 home	root:root	10-Mar-2025 1
🤳 media	root:root	16-May-2022 1
🤳 mnt	root:root	16-May-2022 1
🤳 mnt2	root:root	10-Mar-2025 1-
-		
Create fold	ler:	+
Select fold	er: /afs	Image: A start of the start

Figure 10.3(e) – The *Browse for folder popup*

- Cluster IP addresses: Cluster IP using which you want to share the particular folder.
- Add share protocol settings: CIFS, NFS, and FTP 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**.



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.

- 1. Click on the **Add a shared folder**. This will navigate to a new page.
- 2. Click the **Browse** button to locate the folder to share or enter the path manually if the location is already known.
- 3. 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.
- 4. 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.
- 5. Provide a share name (except for NFS) and other related settings and click **Apply**.

6. If the shared 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 to the <u>High Availability</u> section.

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 two 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)
	O Domain user:
Folder Group:	● Local Group root ✓
	O 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 permission	s
Note: Advanced folder options	can be changed only for mounted shares
🛕 Warning:	
0 11 1 1	ons changes will take effect immediately and any clients connected will
be allowed or denied acc	cess 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

- 1. In the **Share list**, select the folder for which you want to change ownership and permissions.
- 2. Click on the **Edit folder ownership and permissions** button. This will navigate to the folder properties page.
- 3. 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.**
- 4. 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. From the pop-up that appears, select domain and enter a valid group name. NIS/ADS domain should be configured to make a Domain group as the folder group. Click **Apply**.
- 5. Configure global permissions for the folder like Read, Write, and Execute permission for the folder owner, group, and other users (users not belonging to the folder group).
- 6. 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.

7. Click Apply changes.

2. Per shared folder shares list

NFS, CIFS, and FTP shares of the shared folder are listed below the folder's properties.

NFS shares:						
1. /home						
Mode: Standalone						
- Share details	Edit	t sha	are O	Disable	Delete	
Saved settings:			Running settings:			
Saved settings: Shared to:	all:		Running settings: Shared to:	Ca	ching: enabled	
Saved settings: Shared to:	all: caching: enabled		Running settings: Shared to:	ca	ching: enabled eck subtree: no	
Saved settings: Shared to: * Type: all	all: caching: enabled check subtree: no coalesce writes: enabled	Î	Running settings: Shared to: * Type: all	ca ch co file	ching: enabled eck subtree: no alesce writes: enabled system id: 9029	í

Figure 10.3(g) – NFS share with details and actions

CIFS shares:		
2) User's home folder (if applicable)		
Mode: Clustered [offline]		
- Share details	Edit share	Delete
Saved and Running settings:		
Shared to:	Share settings:	
Only Clients allowed:	browseable = no	
All	inherit acls = Yes	
Specific Clients denied:	read only = No	
None	writable = yes	
Only Users allowed:	•	

Figure 10.3(h) – CIFS share with details and actions
(1	FTP shares:			
	3) FTP_Share: /home folder Mode: Standalone			
	- Share details	Edit share	Enable	Delete
	Saved and Running settings:			
	Shared to:			
	authenticated users			

Figure 10.3(i) – FTP share with details and actions

2.1. Per share properties and actions

Current settings and actions for each CIFS, NFS, and FTP share are shown below the share.

2.1.1. Share Actions

• Edit share: This command will display a new page that allows you to edit the shared folder, or the share permissions and settings.

Enable / Disable share: This command sets the current state of the share. If the share is enabled, it is accessible to clients; And it is not accessible to clients when it is 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 the 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: /home
Shared to:	All (*) Constraints boot/subnet/netgroup
Change shared to:	Select 🗸 *
Specific settings for *:	
Permissions:	Read+Write ~
Client Root Access level:	As local Root (no_root_squash) ✔
Client User Access level:	As user (no_all_squash)
Write caching:	Write-back (async) V Note: Caching is only available if the underlying filesystem allows caching
🖶 More settings	
	Apply changes Discard changes

Figure 10.3(j) - Edit NFS share settings



Figure 10.3(k) - 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:	homes (for clients to access as \\< hostname >\< share name >)
Share comment:	Home Directories (Displayed to clients when accessing the share)
Show in shares / browse list:	
Permissions:	Read+Write 🗸
Write caching:	Write-back (async) Note: Caching is only available if the underlying filesystem allows caching
Changes to CIES shar	es will cause a service reload, interrupting any I/O in progress.

