



Overland
Storage

SnapServer®

Administrator's Guide

For GuardianOS™ Version 7.2 on
SnapServer DX-Series Appliances



July 2012
10400317-004



©2010-12 Overland Storage, Inc. All rights reserved.

Overland[®], Overland Data[®], Overland Storage[®], ARCVault[®], DynamicRAID[®], LibraryPro[®], LoaderXpress[®], Multi-SitePAC[®], NEO[®], NEO Series[®], PowerLoader[®], Protection OS[®], REO[®], REO 4000[®], REO Series[®], Snap Appliance[®], Snap Care[®] (EU only), SnapServer[®], StorAssure[®], Ultamus[®], VR2[®], and XchangeNOW[®] are registered trademarks of Overland Storage, Inc.

GuardianOS[™], RAINcloud[™], SnapDisk[™], SnapEDR[™], Snap Enterprise Data Replicator[™], SnapExpansion[™], SnapSAN[™], SnapScale[™], SnapServer DX Series[™], SnapServer Manager[™], SnapWrite[™], and SnapServer Manager[™] are trademarks of Overland Storage, Inc.

All other brand names or trademarks are the property of their respective owners.

The names of companies and individuals used in examples are fictitious and intended to illustrate the use of the software. Any resemblance to actual companies or individuals, whether past or present, is coincidental.

PROPRIETARY NOTICE

All information contained in or disclosed by this document is considered proprietary by Overland Storage. By accepting this material the recipient agrees that this material and the information contained therein are held in confidence and in trust and will not be used, reproduced in whole or in part, nor its contents revealed to others, except to meet the purpose for which it was delivered. It is understood that no right is conveyed to reproduce or have reproduced any item herein disclosed without express permission from Overland Storage.

Overland Storage provides this manual 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. Overland Storage may make improvements or changes in the product(s) or programs described in this manual at any time. These changes will be incorporated in new editions of this publication.

Overland Storage assumes no responsibility for the accuracy, completeness, sufficiency, or usefulness of this manual, nor for any problem that might arise from the use of the information in this manual.

FW 7.2.112.

Overland Storage, Inc.
9112 Spectrum Center Blvd.
San Diego, CA 92123
U.S.A.

Tel: 1.877.654.3429 (toll-free U.S.)
Tel: +1.858.571.5555, Option 5 (International)
Fax: +1.858.571.0982 (general)
Fax: +1.858.571.3664 (sales)
www.overlandstorage.com

Audience and Purpose

This guide is intended for system and network administrators charged with installing and maintaining SnapServers DX appliances running GuardianOS 7.2 on their network. It provides information on the installation, configuration, security, and maintenance of those SnapServers and SnapExpansion units.

It is assumed that the administrator is familiar with the basic concepts and tasks of multi-platform network administration.

This guide also provides information on the following utilities and software components:

- The GuardianOS 7.2 Web Management Interface
- SnapServer Manager (SSM)
- VSS/VDS¹ Hardware Provider
- Computer Associates Antivirus (CA Antivirus)
- Third-party backup agents

GuardianOS 7.2 comes preinstalled on all new SnapServer DX-Series appliances.

Product Documentation

SnapServer product documentation and additional literature are available online, along with the latest release of the GuardianOS 7.2 software.

Point your browser to:

<http://docs.overlandstorage.com/snapserver>

Follow the appropriate link on that page to download the **latest** software file or document. For additional assistance, search at <http://support.overlandstorage.com>.

Overland Technical Support

For help configuring and using your SnapServer with GuardianOS 7.2, search for help at:

<http://support.overlandstorage.com/kb>

You can email our technical support staff at techsupport@overlandstorage.com or get additional technical support information on the [Contact Us](#) web page:

<http://docs.overlandstorage.com/support>

¹ Volume Shadow Copy Services/Virtual Disk Services.

For a complete list of support times depending on the type of coverage, visit our website at:
<http://docs.overlandstorage.com/care>

Conventions

This document exercises several alerts and typographical conventions.

Alerts

Convention	Description & Usage
 IMPORTANT	An <i>Important</i> note is a type of note that provides information essential to the completion of a task or that can impact the product and its function.
 CAUTION	A <i>Caution</i> contains information that the user needs to know to avoid damaging or permanently deleting data or causing physical damage to the hardware or system.
 WARNING	A <i>Warning</i> contains information concerning personal safety. Failure to follow directions in the warning could result in bodily harm or death.
AVERTISSEMENT	Un Canadien avertissement comme celui-ci contient des informations relatives à la sécurité personnelle. Ignorer les instructions dans l'avertissement peut entraîner des lésions corporelles ou la mort.

Typographical Conventions

Convention	Description & Usage
Button_name	Words in this special boldface font indicate the names of command buttons found in the Web Management Interface.
Ctrl-Alt-r	This type of format details the keys you press simultaneously. In this example, hold down the Ctrl and Alt keys and press the r key.
NOTE	A Note indicates neutral or positive information that emphasizes or supplements important points of the main text. A note supplies information that may apply only in special cases, for example, memory limitations or details that apply to specific program versions.
Menu Flow Indicator (>)	Words with a greater than sign between them indicate the flow of actions to accomplish a task. For example, Setup > Passwords > User indicates that you should press the Setup button, then the Passwords button, and finally the User button to accomplish a task.
<i>Courier Italic</i>	A variable for which you must substitute a value
Courier Bold	Commands you enter in a command-line interface (CLI)

Information contained in this guide has been reviewed for accuracy, but not for product warranty because of the various environments, operating systems, or settings involved. Information and specifications may change without notice.

Software Updates

The latest release of the GuardianOS software can be obtained from the Downloads and Resources (NAS Solutions) page at the Overland Storage website:

<http://docs.overlandstorage.com/snapserver>

Follow the appropriate instructions to download the **latest** software file.

For additional assistance, search at <http://support.overlandstorage.com/>

Finding More Information

Product documentation related to GuardianOS SnapServers and expansion arrays are listed below. The current versions of all these documents are always available from the Overland Storage NAS Solutions website (<http://docs.overlandstorage.com/snapserver>).

Source	Location	Content
<i>Quick Start Guide</i>	Product Packaging and Web	Provides complete instructions for installing the server into a rack and connecting the server to the network. Also contains links to warranty registration and information.
<i>Quick Start Guide Translations</i>	Web	Quick Start Guide is translated into French, Italian, German, Spanish, and Russian.
<i>EULA</i>	Product Packaging and Web	End User License Agreement for GuardianOS.
<i>Administrator's Guide</i>	Web	Provides an overview of the configuration, maintenance, and troubleshooting of SnapServers, the administration of the CA Antivirus software, the installation of third-party backup agents, and detailed instructions on using the Web Management Interface.
<i>SnapServer Online Help</i>	Web Management Interface via the Overland Website	Basic troubleshooting information embedded in the software with direct links to additional information from the Administrator's Guide.

Electrostatic Discharge Information

A discharge of static electricity can damage static-sensitive devices. Proper packaging and grounding techniques are necessary precautions to prevent damage. To prevent electrostatic damage, observe the following precautions:

- Transport products in static-safe containers such as conductive tubes, bags, or boxes.
- Cover the appliance with approved static-dissipating material.
- Use a wrist strap connected to the work surface and properly-grounded tools and equipment.
- Keep the work area free of non-conductive materials such as foam packing materials.
- Make sure you are always properly grounded when touching a static-sensitive component or assembly.
- Avoid touching pins, leads, or circuitry.

Preface

Chapter 1 - Overview

GuardianOS Specifications	1-1
What's New in GuardianOS 7.2	1-3
What's New in GuardianOS 7.1	1-4
What's New in GuardianOS 7.0	1-4
Using the SnapServer Manager	1-5
SnapServer Manager Installation	1-6
Launch SnapServer Manager	1-6
SnapServer Manager Help	1-6
Multiserver Administration	1-6
SSM Feature Licensing	1-7
Connecting to the Server for the First Time	1-7
Connect Using the Server Name	1-8
Connect Using SSM	1-8
SnapExtensions	1-8
Wake-on-LAN Support	1-9

Chapter 2 - Server Setup and Options

Initial Setup Wizard	2-1
General Configuration	2-3
TCP/IP Configuration	2-4
RAID Type Selection	2-4
DynamicRAID Setup	2-6
Traditional RAID Setup	2-8
Registration Page	2-10
Server Status and Site Map	2-11
Hardware Information	2-12
Scheduling Data Protection Tasks	2-12
Server Options	2-13
Server Name	2-14
Date/Time	2-15
Secure Shell	2-16
UPS Protection	2-18
Print Server	2-20

Chapter 3 - Network Access

View Network Information	3-2
TCP/IP Options	3-3
Configuring TCP/IP Settings	3-5

Issues in TCP/IP Configuration	3-6
Windows Networking (SMB)	3-8
Support for Windows Networking (SMB)	3-8
Support for Windows Network Authentication	3-9
Connect from a Windows Client	3-10
Connect a Mac OS X Client Using SMB	3-10
Configure Windows/SMB Networking	3-10
Apple Networking (AFP)	3-14
AFP Configuration Guidelines	3-15
AFP Procedures	3-15
NFS Access	3-16
Support for NFS	3-17
NFS Procedures	3-17
NIS Domain	3-20
Guidelines for Configuring NIS	3-20
FTP/FTPS Access	3-21
Supported FTP Clients	3-21
SNMP Configuration	3-22
Default Traps	3-23
Supported Network Manager Applications and MIBs	3-23
Configure SNMP	3-23
Web Access	3-24
Configuring HTTP/HTTPS	3-25
Using Web Root to Configure the SnapServer as a Simple Web Server	3-25
Web View	3-27
iSNS Configuration	3-28

Chapter 4 - DynamicRAID Storage

Storage Pools	4-2
Storage Pool Creation	4-2
Storage Pool Properties	4-5
View Disks from Storage Pool Properties Page	4-7
Storage Pool Deletion	4-7
Parity Management	4-8
Volumes	4-10
Volume Creation	4-10
Volume Properties	4-11
Volume Deletion	4-12

Chapter 5 - Traditional RAID Storage

Storage Guides	5-2
Factors in Choosing a RAID Type	5-2
Local and Global Spares	5-3
RAID Sets	5-4
Create RAID Sets	5-4
Group RAID Sets	5-7
Change RAID Settings	5-11
Manage Global Spares	5-12
RAID Set Properties	5-12
Volumes	5-14
Volumes and the Snapshot Pool	5-14

Volume Creation	5-15
Volume Properties	5-17
Quotas	5-19
Quotas Page	5-20
Enable/Disable Quotas	5-21
Adding Quotas Wizard	5-21
Displaying/Changing Quotas	5-23

Chapter 6 - Other Storage Options

Snapshots	6-2
Creating Snapshots	6-3
Schedule Snapshots	6-5
Snapshot Space	6-6
Snapshot Properties	6-7
iSCSI Disks	6-9
Configuring iSCSI Initiators	6-10
iSCSI Configuration on the SnapServer	6-17
Create iSCSI Disks	6-19
Edit an iSCSI Disk	6-21
Delete an iSCSI Disk	6-22
Configuring VSS/VDS for iSCSI Disks	6-22
Disks	6-25
Replacing Disk Drives	6-26
Adding Disk Drives	6-28
Expansion Units	6-32
Managing Expansion Unit Storage	6-32
Integrating Orphaned Expansion Units	6-34

Chapter 7 - Security Options

Overview	7-1
Guidelines for Local Authentication	7-2
User and Group ID Assignments	7-3
Security Guides	7-3
Windows Active Directory Security Guide	7-4
Entire Volume Security Guide	7-5
Folder on Volume Security Guide	7-5
Shares	7-6
Share Security Overview	7-6
Create Shares	7-7
Edit Share Properties	7-9
Delete Shares	7-10
Configuring Share Access	7-10
Local Users	7-17
Create a User	7-17
Edit User Properties	7-18
User Password Policies	7-19
Assign User to Group	7-21
Delete Local User	7-21
Local Groups	7-22
Create New Group	7-22
Edit Group Properties	7-23

Specify Users in Group	7-24
Delete Group	7-24
Security Models	7-25
Managing Volume Security Models	7-25
Managing Folder Security Models	7-25
ID Mapping	7-26
Add Mapping	7-27
Change Mapping	7-29
Auto Mapping	7-32
Remove Mappings	7-32
Remove Missing ID Mappings	7-34
Filesystem Updates	7-35
Home Directories	7-36
Configure Home Directories	7-37

Chapter 8 - System Monitoring

System Status	8-2
SnapServer Status	8-2
Expansion Unit Status	8-3
Active Users	8-3
Open Files	8-3
Event Log	8-4
Filter the Log	8-4
Erase All Log Entries	8-4
Tape Monitor	8-5

Chapter 9 - Maintenance

Shutdown and Restart	9-2
Manually Powering SnapServers On and Off	9-2
Factory Defaults	9-3
Disaster Recovery	9-4
Backing Up Server and Volume Settings	9-5
The SnapDRImage File and the Volume Files	9-5
Restoring Original Server and Volume Configurations	9-6
Rejoining the Server to a Windows Domain	9-7
SnapDRImage Usage Scenario	9-7
Replacing a Server	9-8
Cloning a Server	9-8
Data Import	9-9
Setting Up a Data Import Job	9-10
Stopping an Import Job	9-12
Recreating an Import Job	9-12
Preserving Permissions	9-12
OS Updates	9-13
Update the GuardianOS Software	9-13
Software Update Notification	9-13
Configuring Update Notification	9-14
Checking for Updates	9-14
Support	9-14
Registering Your Server	9-14
Maintenance Tools	9-15

Email Notification	9-15
Host File Editor	9-17
Checking Filesystems	9-18

Chapter 10 - Misc. Options

Home Pages – Web/Admin	10-1
Web Home	10-2
Admin Home	10-4
SnapExtensions	10-5
Snap Finder	10-6
Snap Finder Properties	10-8
Change Password	10-9
Changing Your Password	10-9
Mgmt. Interface Settings	10-10

Chapter 11 - CA Antivirus Software

Antivirus Dependencies	11-1
Launching the CA Antivirus GUI	11-2
Launching the CA Antivirus Browser Interface	11-2
The Local Scanner View	11-2
Scan Jobs	11-3
Defining Scan Jobs	11-3
Running a Manual Scan Job	11-4
Scheduling a Scan Job	11-4
Signature Updates	11-5
Updating SnapServers with Internet Access	11-5
Updating a SnapServer without Internet Access	11-6
Distributing Updates from One SnapServer to Another	11-6
Verifying Download Events	11-7
Alert Options	11-8
The Move Directory	11-8
Log View	11-9