Figure 10.3(I) - CIFS Shares

b. Access Control: 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.

Guest / Anonymous access:			
Only Guest / Anonymous access allowed:	$\hfill\square$ (Enabling this will disable regular user access and all User ACLs		
+ Clients/Hosts ACL:			
+ Users / Groups ACL:			
Changes to CIFS shares will cause a service reload, interrupting any I/O in progress			

Figure 10.3(m) - 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.

1. Expand the Clients/Hosts ACL section.

- 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.
- 3. Similarly, you can restrict a particular client/host from accessing the CIFS share. To perform 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:

- 1. Expand the Users/Groups ACL section.
- 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 does not 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.

- 3. Similarly, you can also add a local or ADS group.
- 4. Click Add allowed user. The user/group added will appear in Allowed Users/Groups list.
- 5. Repeat steps 2, 3, and 4 to add more users/groups.

6. 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: FTP_Share (for clients to access as ftp://< hostname >/< share name >)				
Guest / Anonymous access:				
Only Guest access allowed:				
Note: This will disable access for regular users, depending on other FTP settings.				
	Apply changes			

Figure 10.3(n) - Edit FTP 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 also be modified.

Local Users and Groups:					
.ocal Users:					
🖋 root (administrator)		User name:	root		
🖌 bacula		Full name:	root		
chrony		UID:	0		
✓ sssd		Primary group:	root 🗸		
polkitd		CIFS mapped:	~		
 systemd-coredump 		User enabled:			
maint		Change password:			
UISS		Set password:			
nobody (quest)		Confirm password:			
System user: bin		😰 Delete User	Save changes		
System user: daemon					
System user: adm	_				

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 the NIS domain, adding a local user with the same name as the 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

- 1. In the Users section, under File sharing, expand the Local Users and Groups.
- 2. 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 log in and access CIFS shares).
- 3. In the Add a local user section, provide the user name.

- 4. 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.
- 5. Click Add User to add the user. The newly added user will appear in the Local Users list.
- 6. You can edit the above configuration and more options by selecting the user in the list.

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.

Status:	Installed, enabled, domain confid	ured, running			
Actions: X Disable Actions					
Current NIS settings:	Domain: nis.t Server: nala	iss.com nda			
Seved NIS Demain:	Domain: nis.uss.com NIS Servers: 10.193.184.188	Domain: Manual Server configuration:	nis.uss.com select		
Saved NIS Domain.	Default Yes Domain:	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 the 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 member	ship:	
Status:	Installed, enabled, domain configure	ed, not joined, running
Actions: X Disable	Restart	
Current ADS settings	Domain: Server:	
Caused ADC Demoins	Domain: TESTING.COM	Domain: TESTING.COM Domain Controller: 2k8r2.testing.com
Saved ADS Domain.	Default Domain: Yes	Delete Domain

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.

- 1. In the **ADS domain membership** section, click on the *Leave Domain now* checkbox located at the bottom right.
- 2. Enter the Administrative username and password.
- 3. Click **Apply changes**. The Administrative credentials will now be verified and the page will refresh automatically. The status box will display errors if any.
- 4. 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.
- 5. 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.

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:	/var/log/messages
Log Filter:	Full log (default)
Filter logs using custom text:	
Number of lines to display per page:	5000
1	/iew Log 🛐 Download Log File 🔯 Clear Log File

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 dropdown box controls which one:
 - Current kernel log equivalent to Linux's dmesg
 - /var/log/messages
 - IPMI BMC System Event Log

The two action buttons to the right of the View Log button operate on the selected log file:

- Download Log File Downloads the entire log file to the viewing pane.
- 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 menu. Choose the relevant filter that fits.



Figure 11.1(b) - Filter criteria for log files

The filter criteria choices details are:

- Full log (default) No filter applied
- iSCSI target Shows only the iSCSI target log entries

- NFS file sharing Shows only the NFS file sharing logs
- CIFS file sharing Shows only the CIFS (Windows) file sharing logs
- FC Target Shows only FC Target log entries
- Chelsio Network driver Shows only the Chelsio driver logs
- 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 select 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 be added to the bottom of 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.

For 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 then 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
CDU Usaga	Average CPU Usage	Average percentage usage of all the CPUs present in the system.
CPU USage	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 resets 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 click the **Upload** button. The Reset configuration to the 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.