Appendix A - DynamicRAID Overview

About DynamicRAID	A-1
Should I use DynamicRAID or Traditional RAID?	A-2
Setting Up DynamicRAID	A-4
Implementation	A-4
Storage Expansion	A-5
Snapshots	A-5
iSCSI Target Volumes	A-5
Indicators	A-5
Additional Information on DynamicRAID Sizing	A-5

Appendix B - Backup Solutions

Backup and Replication Solutions Table	B-1
Integrated Backup Solutions	B-2
Snap Enterprise Data Replicator (Snap EDR)	B-2
Off-the-Shelf Backup Solutions	B-3
Preparing to Install a Backup Solution	B-3
Preinstallation Tasks	B-4

Installing the CA ARCserve Agent	B-5
Installing the Symantec Backup Exec RALUS Agent	B-6
Installing the Symantec NetBackup 6.5 Client	B-8
Installing the EMC NetWorker Client	B-10
iSCSI Disk Backups	B-14
Using Backup Exec for VSS-based Snapshots of SnapServer iSCSI Disks	B-14

Appendix C - Security and Access

Security Model Rules	C-1
Security Model Directories	C-2
Security Model Management	C-2
Special Share Options	C-3
Hiding Shares	C-3
Share Level Permissions	C-3
Where to Place Shares	C-4
File and Share Access	C-4
NFS Share Access	C-4
Snapshot Access	C-4
Snapshot Shares and On Demand File Recovery	C-4
Creating a Snapshot Share	C-5
File-level Security	C-5
Security Personalities and Security Models	C-6
Windows ACLs	C-6

Appendix D - Troubleshooting SnapServers

LED Indicators	D-1
System Reset Options	D-3
Performing System Resets Without Network Access	D-3
Maintenance Mode	D-4
Networking Issues	D-4
Miscellaneous Issues	D-7
Phone Home Support	D-8

Appendix E - Command Line Interface

SnapCLI Syntax	E-1
SnapCLI Procedures	E-3
SnapCLI Commands	E-3
Scripts in SnapCLI	E-9
Running a SnapCLI Script	E-9
Sample Script	E-10

Appendix F - GuardianOS Ports

Master Glossary & Acronym List

Index

SnapServer appliances are designed as flexible, low-maintenance network-attached storage (NAS) file servers optimized for performance and efficiency. They run GuardianOS, an operating system built to maximize file I/O throughput across multi-network protocols. To this end, all unnecessary system control and processing functions that are associated with a general-purpose server have been removed. This guide applies to SnapServer DX appliances running GuardianOS version 7.2 and includes SnapExpansion units.

Topics in Overview:

- [GuardianOS Specifications](#)
- [What's New in GuardianOS 7.2](#)
- [What's New in GuardianOS 7.1](#)
- [What's New in GuardianOS 7.0](#)
- [Using the SnapServer Manager](#)
- [Connecting to the Server for the First Time](#)
- [SnapExtensions](#)
- [Wake-on-LAN Support](#)

GuardianOS Specifications

These specifications apply to all SnapServers running GuardianOS 7.2.

Feature	Specification
Network Transport Protocols	TCP/IP (Transmission Control Protocol/Internet Protocol) UDP/IP (User Datagram Protocol/Internet Protocol)
Network Block Protocols	iSCSI (Internet Small Computer System Interface)
Network File Protocols	Microsoft Networking (CIFS/SMB) UNIX Network Filesystem (NFS) 2.0/3.0/4.0 Apple Filing Protocol (AFP) v2.0/v3.1/3.2* Hypertext Transfer Protocol (HTTP/HTTPS) File Transport Protocol (FTP/FTPS) * AFP v 3.2 ACLs and extended attributes not supported.

Feature	Specification
Network Client Types	Microsoft Windows 2003/2003 R2/2008 SP2/2008 R2 /XP SP3/Vista SP2/7 Mac OS X 10.5/10.6/10.7 Sun Solaris 10 and 11 HP-UX 11 AIX 5.3/6 Red Hat Enterprise Linux (RHEL) 4.x/5.x/6.x Novell SuSE Linux Enterprise Server (SLES) 10.x/11.x
Network Security	<ul style="list-style-type: none">• CA Antivirus software• Microsoft Active Directory Service (ADS) (member server)• UNIX Network Information Service (NIS)• File and Folder Access Control List (ACL) Security for Users and Groups• Secure Sockets Layer (SSL v2/3) 128-bit Encryption• Target Challenge Handshake Authentication Protocol (CHAP) for iSCSI• SMTP Authentication and support for email encryption (STARTTLS and TLS/SSL encryption protocols)
Data Protection	<ul style="list-style-type: none">• Snapshots for immediate or scheduled point-in-time images of the filesystem• Support for local backup with Symantec NetBackup/Backup Exec Remote Media Server for Linux• Support for network backup with Symantec NetBackup/Backup Exec, CA ARCserve, or EMC NetWorker• APC[®] brand Uninterruptible Power Supply (UPS) with Network Management Cards, a USB interface, or a serial interface (with USB-to-Serial adapter) are supported for graceful system shutdown

Feature	Specification
System Management	<ul style="list-style-type: none"> • Browser-based administration tool called the Web Management Interface • SnapCLI for volume system deployment • SnapServer Manager utility (platform independent) • SNMP (MIB II and Host Resource MIB) • User disk quotas for Windows, UNIX/Linux, Mac, FTP/FTPS (Traditional RAID only) • Group disk quotas for UNIX/Linux (Traditional RAID only) • Environmental monitoring • Email event notification and SNMP trap notification • Data importation (migration)
RAID Options with Traditional RAID	<ul style="list-style-type: none"> • RAID 0 (drive striping): Large virtual drive with data striped across all drives of the array to provide maximum performance with no loss in usable capacity. Does not provide data protection. • RAID 1 (drive mirroring): One or more drives duplicate one drive for maximum data protection. Available only on systems with two (2) or more drives. • RAID 5 (drive striping with parity): For each array, the size of one drive is reserved for parity. Provides good performance and space utilization with one-drive fault tolerance. Available only on systems with four (4) or more drives. • RAID 6 (drive striping with two parity drives): Like a RAID 5 except that two drives are used for parity rather than one. Provides moderate performance and reasonable space utilization with two-drive fault tolerance. Available only on systems with four (4) or more drives. • RAID 10 (striped mirroring): A combination of RAID 0 and RAID 1. Provides high performance and fault tolerance. Available only on systems with four (4) or more drives. • Global or local spare support • Instant Capacity Expansion (ICE): Logically groups RAIDs for dynamic online scalability.
DHCP Support	Supports Dynamic Host Configuration Protocol (DHCP) for automatic assignment of IP addresses

What's New in GuardianOS 7.2

NOTE: For details and descriptions of all the new features and a list of other improvements to the firmware, see the [Product Information Bulletin](#) and the [Release Notes](#) on the Overland SnapServer website.

With the release of the latest version of GuardianOS 7.2, the following features and functionality are now available:

Feature	New Functionality
SnapExpansion Units	The capacity of a SnapServer can be increased by attaching expansion units. SnapExpansion units are supported by both DynamicRAID and Traditional RAID. GuardianOS now supports seven (7) SnapExpansion units connected to a SnapServer DX2 or three (3) units connected to a SnapServer DX1. For more information, see "Expansion Units" .
SSD Support	SnapServer servers can now support the use of Solid State Drives (SSD).
Support of Large Domains	The GuardianOS has been enhanced to support domains of more than 60,000 users and groups. The SnapServer administrator can now specifically search for domain users and groups via the Web Management Interface. This replaces the previous Domain Filtering feature. This impacts three areas: Share Access, ID Mapping, and, if using Traditional RAID, Quotas.

What's New in GuardianOS 7.1

NOTE: For details and descriptions of all the new features, see the [Product Information Bulletin](#) and the [Release Notes on the Overland SnapServer website](#).

GuardianOS 7.1 introduces the following features and functionality:

Feature	New Functionality
Domain Filtering	Domain filtering is a new feature (in both the Web Management Interface and CLI), and allows the user to limit the number of domain users and groups that the server imports from the Active Directory domain. This is done by restricting the import to a specific organizational unit, as well as to specific domain users and groups. NOTE: Domain filtering is only available after the server is joined to an Active Directory domain. Domain filtering cannot be used when "Trusted Domains" is enabled.
Display of List Boxes and Tables Improved	The fields are now fixed length and reworked to better display the information. The VSS/VDS Access Control and Snap Finder Properties pages have been changed to match the new Domain Filtering page.

What's New in GuardianOS 7.0

NOTE: For details and descriptions of all the new features, see the [Product Information Bulletin](#) and the [Release Notes on the Overland SnapServer website](#).

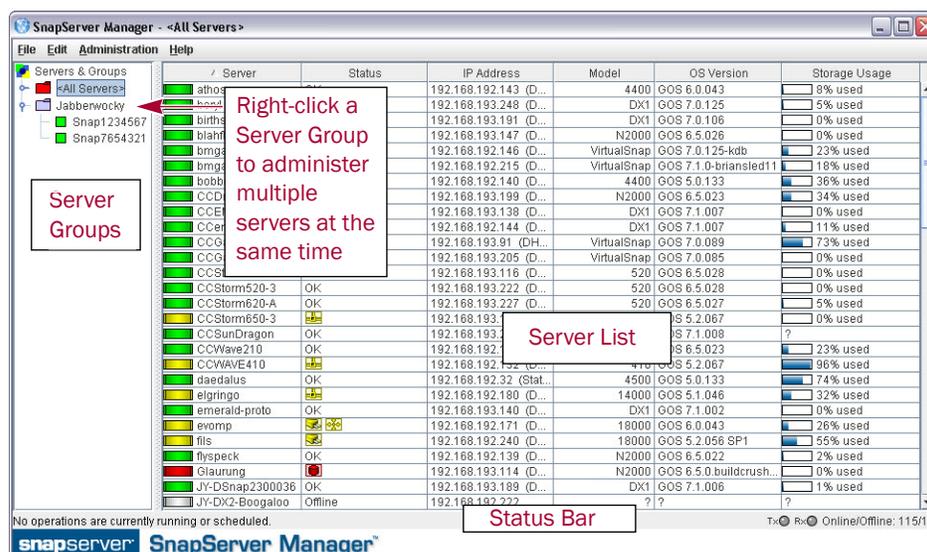
GuardianOS 7.0 introduces the following features and functionality:

Feature	New Functionality
DynamicRAID and Traditional RAID	<p>DynamicRAID simplifies storage management and provides additional configuration options not available in Traditional RAID. A SnapServer can be purchased with any amount of initial storage (or number of drives), and more capacity can be added over time by inserting or replacing drives and incorporating the new drive(s) via the Web Management Interface (see “Storage Pool Creation” for DynamicRAID and “Add a Disk Drive as Local Spare to RAID”). Volumes can be added and removed at will, and all volumes share the same underlying pool of storage. The appliance can also be run in Traditional RAID mode to allow more direct control over storage configuration.</p> <p>Traditional RAID mode allows for the manual management of the RAID sets. RAID sets 0, 1, 5, 6, or 10 can be manually created and configured, snapshots are limited by volume settings, and RAIDS must be manually deleted and recreated to change them.</p> <p>See also Appendix A, “DynamicRAID Overview,” for more information comparing DynamicRAID to Traditional RAID.</p>
64-bit Operating System	This supports volumes over 16TB, RAM over 4GB, and individual disk drives over 2TB.
iSCSI Write Cache Modification	The iSCSI Disk Properties page of the Web Management Interface features an Enable Write Cache checkbox to allow modification of the write cache setting of an existing iSCSI disk.

Using the SnapServer Manager

SnapServer Manager (SSM) is a Java-based, platform-independent, multiserver administrative application that runs on all major platforms. SSM provides a single interface from which administrators can discover, configure, and monitor all GuardianOS SnapServers on their network. With SSM, administrators can compare, copy, and configure

settings for groups of GuardianOS SnapServers in a single operation.



SnapServer Manager Installation

You can download and install SSM by navigating to the Overland Storage NAS website and downloading the [SnapServer Manager executable file](#). SSM can be installed to all client platforms, including Windows, Mac OS X, Linux, and UNIX.

Launch SnapServer Manager

Launch SSM using one of the methods described in the following table:

Operating System	Procedure
Microsoft Windows XP/2003/Vista/2008/7	Click Start . Point to Programs (or All Programs)> SnapServer Manager , then select SnapServer Manager .
Mac OS v10.5 or higher	Open the SnapServer Manager folder and double-click the SnapServer Manager icon.
UNIX/Linux	For default options, cd to home directory, then run the SnapServer Manager command: <pre>./Snap_Server_Manager</pre> If you selected not to create links, cd to home directory, then cd to the SnapServer Manager directory, and run the SnapServer Manager command: <pre>./Snap_Server_Manager</pre>

SnapServer Manager Help

For detailed information on SSM, select **Help > Help Topics** from within the application.

Multiserver Administration

Multiserver administration provides the following:

- **Simultaneous application of settings to server groups** – You can organize GuardianOS servers into functional groups and apply settings to all servers in the group simultaneously.
- **Comparing settings across servers** – SSM can compare settings across any number of GuardianOS servers and identify when settings differ among servers. For example, comparing protocol access configuration for a group of servers may reveal that settings are consistent for Windows, NFS, and AFP but that differences exist among servers in HTTP/HTTPS and FTP/FTPS settings.
- **Copying settings from one server to one or more different servers** – SSM can copy selected settings (TCP/IP, SNMP, SMB, etc.) from any GuardianOS server to one or more different GuardianOS servers.
- **Scheduling operations to run during off-peak hours** – Operations can be scheduled to run on multiple GuardianOS servers during off-peak hours.
- **Automatic email notification of completed operations** – You can configure SSM to send an operations report (CSV format) upon completion of any operation.
- **Automatic notification of available GuardianOS updates** – SSM is by default configured to check daily for applicable updates to the servers it has discovered and display an alert, notifying the administrator of the available updates.

SSM Feature Licensing

Use the SSM Feature Licensing menu to apply SnapExtension license keys to one or more servers. There is no limit to the number of licenses that can be entered using this dialog box.

1. Start SSM and select the GuardianOS servers to be licensed.
2. Navigate to **Administration > Feature Licensing**. If you have not already obtained your licenses, in the **License Required** dialog box, select **Click here to purchase SnapExtension license keys at www.SnapServer.com**.
3. Once you have obtained the license keys, enter one license key per line (or multiple keys per line, separated by spaces), click **Enter License**, then click **OK**.

The Feature License dialog box does not display any pre-existing SnapExtension licenses. Only licenses that have been applied while the current dialog box is open will be displayed.

Connecting to the Server for the First Time

SnapServers are configured to acquire an IP address from a DHCP server. If no DHCP server is found on the network, the SnapServer defaults to an IP address in the range of 169.254.xxx.xxx and is labeled “ZeroConf” in SSM. While you may not be able to see the server on your network, you can discover the SnapServer using either the default server name or the SSM utility. Use the server name method if you are installing one SnapServer on the network. Use SSM if you are installing two or more SnapServers, or if your network does not have IP-to-name resolution services.

Connect Using the Server Name

This procedure requires that name resolution services (via WINS or an equivalent service) be operational.

1. Find the **server name**.
The default server name is “SNAP n ”, where n is the server number. For example, the name of a SnapServer with a server number of 6100191 is SNAP6100191. The server number is a unique, numeric-only string that appears on a label affixed to the top of the server in the left front corner.
2. In a Web browser, enter the **server URL**.
For example, enter `http://SNAP n` (where n is the server name).
3. Press **Enter** to open the Web View page.
4. Log into the Web Management Interface.
In the login dialog box, enter **admin** as the user name and **admin** as the password, then click **OK**.
5. Complete the **Initial Setup Wizard**.

Connect Using SSM

1. Launch **SSM**.
SSM discovers all SnapServers on its local network segment and displays their server names, IP addresses, and other status information in the main console. If you do not have a DHCP server, there might be a delay before the server appears on the network.

NOTE: To distinguish multiple SnapServers, you may need to find their default server names as explained in the previous procedure.
2. If using a DHCP server, proceed to [Step 3](#); otherwise, assign an **IP address** to the new server.
 - a. In SSM, right-click the **server name**.
 - b. Select **Set IP Address**.
 - c. Enter an IP address and a subnet mask, then click **OK**.
3. In SSM, right-click the server name and select **Launch Web Administration**.
4. Log into the Web Management Interface.
In the login dialog box, enter **admin** as the user name and **admin** as the password, then click **OK**.
5. Complete the **Initial Setup Wizard**.

At this point, your SnapServer is ready to be configured for your specific environment and needs.

SnapExtensions

SnapExtensions are software applications, agents, and utilities that extend the capabilities of a SnapServer (for additional details, see “[SnapExtensions](#)”). Some SnapExtensions are fully functional out-of-the-box; others may require a download and/or the purchase of a license for full operation. For up-to-date information on feature availability, contact [Overland Storage](#).

Wake-on-LAN Support

Wake-on-LAN, the Ethernet computer networking standard that allows a powered-off computer to be powered on by a network signal, is automatically enabled (and cannot be disabled) for Ethernet 1 (Management) and Ethernet 2 ports. Wake-on-LAN is activated when another computer on the same LAN sends a “magic packet” to the SnapServer using the SnapServer Manager or other program designed to send magic packets. Wake-on-LAN only works for SnapServers and does not work with any expansion units that may be attached to the system.

Server Setup and Options

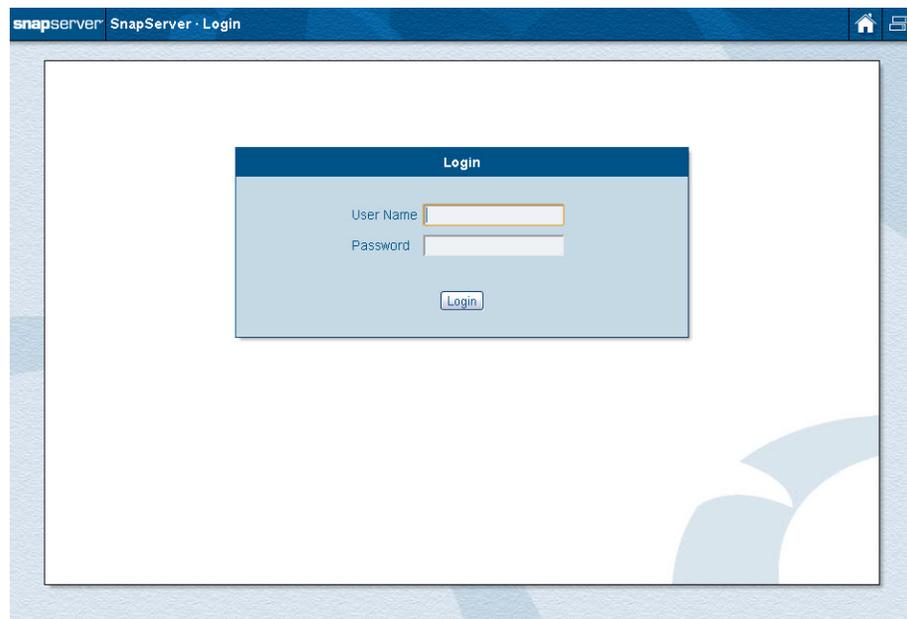
This section covers the initial steps needed to set up and configure a SnapServer running GuardianOS 7.2.

Topics in Server Setup and Options:

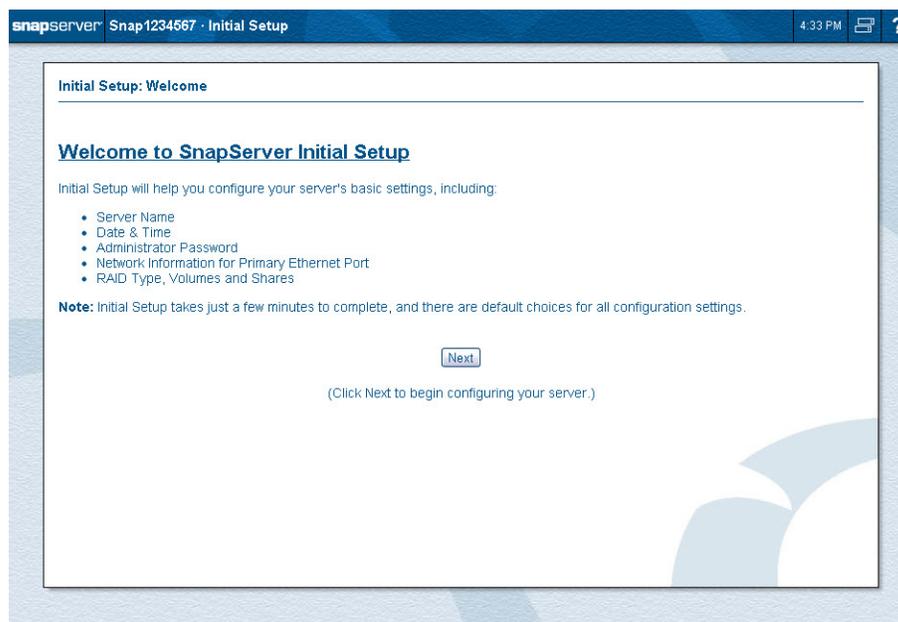
- [Initial Setup Wizard](#)
- [Server Status and Site Map](#)
- [Scheduling Data Protection Tasks](#)
- [Server Options](#)

Initial Setup Wizard

The first time you connect to a SnapServer via the Web Management Interface, you are prompted to log in. Log in using the default administrator user name **admin** and password **admin**.



Once you log in, the Initial Setup Wizard runs displaying the Welcome page.



The Initial Setup Wizard consists of several web pages that help you configure your server's basic settings, including:

- Server name
- Date and time
- Administrator password
- Network information for the Primary Ethernet port (default = Ethernet 1)
- RAID type, volumes, and shares

The basic pages include:

- General Configuration
- TCP/IP Configuration
- RAID Type Selection
 - DynamicRAID Setup
 - Detecting Available Disks
 - Configure Storage – Head Unit
 - Traditional RAID Setup
 - Detecting Available Disks
 - Configure Storage – Head Unit
- Registration Page

Click **Next** to start the Initial Setup wizard.

General Configuration

Clicking Next on the Welcome page displays the General Information page of the Initial Setup wizard. This page allows you to change the basic information for the SnapServer. It is recommended that you set your Administrator password at this time to something other than the default.

The screenshot shows the 'Initial Setup: General Information' page. At the top, it says 'SnapServer Snap1234567 · Initial Setup' and '4:34 PM'. The main content area is titled 'Initial Setup: General Information' and contains the following fields and options:

- Server Name:** Snap1234567
- Date and Time:**
 - Date:** 2011 - 12 - 13 (YYYY-MM-DD)
 - Time:** 04 : 34 : 07 PM
 - Time Zone:** (UTC-08:00) Pacific Time (US & Canada)
- Note:** Changing the date, time or time zone may require you to re-login to this Web Management Interface.
- Administrator Password:** Two input fields for 'Admin Password' and 'Confirm Password'. A note says '(Leave blank to keep existing administrator password)'.
- Warning:** You should assign an administrator password if you want to protect your server from unauthorized changes.
- Next** button.

1. Enter (or accept) the **Server Name**.

The default server name is `SNAPnnnnnnnn`, where `nnnnnnnn` is the server number. If desired, a unique server name of up to 15 alphanumeric characters can be used. In addition to letters and numbers, you can also use a dash (-) between characters, but spaces are not allowed.

2. Enter (or accept) the **Date/Time Settings**.

The SnapServer time stamp applies when recording server activity in the event log (Monitor Menu), setting the create/modify time on a file, and when scheduling snapshot, antivirus, or Snap Enterprise Data Replicator (EDR) operations. Edit the settings according to local conditions.

NOTE: GuardianOS automatically adjusts for Daylight Saving Time, based on the selected time zone.

3. Change the **Administrator Password**.

The default administrator user name is `admin`, and the default password is also `admin`. To prevent unauthorized access to the SnapServer, enter a new secure password immediately in the fields provided.

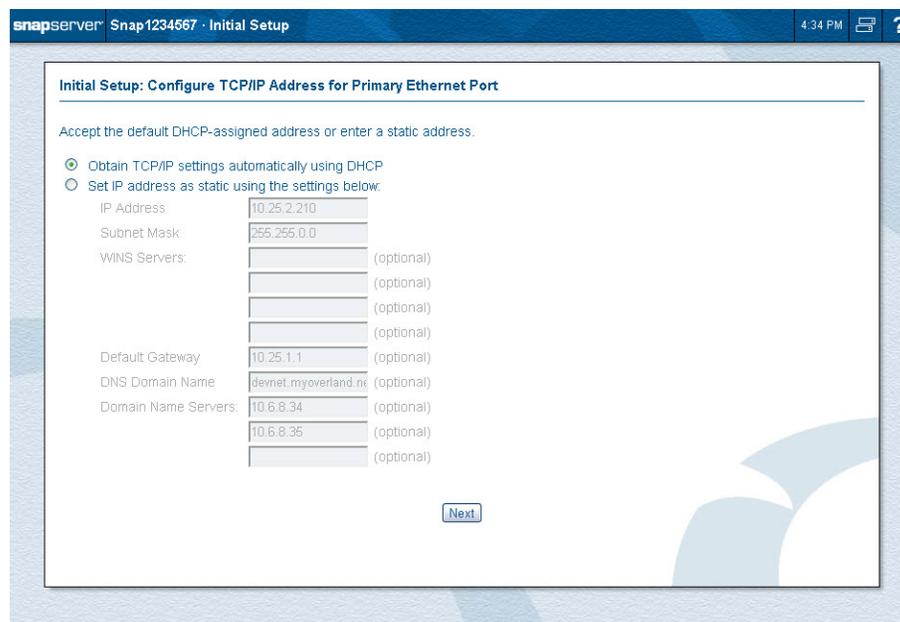
NOTE: Passwords consist of 1 to 15 alphanumeric characters and are case-sensitive.

4. To continue to the next page, click **Next**.

If you have changed the date, time, or time zone settings in the [General Configuration](#) window above, you may be prompted to log in again before continuing the setup.

TCP/IP Configuration

The next wizard page shows the current TCP/IP information for this SnapServer. All SnapServers come preset to acquire an IP address from a DHCP server.



The screenshot shows the 'Initial Setup: Configure TCP/IP Address for Primary Ethernet Port' window. It has a title bar with 'snapserver Snap1234567 · Initial Setup' and a system tray with '4:34 PM' and a help icon. The main content area has the title 'Initial Setup: Configure TCP/IP Address for Primary Ethernet Port' and the instruction 'Accept the default DHCP-assigned address or enter a static address.' There are two radio buttons: 'Obtain TCP/IP settings automatically using DHCP' (selected) and 'Set IP address as static using the settings below:'. Below the second option are input fields for: IP Address (10.25.2.210), Subnet Mask (255.255.0.0), WINS Servers (three optional fields), Default Gateway (10.25.1.1), DNS Domain Name (devnet.myoverland.net), and Domain Name Servers (two optional fields: 10.6.8.34 and 10.6.8.35). A 'Next' button is at the bottom right.

1. If you wish to assign a **static IP** instead, check the box for obtaining a static address and enter the following information:
 - The IP address for the SnapServer (required)
 - The subnet mask (required)
 - Any WINS server IP addresses
 - The default gateway IP address
 - The DNS domain name and IP addresses
2. Click **Next** to configure the type of RAID storage you want to use.