Bac	kup System configuration:
	Download the backup of System configuration data as of now (Wed Mar 12 21:22:38 2025):
Rest	tore System configuration:
۲	Upload a previous backup of Choose File No file chosen System configuration data:
0	Reset configuration to installation defaults:
Note volu	Restoring configuration from a backup or to installation defaults will not affect any data on any data mes, but it will stop all services and restart the system.
Note	A system restart is required to apply all restored settings correctly!
-:~	ro 11 1(a) Sustam backup and restare of configuration data

Figure 11.1(e) - System backup and restore of configuration data

4. System software update

Since this is the initial release, the system software update will be included in upcoming releases.

System software update				
Upload System update file / ISO: Vpload System Update file and apply update				

Figure 11.1(f) - System software update

5. 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 the Chelsio USS QA team using this button. If not, contact Chelsio support team with the hardware profile info file at <u>support@chelsio.com</u>.

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.

Vendor:	Supermicro	Vendor:	American Megatrends Inc.
Model:	X9DR3-F	Version:	3.3 1.0
Serial Number:	VM2BS57383	Release date:	07/12/2018
Processor:		Chassis:	
Vendor:	Intel	Vendor:	Supermicro
Model:	Intel(R) Xeon(R) CPU E5-2687W v2 @ 3.40GHz	Asset Tag:	
Count:	2	Thermal status:	Safe
Cache:	CPU Internal L1: 512 kB, CPU Internal L2: 2 MB, CPU Internal L3: 25 MB	Power supply status:	Safe
Current speed:	3400 MHz	Туре:	Desktop
Memory:		Software:	
Туре:	DDR3 1600 MT/s	Operating system:	Unified Storage Server 4.0 29
Count:	: 497 8 GB: 4 No Module Installed: 12	Kernel:	Linux 6.6.33 #1 SMP PREEMPT_DYNAMIC Thi Jun 13 20:27:58 EDT 2024

Figure 11.1(j) – Hardware Profile Information

6. 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 and Power ON).

- Summary:		
BMC details:		
IPMI version: BMC firmware: Reset BMC Firmware:	2.0 1.30 Warm reset	✓ 😝 Reset BMC
(Only required if the BMC	is not responding)	
IPMI drivers status:		
ipmi_devintf ipmi_si	loaded loaded	
BMC sensors status:		
CPU1_DIMM CPU1_Temp CPU1_Vcore CPU2_DIMM CPU2_Temp CPU2_Vcore Fan1 Fan2 Fan3 Fan4 Fan5 Fan6 Fan6 Fan7 Fan8 Intrusion P1-DIMM1A_Temp P1-DIMM1B_Temp P1-DIMM1B_Temp P1-DIMM1B_Temp	1.560 Volts 0x0 discrete 0.944 Volts 1.560 Volts 0x0 discrete 0.936 Volts 3645.000 RPM 3510.000 RPM 3510.000 RPM 3780.000 RPM 3780.000 RPM 3780.000 RPM 3780.000 RPM 0x0 discrete 24.000 degrees C 24.000 degrees C	

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. The available options are always-off, always-on, and previous.

Chassis:	
Power restore policy: (after power loss)	always-off
, i ,	always-on always-on (recommended)
	previous

Figure 11.1(I) - *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:	00:25:90:0f:a2:51	
IP address type:	Static IP 🗸	
IP address:	10.193.185.191	
IP netmask:	255.255.252.0	
IP gateway:	10.193.184.5	
VLAN:	Enabled: DID:	(1 - 4094)
	apply	

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.

Maximum users:	10			
	ID: 6 - User: true - Level ID: 7 - User: true - Level	access access		
Users:	ID: 8 - User: true - Level ID: 9 - User: true - Level ID: 10 - User: true - Lev	: access : access el: access	•	
	User ID:			
	User name:			
	Set password:			
	Password:			Show passwo
	Privilege level:	Not set	\sim	
Edit user:	Callin / Callback:	Not set	\sim	
	Link authentication:	Not set 🗸		
	IPMI messaging:	Not set 🗸		
	User status:	Not set ✔		

Figure 11.1(n) - IPMI Users settings

7. 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, click the **Apply** button to apply the changes.