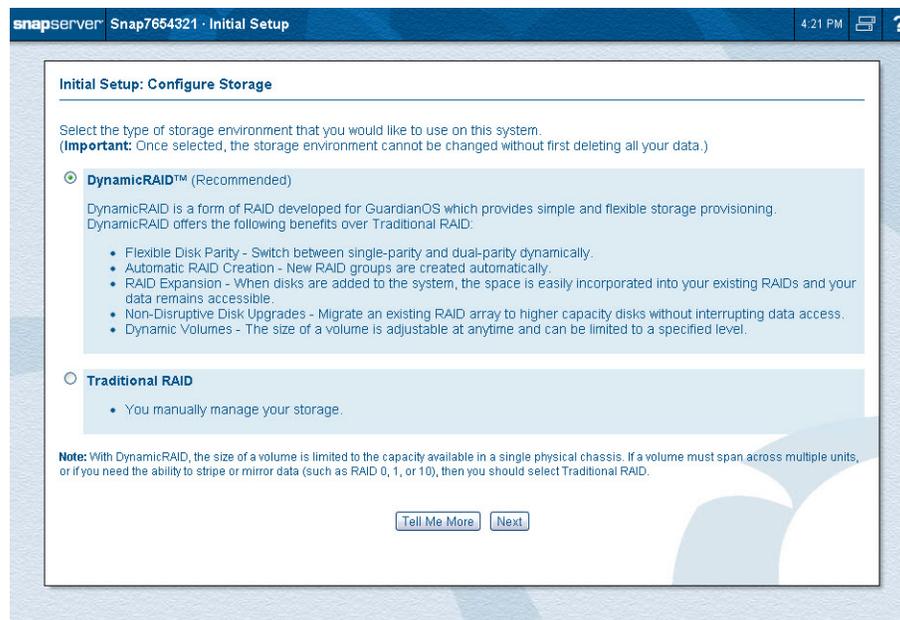
NOTE: If a Static IP address was entered, the network is restarted automatically (without confirmation from the user) when Next is clicked.

RAID Type Selection

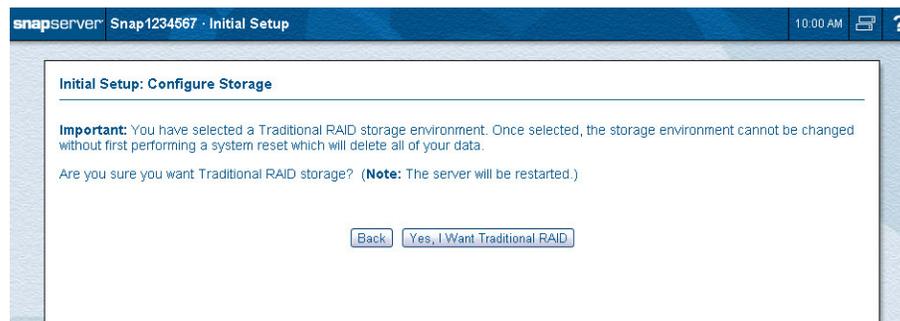
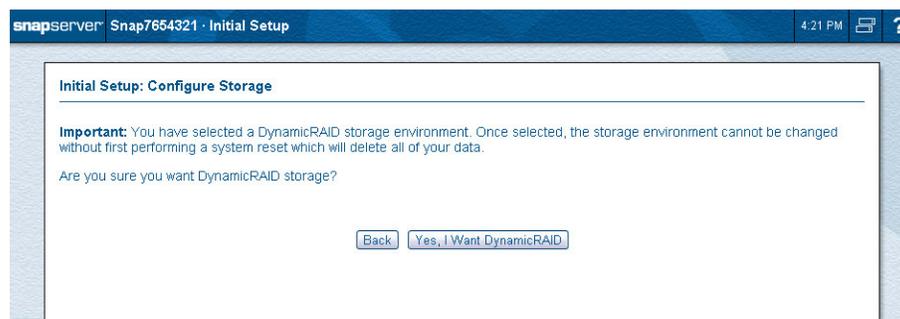
GuardianOS 7.2 offers the new, powerful DynamicRAID feature that simplifies management of disk additions and replacements in a RAID environment. You can also manually manage the RAIDs using the Traditional RAID option.

To determine which RAID configuration is appropriate for your needs, see [“Should I use DynamicRAID or Traditional RAID?”](#).

1. Click the button to choose either DynamicRAID or Traditional RAID.



2. After you have made your selection, click **Next**. You will be prompted to confirm your selection of either DynamicRAID or Traditional RAID:



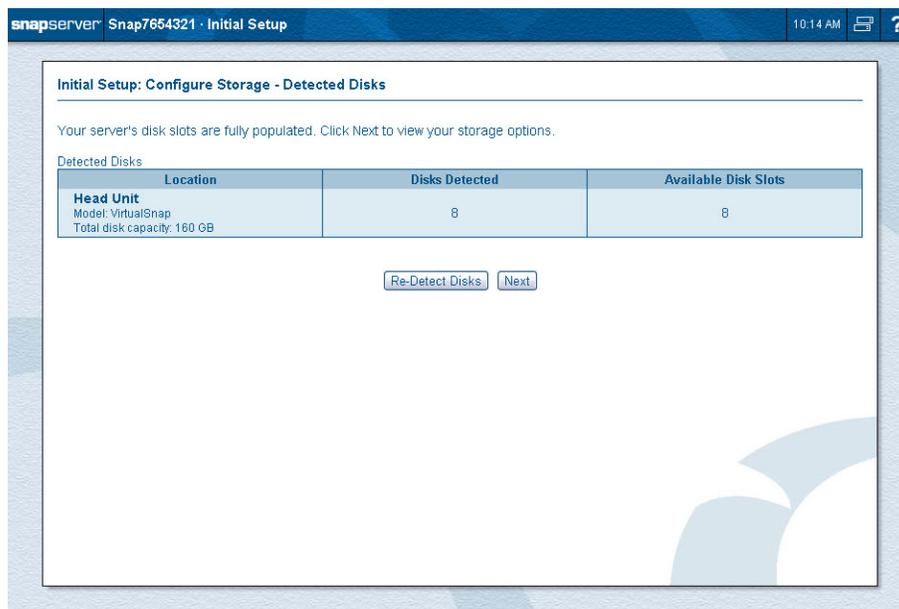
3. Click the **Yes** option.
4. Continue configuring the SnapServer based on the **RAID mode** you selected:
- If you selected DynamicRAID, proceed to [DynamicRAID Setup](#).
 - If you selected Traditional RAID, the server will be restarted. Once the server has restarted and you have logged back in, proceed to [Traditional RAID Setup](#).

DynamicRAID Setup

Once DynamicRAID is selected, the wizard continues with the setup of the head unit. The first step is to detect the installed disk drives.

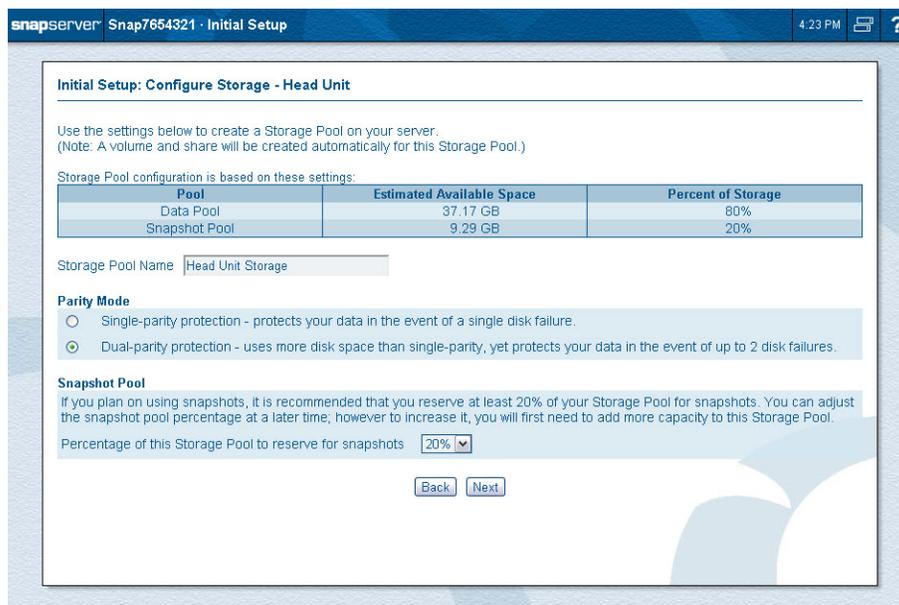
Detecting Available Disks

Available disks are detected and displayed on the Detected Disks setup page. You will have an opportunity to add more disks if empty slots exist.



Configure Storage – Head Unit

The next page lets you configure your head unit storage by choosing the parity mode and snapshot pool size.



You choose the parity mode to set the data pool size. Options presented are based on the number and available space of disk drives detected in the previous step. See [“Parity Management”](#) for full details.

In addition, snapshot space can be reserved using the following guidelines:

- For typical usage, at least 20% snapshot space should be reserved from each storage pool.
- Once snapshot space is set up under DynamicRAID, it can be decreased at any time. However, to increase the size of the snapshot space, either the storage pool must be deleted and re-created, or you must add more storage capacity to your storage pool. See [“Snapshots”](#) for more information.

To proceed with DynamicRAID storage configuration:

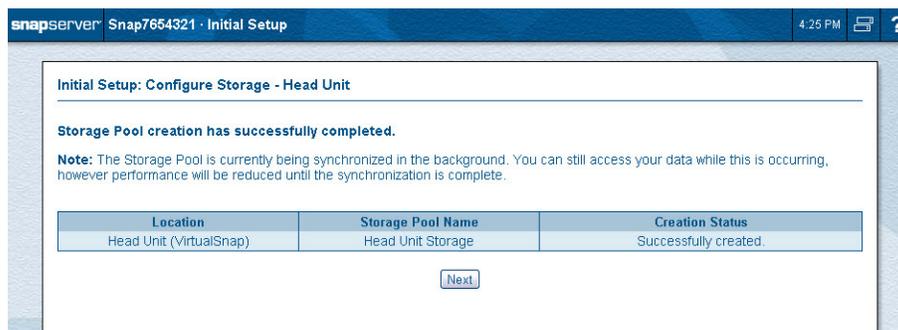
1. Select the **parity mode** from the options provided.
2. Use the drop-down list to choose the size of the **snapshot pool**.
3. After you have made your selections, click **Next**.
4. At the confirmation page, click **Create Storage Pool**.

NOTE: If a disk in the storage pool has previously been used in a different system, it will be reformatted and all data on the disk will be deleted.



The hard drives are configured to create a usable storage pool that can be divided into volumes for different applications or user groups. See [“Additional Information on DynamicRAID Sizing”](#) for additional details.

- When the storage pool has been successfully created on the head unit, a summary page is shown. Click **Next** to continue.

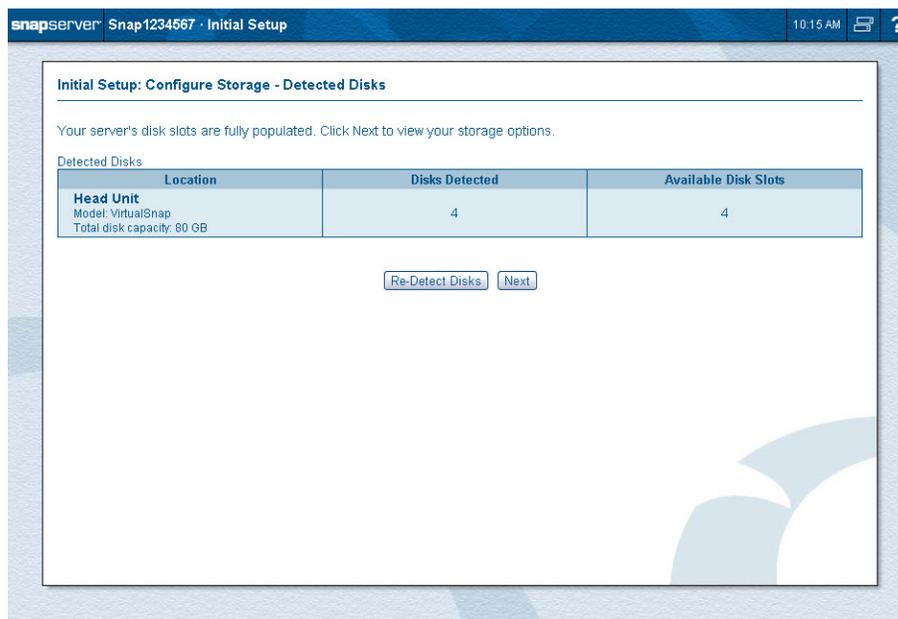


Traditional RAID Setup

Once Traditional RAID is selected, the wizard continues with the setup of the head unit. The first step is to detect the installed disk drives.

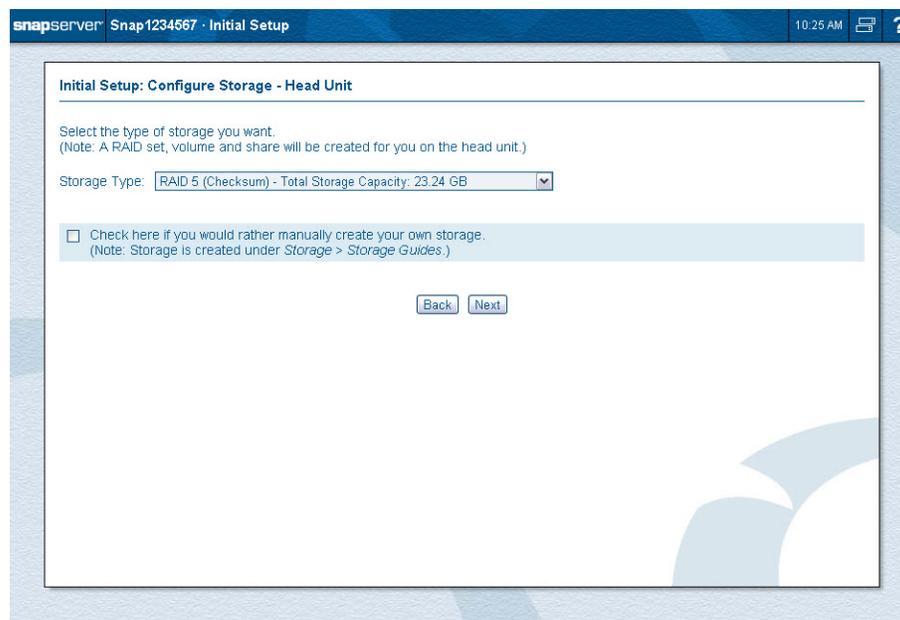
Detecting Available Disks

Available disks are detected and displayed on the Detected Disks setup page. You will have an opportunity to add more disks if empty slots exist. Click **Next** to continue.



Configure Storage – Head Unit

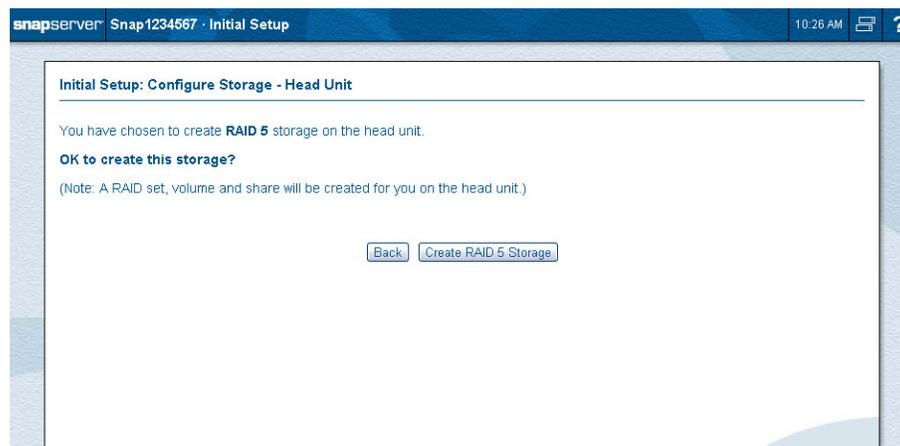
The next page lets you manually configure your head unit storage by selecting the RAID type based on the installed disk drives. See “Factors in Choosing a RAID Type”.



1. From the drop-down list, select the **storage type** (RAID) you want.
2. If you want to manually configure storage, check **Check here if you would rather manually create your own storage**.
3. Click **Next** to accept the settings.

NOTE: If the manual storage creation option was checked, you will be asked to confirm this choice at the next prompt and the wizard exits. You must go to Storage > Storage Guides to complete the setup.

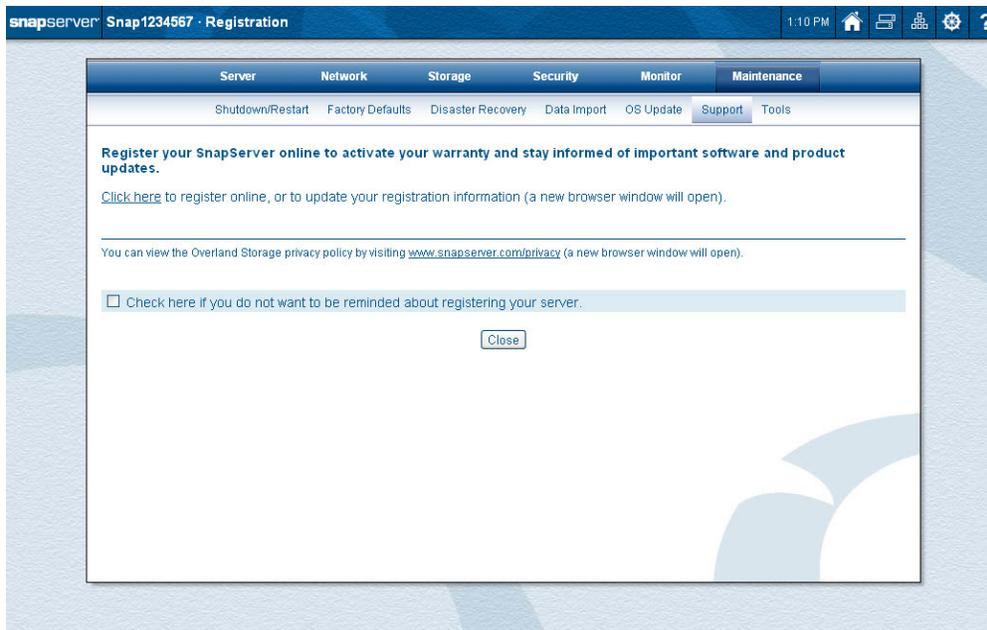
4. At the confirmation page, click **Create RAID *n* Storage** to proceed.



A RAID set, volume, and share are created automatically on the head unit with the default space reserved for snapshots equal to 20% of the volume's size.

Registration Page

After the setup wizard is done, you are given a chance to register your SnapServer. The **Registration** page appears (this page can also be accessed by clicking **Maintenance > Support > Registration**). Click as indicated to launch the Overland Storage Support website.



IMPORTANT: Because technical and warranty service are not available until your appliance is registered, it is recommended that you do so at this time. Registration is quick and easy.

Once registration is complete, you can check the box to no longer be reminded. Then click **Close** to finish your setup.

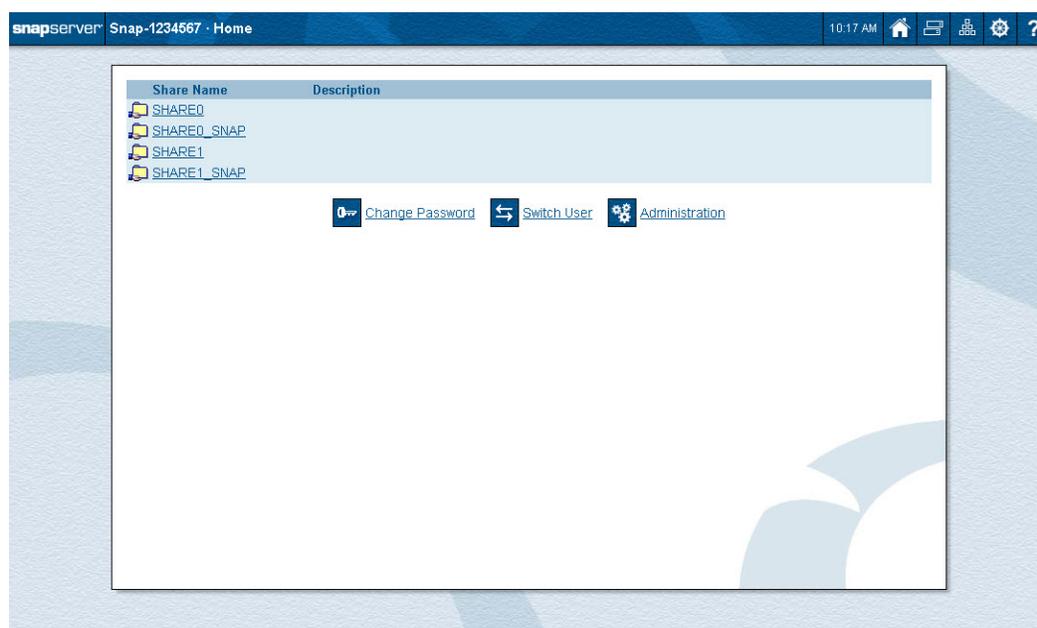


Server Status and Site Map

SnapServer appliances running GuardianOS 7.2 use a web-based graphical user interface (GUI) called the Web Management Interface. It supports most common web browsers, including Internet Explorer 7 and higher, Firefox 3 and higher, Apple Safari 5, and Google Chrome 9 and higher. JavaScript must be enabled in the browser.

When connecting to the server with a web browser, the Home page of the Web Management Interface is displayed. This page shows any shares at the top, the three primary options below the shares list, and has special navigation buttons displayed on the right side of the title bar (see the next table).

NOTE: If you haven't gone through the initial setup or authentication is required, you may be prompted to log in when you first access the Web Management Interface.



The **Home** page displays the following icons and options:

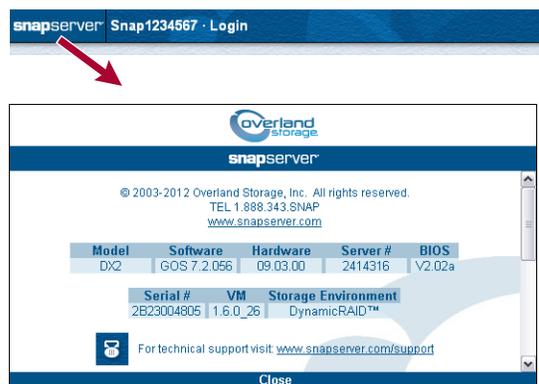
Icons & Options	Description
Change Password 	Click to access the password change page. Passwords are case sensitive. Use up to 15 alphanumeric characters.
Switch User 	Click to log out and open the login dialog box to log in as a different user.
Administration 	Click to administer the server. If you are not yet logged in, you are prompted to do so.
Navigation Buttons 	The following Navigation buttons are present in the upper right on every Web Management Interface page: Home – Click this to switch between the Web Home page and the Admin Home page. If you have not yet logged in to the Admin Home page, only the Web Home page is available.

Icons & Options	Description
	Snap Finder – Click this to view a list of all SnapServers found on your network and to specify a list of remote servers that will be used to discover SnapServers on other subnets. You can access these servers by clicking the listed name or IP address.
	SnapExtensions – Click this to view the SnapExtensions page, where you can acquire licenses for and configure third-party applications.
	Site Map – Click this to view a Site Map of the available options in the Web Management Interface, where you can navigate directly to all the major utility pages. The current page is shown in orange text.
	Help – Click this to access the online help for the UI page you are viewing.
UI Appearance	Click the Mgmt. Interface Settings link in the Site Map to choose a background for the Web Management Interface. You can select either a solid-colored background or a textured-graphic background.

For more information, see “Home Pages – Web/Admin.”

Hardware Information

From the Web Management Interface, click the SnapServer logo in the upper left corner to display the pertinent hardware information and contact links:



Scroll down to view additional contact information. Click outside the box to dismiss.

Scheduling Data Protection Tasks

Scheduling backups, snapshots, and antivirus scans, and creating a disaster recovery image preserves your server configuration and protects your data from loss or corruption.

Snapshots can be taken to provide a point-in-time image of files and changes to files to help in quickly recovering from accidental deletion or modification, or to facilitate performing an offline tape backup of an active data partition.

Navigate to **Storage > Snapshots** in the browser-based Web Management Interface to create or schedule snapshots. Snapshots should be taken when the system is idle or under low data traffic. To modify the space available for storing snapshots:

- In DynamicRAID mode, go to **Storage > Storage Pools**.
- In Traditional RAID mode, go to **Storage > Snapshots > Snapshot Space**.

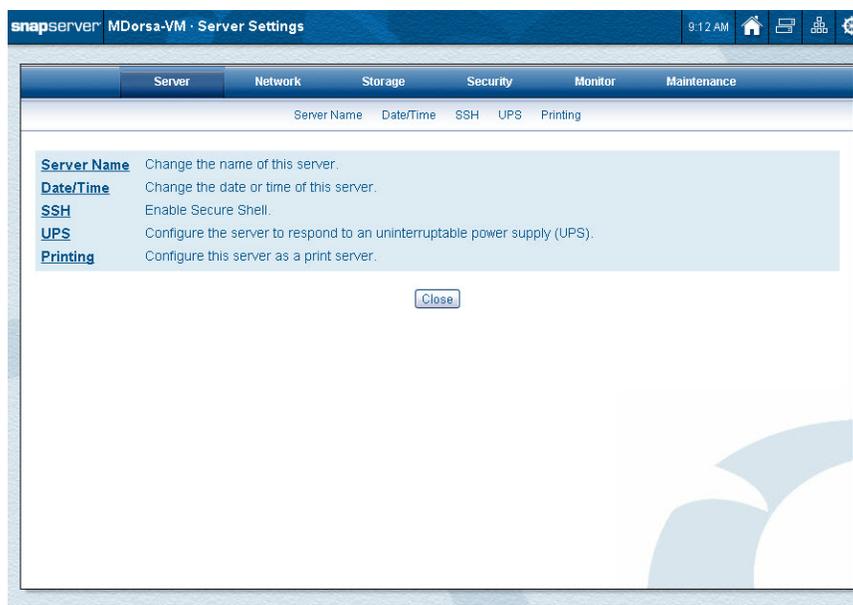
Set up antivirus protection by clicking the **SnapExtensions** icon (🔧), and then clicking **CA Antivirus**. Click the checkbox to enable antivirus, and click **OK**. When the configuration link appears, click it to launch the administration user interface for configuration and scheduling of virus scans and virus signature file updates.

Create a disaster recovery image on the **Maintenance > Disaster Recovery** page. This image should be created after the server configuration is complete, and can be used to recover the server or a replacement server to the configured state. See [Chapter 9, “Disaster Recovery,”](#) for detailed information on creating and using disaster recovery images.

GuardianOS contains built-in support for Snap EDR (trial mode) to synchronize and back up to and from other SnapServers. GuardianOS also supports several third-party backup agents. For information on using these backup methods to help protect your data, see [Appendix B, “Backup Solutions.”](#)

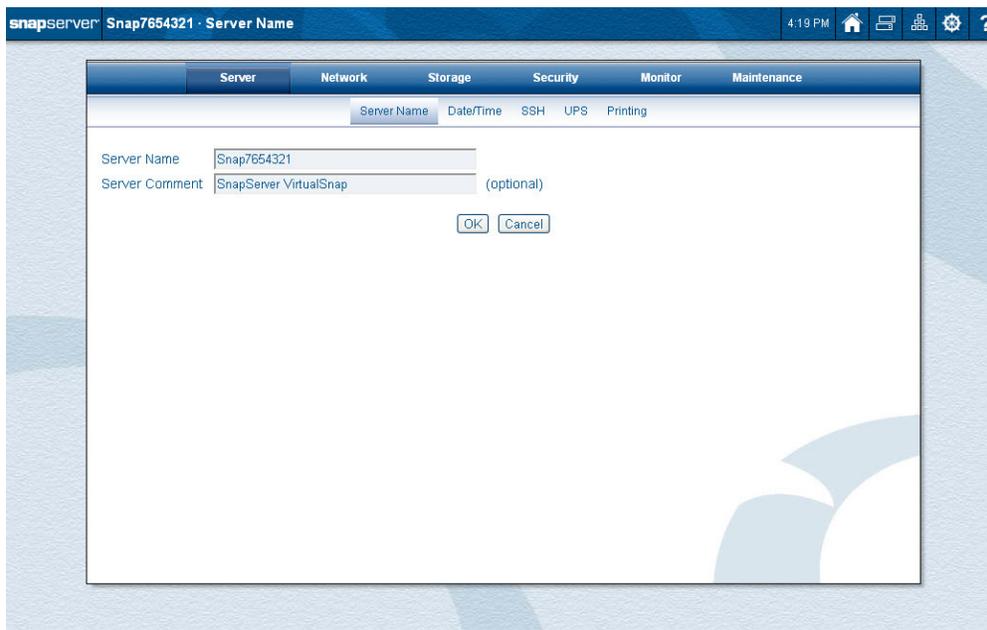
Server Options

By clicking either the Administration link (or using the GuardianOS 7.2 site map link (🔧)), you have immediate access to five editable server options.