Description	Current setting	Change Setting	
Date:	12 March 2025	New Date: 12 March 2025 / Edit	
Time:	21:35:40	New Time: 21 - : 35 - : 40 - (hh:mm:ss)	
Timezone:	Asia - Kolkata +0530 IST	(UTC/GMT offset Code Continent/City) Timezone ✓ -1100 SST Pacific/Pago_Pago ✓ -1100 SST Pacific/Nidway ✓ -1100 -11 Pacific/Niue ✓ -1000 HST Pacific/Honolulu ✓ -1000 -10 Pacific/Honolulu ✓ -1000 -10 Pacific/Rarotonga ✓ -0930 -0930 Pacific/Marquesas ✓ -0900 HDT America/Adak	
Network Time synchronization:	Disabled V Currently stopped	NTP servers: 2.rhel.pool.ntp.org	
		Apply	

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 click **Apply**. This will fill in the **New Date** field under the **Change Setting** column.

Note: An option to change the date is available only if the *Network Time synchronization* option is disabled.
Description	Current setting	Change Setting
Date:	12 March 2025	New Date: 12 March 2025
Time:	21:36:15	New Time: 21 ↓ : 36 ↓ : 15 ↓ (h
Timezone:	Asia - Kolkata +0530 IST	UTC/GMT offset Code Contine Su Mo Tu We Th Fr Sa Timezone 1 -1100 SST Pacific/Pago_F 1 -1100 SST Pacific/Midway 10 11 12 13 14 15 -1100 -11 Pacific/Niue 2 3 24 25 26 27 28 29 -1000 -10 Pacific/Taniti 30 31 -1000 -10 Pacific/Raroton 2025 -0930 -0930 Pacific/Marqu Apply
Network Time synchronization:	Disabled V Currently stopped	NTP servers: 2.rhel.pool.ntp.org

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: An option to change time is available only if the *Network Time synchronization* option 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.

8. Alerting

Email alerting is provided for different types of failures or events that may occur in 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. **Auto Start** is enabled by default, hence the alerting service will run automatically when the system starts.

- Summary			
Service details:			
Alert service status:	Not installed		
Actions:			
SNMP Traps:	Enable		
Email alerts:	Enable (Recommended)		
Email Recipients:	aaaaaa@chelsio.com		
Sender Email Address:	Vertesx1@chelsio.com		
Mail Server:	stargate3.asicdesigne		
SMTP Port:	25		
SMTP Authentication:	Enable		
💉 Арр	ly		

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 the management station and the events for which alerts are to be received in the **Settings** section.

SNMP Traps:	Enable	
Trap Sink Hostname / IP Address:	102.55.55.56	
Trap Sink Community Port:	162	
	Add	
	Sink List	
Trap Sink List:		Remove

Figure 11.1(t) - Setting up SNMP alerts

Example: Setting up SNMP traps

Follow the steps mentioned below to set up SNMP traps:

- 1. Enable SNMP traps by clicking on the check box adjacent to **SNMP traps**.
- 2. Specify the hostname /IP address and port id of the management station and click **Add**. Perform this step to add more devices.
- 3. Click **Apply**.
- 4. Now under the **Settings** section, select the events for which alerts should be sent.
- 5. 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 username and password should be provided.

Email alerts:	Enable (Recommended)
Email Recipients:	aaaaaaa@chelsio.com
Sender Email Address:	Vertesx1@chelsio.com
Mail Server:	stargate3.asicdesigne
SMTP Port:	25
SMTP Authentication:	Enable
	Apply

Figure 11.1(u) - setting up Email alerts

Example: Setting up Email alerts

Follow the steps mentioned below to set up email alerts:

- 1. Enable email alerts by clicking on the check box adjacent to **Email alerts**.
- 2. Enter the email id of the recipients separated by commas.
- 3. Enter the email address from which you would like to receive the alerts in **Sender Email Address** in the format of<name>@<domain>.<extension>.
- 4. Enter the IP or host name of the mail server and SMTP port number.

- 5. If your SMTP server requires authentication, then enable SMTP authentication here and provide the username and password.
- 6. Click Apply.
- 7. Now under the **Settings** section, select the events for which alerts should be sent.
- 8. Click **Apply**.

2. Settings

Using the **Settings** section, user can customize various options, based on which SNMP traps/Email alerts will be sent.



Figure 11.1 (v) - Configuring alert settings

The available options are dependent on the hardware.

• 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.

- *IPMI* provides alerts on system health.
- *Network Devices* alerts will alert the user on a device losing or acquiring an Ethernet link.
- *Fibre Channel* alerts will alert the user on a device losing or acquiring the link on the FC port.
- *iSCSI target* alerts for an initiator login or logout, and security events such as an initiator ACL deny or CHAP authentication failure.
- *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.
- *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.
- *Software* or *Hardware RAID (SAS Raid)* alerting is provided for RAID array degrade / failure and rebuild events.
- *Cluster* provides alerts when cluster membership changes.
- *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.