Server Name

Use this option to change the server name and add a comment.

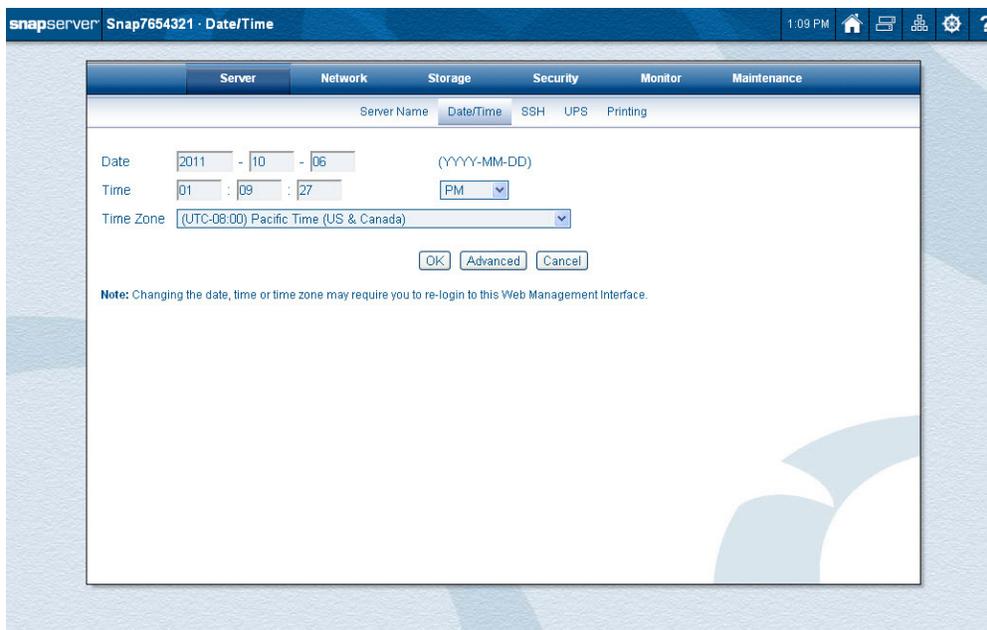


Edit the following fields and click **OK**.

Option	Description
Server Name	<p>The default server name is SNAPnnnnnnn, where nnnnnnn is your server number. For example, the default name for a SnapServer with the serial number 1234567 would be Snap1234567. If desired, enter a unique server name of up to 15 alphanumeric characters. In addition to letters and numbers, you can also use a dash (-) between characters, but spaces are not allowed.</p> <p>NOTE: The server number can be found on the Monitor > System Status page.</p>
Server Comment	<p>Optionally, add a comment (for example, server location) specific to the server.</p>

Date/Time

Use this page to configure date and time settings in ISO 8601 formatting. The time stamp applies when recording server activity in the Event Log (Monitor tab), when creating or modifying files, and when scheduling snapshot or antivirus operations.



The screenshot displays the 'Date/Time' configuration page in the SnapServer GuardianOS 7.2 web management interface. The page title is 'SnapServer Snap7654321 · Date/Time'. The interface includes a navigation bar with tabs for 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Below the navigation bar, there are sub-tabs for 'Server Name', 'Date/Time', 'SSH', 'UPS', and 'Printing'. The 'Date/Time' sub-tab is active, showing the following configuration options:

- Date:** 2011 - 10 - 06 (format: YYYY-MM-DD)
- Time:** 01 : 09 : 27 (format: HH : MM : SS) with a dropdown menu set to PM.
- Time Zone:** (UTC-08:00) Pacific Time (US & Canada) (dropdown menu)

At the bottom of the configuration area, there are three buttons: 'OK', 'Advanced', and 'Cancel'. Below the buttons, a note reads: 'Note: Changing the date, time or time zone may require you to re-login to this Web Management Interface.'



CAUTION: If the current date and time are reset to an earlier date and time, the change does not automatically propagate to any scheduled events you have already set up for snapshot, antivirus, or Snap EDR operations. These operations will continue to run based on the previous date and time setting. To synchronize these operations with the new date and time settings, you must reschedule each operation.

To Configure Basic Date and Time Settings

Edit date and time settings as described in the following table, and then click **OK**.

Option	Description
Date	Enter the current date in the format indicated.
Time	Enter the current time in the format indicated.
Time Zone	Select the time zone that you want to use for this server.

NOTE: GuardianOS automatically adjusts for Daylight Saving Time, depending on your time zone.

To Configure Advanced Date and Time Settings

If a Network Time Protocol (NTP) server is to be used to manage the date and time, use the Advanced page to set up those parameters.

1. Click the **Advanced** button for the additional options:

Option	Description
Set the server's date & time (blue highlight)	<ul style="list-style-type: none"> • Set the server's date & time to the following – Select this option and enter the date and time in the format indicated. • Automatically synchronize this server's date & time to the following NTP servers – Select the option and enter the server name of one or more NTP servers. You can find a list of public NTP servers at http://www.ntp.org. <p>NOTE: Due to security restrictions, NTP cannot be used when a server is joined to an Active Directory domain.</p>
Enable this server as an NTP server	Check this box to enable the SnapServer as an NTP server.
Time Zone	Select the time zone that you want to use for this server.

2. Click **OK** when finished.

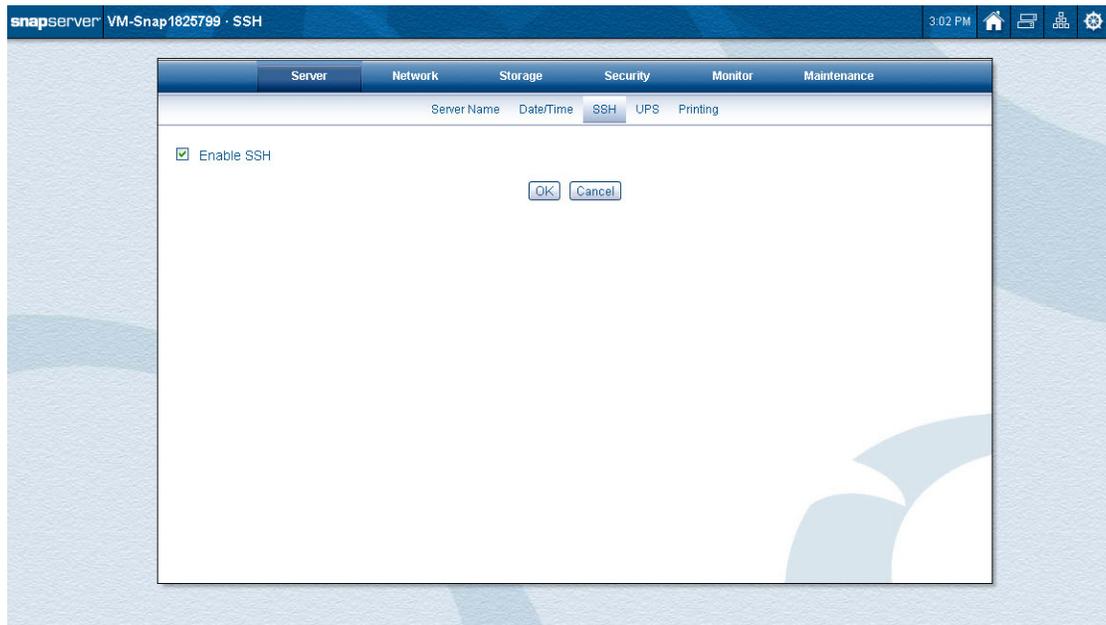
Secure Shell

Secure Shell (SSH) is a service that provides a remote console to access a command line shell that allows the user to perform basic management and update functions outside the GuardianOS Web Management Interface. See “[Command Line Interface](#)” on page E-1 for more information. The SSH implementation requires SSH v2.

NOTE: To maintain security, consider disabling SSH when not in use.

Disable SSH

SSH is enabled by default. To disable SSH, uncheck the **Enable SSH** check box, and click **OK**.



Connect to the CLI using SSH

1. Verify that your remote machine has an **SSH client application** installed.
Free or low-cost SSH applications are available from the Internet.
2. Connect to the server using its **IP address**.
Before the Initial Setup Wizard is completed and storage is configured, SnapCLI disables and hides all standard commands and makes only the system command available. See “[Before You Begin](#)”.
3. Log in as **admin**.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.

You will automatically be placed in the CLI shell.

For information about starting a supported backup agent using SSH, see “[Backup and Replication Solutions Table](#)” on page B-1.

UPS Protection

SnapServers support automatic shutdowns when receiving a low-power warning from an APC uninterruptible power supply (UPS). Use this page to manage this feature:

NOTE: If you are not using a UPS and your server supports disabling write cache, consider disabling write cache to help protect your data in case of a power outage.

An APC Smart-UPS® series device allows the SnapServer to shut down gracefully in the event of an unexpected power interruption. You can configure the server to automatically shut down when a low power warning is sent from an APC network-enabled or USB-based UPS device (some serial-only APC UPS units are also supported by using the IOGear GUC232A USB to Serial Adapter Cable). To do this, you must enable UPS support on the SnapServer, as described in this section, to listen to the IP address of one or two APC UPS units, and you must supply the proper authentication phrase configured on the UPS. Some SnapServer products have a single power supply, allowing you to attach a single UPS device. Other products have dual power supplies, allowing you to attach two UPS devices.

NOTE: Select a UPS capable of providing power to the SnapServer for at least ten minutes. In addition, in order to allow the SnapServer sufficient time to shut down cleanly, the UPS should be able to provide power for at least five minutes after entering a low battery condition.

Procedure to Configure UPS Protection

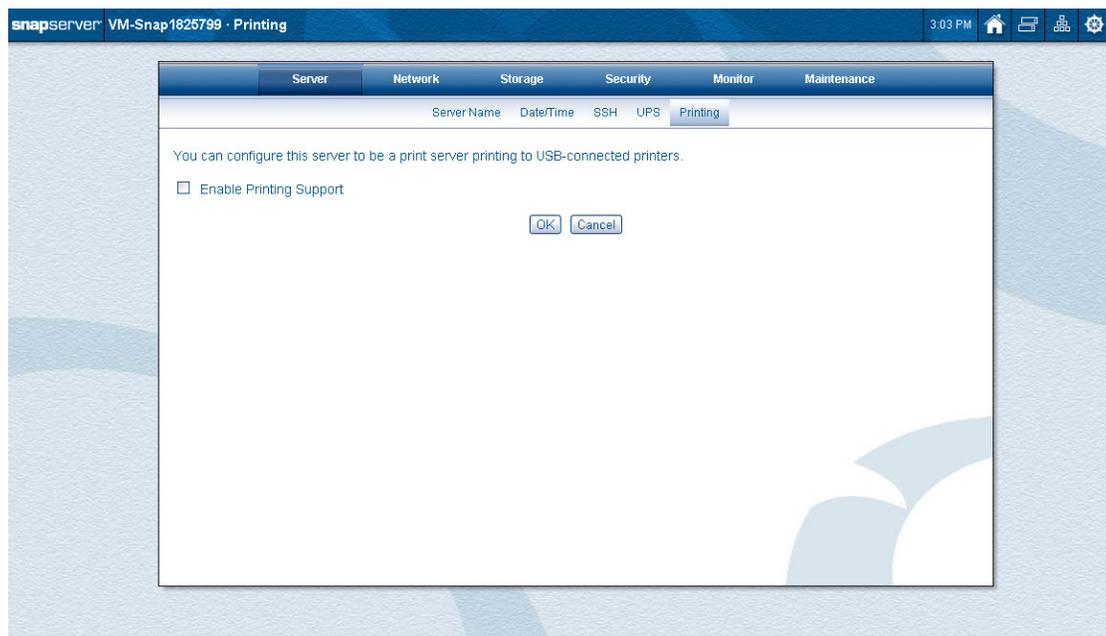
1. Complete the following fields:

Option	Description
Enable UPS Support	Check the Enable UPS Support check box to enable; leave the check box blank to disable UPS support.
Automatically restart server...	Check this box to automatically restart the server when power has been restored or the UPS comes back online. Leave the check box blank to manually start the server after a power failure.
Use a single USB-connected UPS device	Select this option button to use a USB-connected APC UPS device or serial UPS with USB to serial adapter cable. NOTE: If using a serial UPS with a USB-to-serial adapter cable, reboot the SnapServer after connecting the cable to the server to properly initialize the connection to the UPS.
APC Status	Under the selected UPS connection type, an APC status field will display the following possible values: Unknown, No Connection, Low Battery, On Battery, and Online.
Use the following network-connected UPS devices	Select this option button to use up to two network-connected APC UPS devices.
IP Address	Enter the IP address of the network UPS device.
APC User Name	Enter the APC Administrator user name. NOTE: The APC user name entered must be the APC Administrator name for the UPS (by default, apc).
APC Authentication Phrase	Enter the authentication phrase configured for shutdown behavior on the UPS (in the UPS Web UI, this can be configured in PowerChute settings or, for older firmware, in the User Manager for the administrator user). NOTE: This password phrase is not the same as the user's password.
Secondary UPS device (optional)	Check the Secondary UPS device check box to enable; leave the check box blank to disable secondary UPS support.
Low Battery Alert	Select one of the following: <ul style="list-style-type: none"> • Either UPS Device: Select this option to allow shutdown upon receipt of a message from either of the two specified UPS servers. • Both UPS Devices: Select this option to allow shutdown only upon receipt of one message from each of the two specified UPS servers.

2. Click **OK** to finish.

Print Server

The SnapServer can be configured to emulate a Windows print server for locally-attached USB printers or IPP printing. Client machines connect to the SnapServer over the network and use the printer similarly to using a printer shared by a Windows or CUPS server. You can pause or resume the printer, and monitor or cancel print jobs using the Web Management Interface.



Configuring your as a print server is a two-part process:

- Step 1:** Configuring the **printer** on the SnapServer.
- Step 2:** Configuring the **client** to print via the SnapServer.

Procedure to Configure the Printer

First, you need to configure the printer connected to the SnapServer.

1. Check **Enable printing support**.
2. Connect a printer to one of the USB ports on the SnapServer.
3. Power ON the printer.
4. In the SnapServer's Web Management Interface, navigate to **Server > Printing**. A list of currently defined USB printers is displayed. To add the new printer, click **Add Local Printer**.
5. The SnapServer will detect the new printer and it should appear as an option in the **Local Printer Device** drop-down list. Select that printer.
6. Give the printer a name, and complete Description and Location information as desired. Click **OK**. The printer will appear in the list on the main printing page.

Procedure to Configure the Client

Next, add the printer to a Windows, Mac, or Linux client, enabling you to print via the SnapServer. The SnapServer supports both Windows SMB and IPP printing protocols.

NOTE: To make printer drivers easily accessible to users, copy them to a share that everyone can access on the SnapServer. The SnapServer cannot be configured to automatically provide printer drivers to clients.

Adding the Network Printer to a Windows Client

Windows offers several methods for adding a printer. Follow your usual printer configuration method to add a printer shared on a SnapServer. When asked to locate the printer:

- To use SMB, enter the SnapServer name or IP address, or browse to the server to choose the printer share.
- To use IPP, enter the exact path as follows in the URL field:
http://servername:631/printers/sharename
where *servername* is the name or IP address of your SnapServer and *sharename* is the name of the printer.

NOTE: 631 is the IPP port number.

If you experience difficulty adding the printer, try the following:

1. Navigate to **Start > Run** and enter the server name as follows:
\\servername
2. After a delay, you may be prompted for a user name and password. Log in as a user with access to the SnapServer.
3. A Windows Explorer window opens displaying all shares and printers on the server. Right-click the server and choose **Connect**.
4. Follow the instructions to provide the printer driver and complete the setup.

To Add a Network Printer to a Mac OS X Client

Add a printer using your usual method. If you are using SMB, you will need to know the SnapServer name. If you are using IPP, you will need to enter the IP address in the **Type** field and the printer and sharename in the **Queue** field.

To Add a Network Printer to a Linux Client

Add a printer using your usual method. If you are using SMB, you will need to know the SnapServer name. If you are using IPP, enter the exact path as follows in the URL field:

http://servername:631/printers/sharename

where *servername* is the name or IP address of your SnapServer and *sharename* is the name of the printer.

NOTE: 631 is the IPP port number.

To Monitor Print Jobs Remotely

Pause or resume the printer, and check the status of or cancel print jobs from the SnapServer's Web Management Interface.

To Pause the Printer

Use this procedure to pause the printer:

1. Navigate to **Server > Printing**.

2. Click the **Status** link next to your printer to open the **Job Status** window and see your print job queue.
3. Click the **Pause Printer** button to pause all print jobs.
When the printer is paused, the button becomes a **Resume Printer** button, which you can click to resume printing.

To Cancel or Check the Status of Print Jobs

Use this procedure to cancel or check the status of a print job:

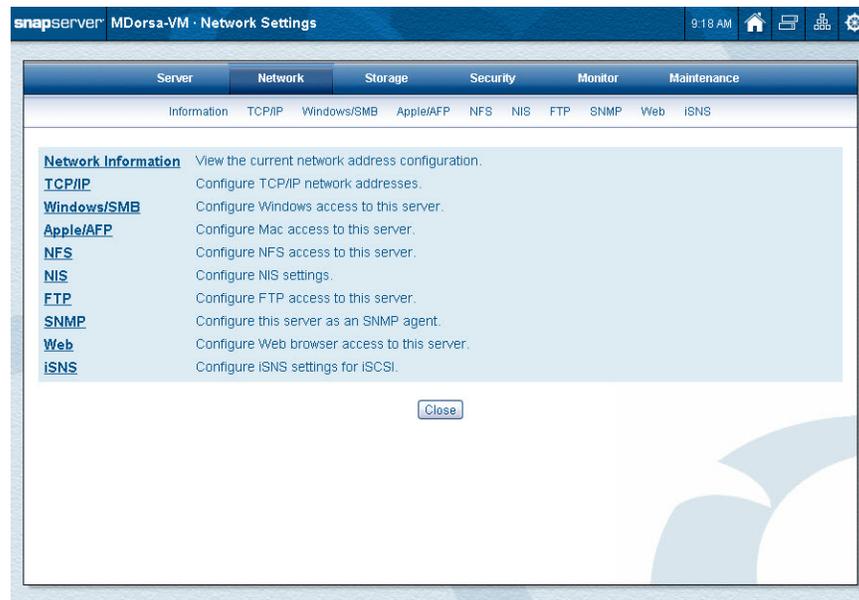
1. Navigate to **Server > Printing** and click the **Status** link next to your printer to open the Job Status window and see your print job queue.
2. To cancel a print job, click to put a check in the box next to the job you want to remove and click **Cancel Selected Jobs**. You can select to cancel multiple jobs. If you want to cancel all the listed print jobs, click the **Cancel All Jobs** button. Click the **Refresh** button to update the page with the current list of print jobs.

To Delete a Printer

When you remove a printer, remember to remove its information from both the Web Management Interface and the client machines.

1. Disconnect the printer cable from the SnapServer.
2. In the Web Management Interface, navigate to **Server > Printing**. In the list of printers, the status of printer you just removed should appear as **Offline**.
3. Click the printer link to open the **Edit Printer** page, then click the **Delete** button to delete the printer.

SnapServers are preconfigured to use DHCP, autonegotiate network settings, and allow access to the server for Windows (CIFS/SMB), UNIX (NFS), Mac (AFP), FTP/FTPS, and HTTP/HTTPS clients. This chapter addresses the options for configuring TCP/IP addressing, network bonding, and access protocols. Network bonding options allow you to configure the SnapServer for load balancing and failover. Network protocols control which network clients can access the server.



Topics in Network Access:

- [View Network Information](#)
- [TCP/IP Options](#)
- [Windows Networking \(SMB\)](#)
- [Apple Networking \(AFP\)](#)
- [NFS Access](#)
- [NIS Domain](#)
- [FTP/FTPS Access](#)
- [SNMP Configuration](#)
- [Web Access](#)
- [iSNS Configuration](#)

 **IMPORTANT:** The default settings enable access to the SnapServer via all protocols supported by the SnapServer. As a security measure, disable any protocols not in use. For example, if no Mac or FTP clients need access to the SnapServer, disable these protocols in the Web Management Interface.

View Network Information

The **Network > Information** page displays the server's current network settings. One column appears for each Ethernet port in use.



Ethernet Interface Information		
Port Name	Ethernet 1 (Primary Interface)	Ethernet 2
Enabled	Yes	Yes
TCP/IP Mode	Static	Static
IP Address	10.20.16.60	172.10.16.60
Subnet Mask	255.255.0.0	255.255.0.0
Primary WINS Server	10.20.18.172	-
Secondary WINS Server #1	-	-
Secondary WINS Server #2	-	-
Secondary WINS Server #3	-	-
Ethernet Address	00:C0:B6:24:CA:28	00:C0:B6:24:CA:29
Speed Status	100 Mbps (Auto)	1000 Mbps (Auto)
Duplex Status	Full Duplex (Auto)	Full Duplex (Auto)
Bonding Status	Standalone	Standalone

Gateway Information	
Default Gateway	10.20.8.1

DNS Information	
Domain Name	devnet.myoverland.net
Primary DNS	10.20.18.172
Secondary DNS #1	-
Secondary DNS #2	-

Field definitions are given in the following table:

Ethernet Interface Information	
Port Name	The name of the Ethernet interface.
Enabled	Yes or No.
TCP/IP Settings Obtained From	DHCP or Static.
IP Address	The unique 32-bit value that identifies the server on a network subnet.
Subnet Mask	Combines with the IP address to identify the subnet on which the server is located.
Primary WINS Server	The Windows Internet Naming Service server, which locates network resources in a TCP/IP-based Windows network by automatically configuring and maintaining the name and IP address mapping tables.
Secondary WINS Servers	Secondary Windows Internet Naming Service servers. Up to three secondary servers can be used.

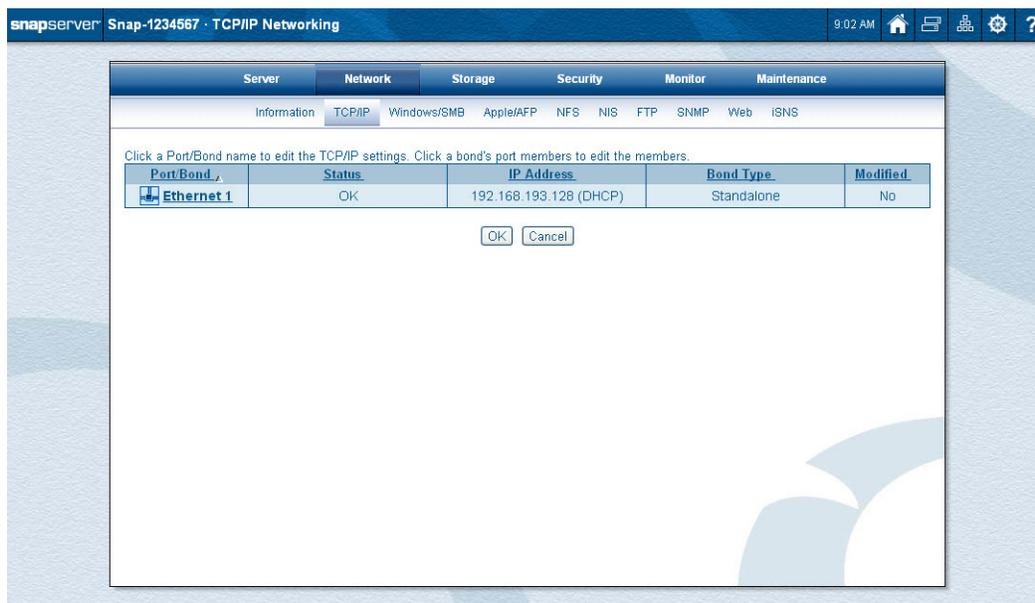
Ethernet Interface Information	
Ethernet Address	The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet port.
Speed Status	10 Mbps, 100 Mbps, or 1000 Mbps.
Duplex Status	Half-duplex: two-way data flow, only one way at a time. Full-duplex: two-way data flow simultaneously.
Bonding Status	Standalone, Load balance, Failover, Switch Trunking, or Link Aggregation.

Gateway Information	
Default Gateway	The network address of the gateway is the hardware or software that bridges the gap between two otherwise unroutable networks. It allows data to be transferred among computers that are on different subnets.

DNS Information	
Domain Name	The ASCII name that identifies the Internet domain for a group of computers within a network.
Primary DNS	The IP address of the primary Domain Name System server that maintains the list of all host names.
Secondary DNS (servers)	Up to two secondary Domain Name System servers can be used.

TCP/IP Options

SnapServers ship with one or more Gigabit Ethernet (GbE) ports. The information about those ports is displayed on the primary TCP/IP page:

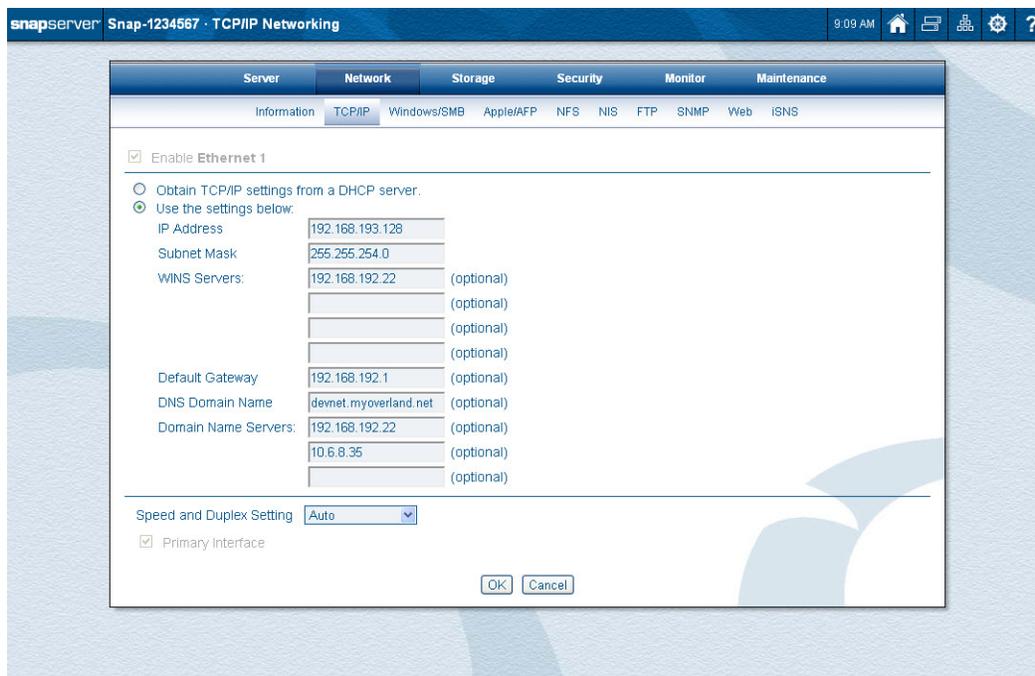


The following table describes the port information:

Column	Description
Port/Bond	A list of the Ethernet Ports or Bonds on the server. Click a port or bond name to display or modify configuration details. See "Configuring TCP/IP Settings" .
Status	<ul style="list-style-type: none"> • OK – Port is connected and active. • No link – Port is not connected. • Failed – Port has failed.
IP Address	<ul style="list-style-type: none"> • The IP address for the NIC or bond if known or not available if unknown. • Whether the IP address was obtained by DHCP or is Static.
Bond Type	<p>NOTE: If you have more than two ports, you can have a mixture of standalone and bonded ports. For example, on a 4-port system, one port can be a standalone and the other three ports can be bonded into a load balanced configuration.</p> <ul style="list-style-type: none"> • Standalone – The default state <i>Standalone</i> is the absence of network bonding and treats each port as a separate interface. • Load Balance (ALB) – An intelligent software adaptive agent repeatedly analyzes the traffic flow from the server and distributes the packets based on destination addresses, evenly distributing network traffic for optimal network performance. All ports in the same ALB configuration need to be connected to the same switch. • Failover – This mode uses one Ethernet port (by default, <i>Ethernet 1</i>) as the primary network interface and a one or more Ethernet ports are held in reserve as the backup interface. Redundant network interfaces ensure that an active port is available at all times. If the primary port fails due to a hardware or cable problem, the second port assumes its network identity. The ports should be connected to different switches (though this is not required). <p>NOTE: Failover mode provides switch fault tolerance, as long as ports are connected to different switches.</p> <ul style="list-style-type: none"> • Switch Trunking – This mode groups multiple physical Ethernet links to create one logical interface. Provides high fault tolerance and fast performance between switches, routers, and servers. • Link Aggregation (802.3ad) – Like Switch Trunking, this mode groups multiple physical Ethernet interfaces to create one logical interface, and provides high fault tolerance and fast performance between switches, routers, and servers. Uses Link Aggregation Control Protocol (LACP) to autonegotiate trunk settings.
Modified	<p>Indicates whether configuration for one or more interfaces has been changed and needs to be applied to take effect:</p> <ul style="list-style-type: none"> • Yes – One or more parameters for the interface have been modified. • No – None of the parameters for the interface have been modified.