9. 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 click the **Change password** button.

Old password:	
New password:	
Confirm password:	
	Change password

Figure 11.1 (w) - Administrator's account password change

10. 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 number of minutes from the drop-down dialog box (00 minutes being immediate), and click the appropriate button for either Halt / Shutdown or Reboot / Restart. If it is not an immediate action (set to 00 minutes), then either action can be canceled by clicking the **Cancel scheduled Shutdown** button. In HA mode, command to shutdown or reboot peer node from the local node is available.

O Shutdown / halt	immediately or after 00 🗸 minutes
😡 Reboot / restart	immediately or after 00 🗸 minutes
X Cancel schedule	shutdown / reboot

Figure 11.1 (x) - System Control for shutdown or reboot

The granularity of the number 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 does not exist.

This error occurs if your USS appliance is part of an NIS domain, and the name of the local user you are 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

Copyright ©2025. Chelsio Communications. All Rights Reserved.

337 | Page

The above command will list all the users in NIS domain database. Now try creating a local user again but with a name which does not appear on the list.

2. While running I/O operations, the USS appliance was rebooted abruptly (example: 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 the below steps to resolve:

- 1. Reseat the drive on which I/O operations were running.
- 2. Access the **Disk devices** section under **Storage**, and use the **Rescan all storage devices** button.

The drive appears in the list again.

Note: It is highly recommended that the appliance must be 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 a 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. 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:
 - 1. Reboot/boot-up USS machine.
 - 2. During boot, press any key when prompted to enter GRUB menu.
 - 3. Select the third option. That is *Chelsio Unified Storage vX.X.X (recovery)*.
 - 4. Now, at the BASH prompt, run the following command:

bash-3.2# mount -o remount,rw /sysroot

5. Run the following command. This will change the root directory to /sysroot.

bash-3.2# chroot /sysroot

6. Finally, change the password:

sh-3.2# /usr/bin/passwd

7. Enter the new password:

Changing password for user root New UNIX password:

8. Confirm the new password:

Retype new UNIX password:

9. Reboot for changes to take effect.

6. 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.

7. 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. That is, 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:

- 1. Access the Volume Management module
- 2. Next, expand the **Settings** section.
- 3. Select the appropriate option in the **Encryption Settings** drop-down list depending on the type of pass phrase you want to reset. Click **Apply**.

- 4. Enter the new pass phrase consisting of minimum six characters in the New Pass Phrase field.
- 5. Re-enter the same pass phrase again in the *Confirm Pass Phrase* field.
- 6. If you selected *Reset volume pass phrase* in step 3, then select the pool on which the volume was created in the **Pool** drop-down.
- 7. Select the volume for which you want to reset the pass phrase.
- 8. Depending on the type of pass phrase being reset, copy and paste the correct recovery key.
- 9. Click Change Pass Phrase.

The pass phrase is now reset.

8. 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 (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
 - 1. Click the **Start** button.

- 2. 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* (Windows XP). This will open the **Registry Editor**.
- 3. Locate the *HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk* registry key.
- 4. Locate the *TimeOutValue* entry. Right-click and select **Modify**.
- 5. Select **Decimal** and set the *Value data* field to a value greater than 60. Click **OK**.
- 6. Reboot the 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

- 1. Click the **Start** button.
- 2. Depending on the Windows version, either type *regedit.exe* in the **Search box** (Windows 7) or click **Run** and then type *regedit.exe* (Windows XP). This will open the **Registry Editor**.
- 3. Click **Computer.**

- 4. On the Edit menu click Find or use the keyboard shortcut Ctrl+F.
- 5. 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

- 6. Right-click on the entry and select **Modify**.
- 7. Select the Decimal option and set the Value data field to 180. Click OK.
- 8. Reboot machine for changes to take effect.
- Increasing maximum request hold time on Linux client
 - 1. Edit the iSCSI configuration file, *iscsid.conf*.

[root@host] # vim /etc/iscsi/iscsid.conf

2. Set the value of *node.session.timeo.replacement_timeout* to 180.

node.session.timeo.replacement_timeout = 180

- 3. Save changes and exit.
- 4. 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.

209 North Fair Oaks Avenue,

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.

Copyright © 2025. Chelsio Communications. All Rights Reserved.

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.

Copyright © 2025. Chelsio Communications. All Rights Reserved.
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.