Configuring TCP/IP Settings

To configure the TCP/IP settings of a specific port or bond, click the name on the primary TCP/IP page. A secondary TCP/IP page displays the configuration options for the Ethernet port selected.



The following table describes these TCP/IP options.

Option	Setting	Description
Enable Ethernet 1	Checked	By default, all Ethernet ports are enabled, whether they are used or not.
	Unchecked	Ports other than the Primary Interface (by default, <i>Ethernet 1</i>) can be disabled by selecting the port and unchecking the Enable Ethernet 1 checkbox. However, a bonded Ethernet port cannot be disabled, nor can a disabled Ethernet port be placed in bonded mode. NOTE: The primary Ethernet port must always be enabled. GuardianOS will not allow you to disable it.
TCP/IP Addressing	DHCP	By default, SnapServers acquire an IP address from the DHCP server on the network.
	Static	Administrators may assign a fixed IP address or other IP settings as needed.

Option	Setting	Description
Speed and Duplex Setting	Auto	The default setting of Auto enables automatic negotiation of the speed and duplex settings based on the physical port connection to a switch. The speed setting establishes the rate of transmission and reception of data. The duplex setting allows the Ethernet port to transmit and receive network packets simultaneously. NOTE: Auto is the only allowable setting for a Gigabit port.
	Fixed	The SnapServer may also be set to fixed speed/duplex setting: 10Mbps/half; 10Mbps/full; 100Mbps/half; 100Mbps/full. NOTE: To prevent connectivity problems when changing to a fixed setting, see “Changing from Auto to a Fixed Link Setting” on page 3-7.
Primary Interface	Checked or Unchecked	By default, the primary Ethernet port is Ethernet 1 and it cannot be disabled. However, the Primary Interface can be changed to a different Ethernet port by selecting the Ethernet port you want as the Primary port and checking the Primary Interface box. The Primary Interface is prioritized for various network configuration parameters that apply to the server as a whole (for example, DNS IP address, hostname, and default gateway). In addition, the IP address of the Primary Interface is preferred to identify the server for various services and circumstances that require a single IP address.

Issues in TCP/IP Configuration

Consider the following guidelines when connecting a SnapServer to the network.

Cabling for Single-Subnet, Multihomed, or Network Bonding Configurations

- For a **Single Subnet** or **Multihomed** Configuration (Standalone) – Standalone treats each port as a separate interface. In a single-subnet configuration, only the primary port is connected to the switch. In a multihomed configuration, each port is cabled to a different switch and the network connections lead to separate subnets.

 **CAUTION:** Do not connect multiple Ethernet ports to the same network segment in Standalone mode, except for iSCSI MPIO configurations. This configuration is not supported by most network file protocols and can lead to unexpected results.

If you connect only one port, use the default primary port (*Ethernet 1*). If you use Ethernet 2 or any other non-primary port, some services may not function properly.

- For a **Network Bonding** Configuration (Load Balancing, Failover, Switch Trunking, or Link Aggregation) – Network bonding technology treats multiple ports as a single channel, with the network using one IP address for the server.

NOTE: This network bonding configuration is only applicable to SnapServers with more than one Ethernet port. To take advantage of network bonding, all ports in the bonded team must be physically connected to the same network:

- For load balancing, Switch Trunking, or Link Aggregation, these parts are connected to the same switch on the same subnet.

- For failover, these parts are connected to a different switch on the same subnet (in case one switch fails).

Make Sure the Switch is Set to Autonegotiate Speed/Duplex Settings

When the server is shipped from the factory, both ports are set to autonegotiate. This setting allows the SnapServer to base speed and duplex settings on the physical port connection to a switch. Thus, the switch/hub to which the SnapServer is cabled *must* be set to autonegotiate to initially connect to the server; otherwise, network throughput or connectivity to the server may be seriously impacted.

To use fixed duplex settings (not applicable to gigabit), the same fixed setting must be set on the server and switch.

Configure the Switch for Load Balancing

If you select either the Switch Trunking or Link Aggregation network bonding configuration, be sure the switch is configured correctly for that bonding method. No switch configuration is required for Adaptive Load Balancing (ALB).

Changing from Auto to a Fixed Link Setting

You can configure a fixed link speed and duplex setting on the **Network > TCP/IP** page in the browser-based Web Management Interface. If you change this setting, you must:

1. Configure the **fixed setting** in the Web Management Interface first.
2. Configure the **switch** to the same fixed setting.



IMPORTANT: If you change the switch setting before you change the setting in the Web Management Interface, the SnapServer may not connect to the network. The **Link LED** on the SnapServer front panel will be off or amber if the server is not connected to the network.

Procedure to Create a Bond

On a SnapServer with two or more Ethernet ports, a bond can be created:

1. Click the **Create Bond** button.
 2. On the page that opens, select a bonding mode:
 - a. **Bond Type** – Choose one of the following settings:
 - *Load Balance (ALB)* – enables all selected ports to share the network load.
 - *Failover* – enables other selected ports to automatically take over the connection if the primary port fails. Only one port is active at any given time.
 - *Switch Trunking* and *Link Aggregation (802.3ad)* – these two options group multiple Ethernet ports into one logical Ethernet port for high speed and fault tolerance.
- Ports not joined to a bond are configured as Standalone and have separate interfaces (one IP address per port).
- b. **Select Ports to include in Bond** – Select the ports you want to include in the Bond from the Ports not in Bond n column and use the **Add** button to move them to the Ports in Bond n column.
3. Click **OK** to close the Configuration page, then click **OK** again to apply the changes.

Windows Networking (SMB)

Windows SMB and security settings are configured on the **Network > Windows/SMB** page of the Web Management Interface. You can configure these settings as a member of a **Workgroup** or an **Active Directory Domain**, as shown below:

The screenshot shows the SnapServer Web Management Interface for the 'Windows/SMB' configuration page. The page is titled 'SnapServer Snap7654321 · Windows/SMB' and shows the time as 1:41 PM. The navigation menu includes Server, Network, Storage, Security, Monitor, and Maintenance. The 'Network' menu is expanded to show sub-options: Information, TCP/IP, Windows/SMB, Apple/AFP, NFS, NIS, FTP, SNMP, Web, and iSNS. The 'Windows/SMB' sub-page is active, showing the following configuration options:

- Enable Windows Networking (SMB)
- Member of:
- Workgroup Name:
- Enable Guest Account
- Enable Opportunistic Locking
- Enable this server as the Master Browser
- Allow root authentication

Buttons for 'OK' and 'Cancel' are located at the bottom right of the configuration area.

The screenshot shows the SnapServer Web Management Interface for the 'Windows/SMB' configuration page. The page is titled 'SnapServer Snap7654321 · Windows/SMB' and shows the time as 1:39 PM. The navigation menu is the same as in the previous screenshot. The 'Windows/SMB' sub-page is active, showing the following configuration options:

- Enable Windows Networking (SMB)
- Member of:
- Domain Name:
- Administrator Name: (Required to join a domain)
- Administrator Password: (Required to join a domain)
- Organizational Unit:
- LDAP Signing:
- Enable Guest Account
- Enable Opportunistic Locking
- Enable this server as the Master Browser
- Allow root authentication
- Disable NetBIOS over TCP/IP
- Enable Trusted Domains

Buttons for 'OK' and 'Cancel' are located at the bottom right of the configuration area.

Support for Windows Networking (SMB)

The default settings make the SnapServer available to SMB clients in the workgroup named *Workgroup*. Opportunistic locking is enabled, as is participation in master browser elections.

Consider the following when configuring access for your Windows networking clients:

Support for Microsoft Name Resolution Servers

The SnapServer supports NetBIOS, WINS, and DNS name resolution services. However, when you use a domain name server with a Windows Active Directory (ADS) server, make sure the forward and reverse name lookup are correctly set up. ADS can use a UNIX BIND server for DNS as well.

ShareName\$ Support

GuardianOS supports appending the dollar-sign character (\$) to the name of a share in order to hide the share from SMB clients accessing the SnapServer.

NOTE: As with Windows servers, shares ending in '\$' are not truly hidden, but rather are filtered out by the Windows client. As a result, some clients and protocols can still see these shares.

To completely hide shares from visibility from any protocols, the **Security > Shares** page gives you access to a separate and distinct hidden share option that hides a share from SMB, AFP, HTTP, HTTPS, and FTP clients. However, shares are not hidden from NFS clients, which cannot connect to shares that aren't visible. To hide shares from NFS clients, consider disabling NFS access on hidden shares.

For new shares, select **Create Share** and click the **Advanced Share Properties** button to access the Hidden share option. For existing shares, select the share, click **Properties**, and click **Advanced Share Properties** to access the Hidden share option.

Support for Windows Network Authentication

This section summarizes important facts regarding the GuardianOS implementation of Windows network authentication.

Windows Networking Options

Windows environments operate in either workgroup mode, where each server contains a list of local users it authenticates on its own, or Active Directory (ADS) domain mode, where domain controllers centrally authenticate users for all domain members.

Option	Description
Workgroup	In a workgroup environment, users and groups are stored and managed separately on each server in the workgroup.
Active Directory (ADS)	<p>When operating in a Windows Active Directory domain environment, the SnapServer is a member of the domain and the domain controller is the repository of all account information. Client machines are also members of the domain and users log into the domain through their Windows-based client machines. Active Directory domains resolve user authentication and group membership through the domain controller.</p> <p>Once joined to a Windows Active Directory domain, the SnapServer imports and then maintains a current list of the users and groups on the domain. Thus, you must use the domain controller to make modifications to user or group accounts. Changes you make on the domain controller appear automatically on the SnapServer.</p> <p>NOTE: Windows 2000 domain controllers must run SP2 or later.</p>

Kerberos Authentication

Kerberos is a secure method for authenticating a request for a service in a network. Kerberos lets a user request an encrypted “ticket” from an authentication process that can then be used to request a service from a server. The user credentials are always encrypted before they are transmitted over the network.

The SnapServer supports the Microsoft Windows implementation of Kerberos. In Windows Active Directory (ADS), the domain controller is also the directory server, the Kerberos key distribution center (KDC), and the origin of group policies that are applied to the domain.

NOTE: Kerberos requires the server's time to be closely synchronized to the domain controller's time. This means that (1) the server automatically synchronizes its time to the domain controller's and (2) NTP cannot be enabled when joined to an ADS domain.

Interoperability with Active Directory Authentication

The SnapServer supports the Microsoft Windows 2000/2003/2008 family of servers that run in ADS mode. any SnapServer can join Active Directory domains as a member server. References to the SnapServer's shares can be added to organizational units (OU) as shared folder objects.

NOTE: Windows 2000 domain controllers must run SP2 or later.

Guest Account Access to the SnapServer

The **Network > Windows/SMB** page in the Web Management Interface contains an option that allows unknown users to access the SnapServer using the guest account.

Connect from a Windows Client

Windows clients can connect to the SnapServer using either the server name or IP address. To navigate to the server using Windows Explorer, use one of these procedures:

- For Microsoft Windows Vista, 2008, and 7 clients, navigate to **Network > server_name**.
- For Microsoft Windows 2003, 2000, or XP clients, navigate to **My Network Places > workgroup_name > server_name**.

Connect a Mac OS X Client Using SMB

Mac OS X clients can connect using SMB as well as AFP. Specify the server in the Connect to Server window (from **Finder** press **Cmd + K**, or select **Finder > Go > Connect to Server**) as one of the following:

- **smb://servername**
- **smb://ipaddress**

Tip: To disconnect from the SnapServer, drag its icon into the Trash.

You can also browse the servers in the Finder file window, under the Shared tab.

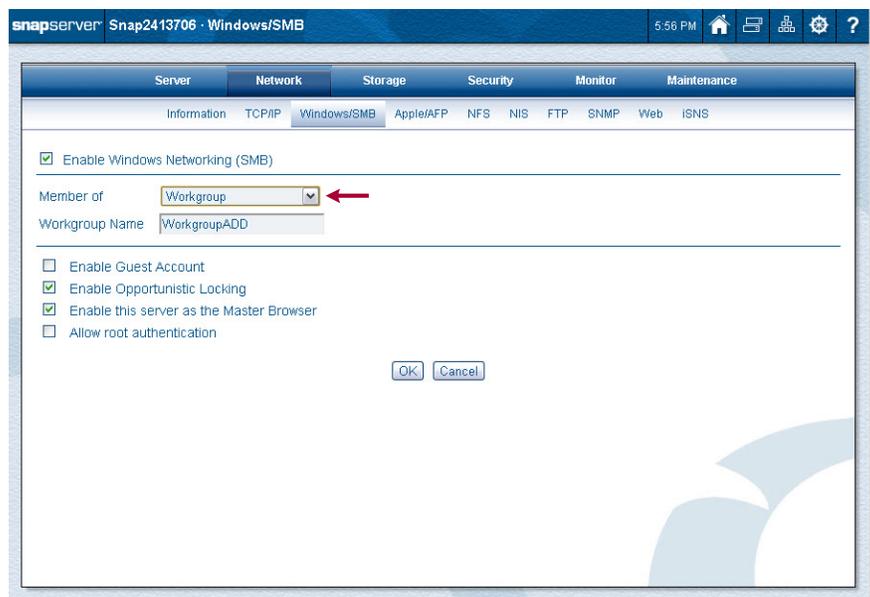
Configure Windows/SMB Networking

Windows SMB and security settings are configured from this page. The server can be configured as part of a Workgroup or an Active Directory Domain.

Before performing the configuration procedures provided here, be sure you are familiar with the information provided in “[Support for Windows Networking \(SMB\)](#)” and “[Support for Windows Network Authentication](#)”.

To Join a Workgroup

1. Go to **Network > Windows/SMB**.
2. At the Member list, verify that the default **Workgroup** is selected.



3. Edit the **fields** shown in the following table:

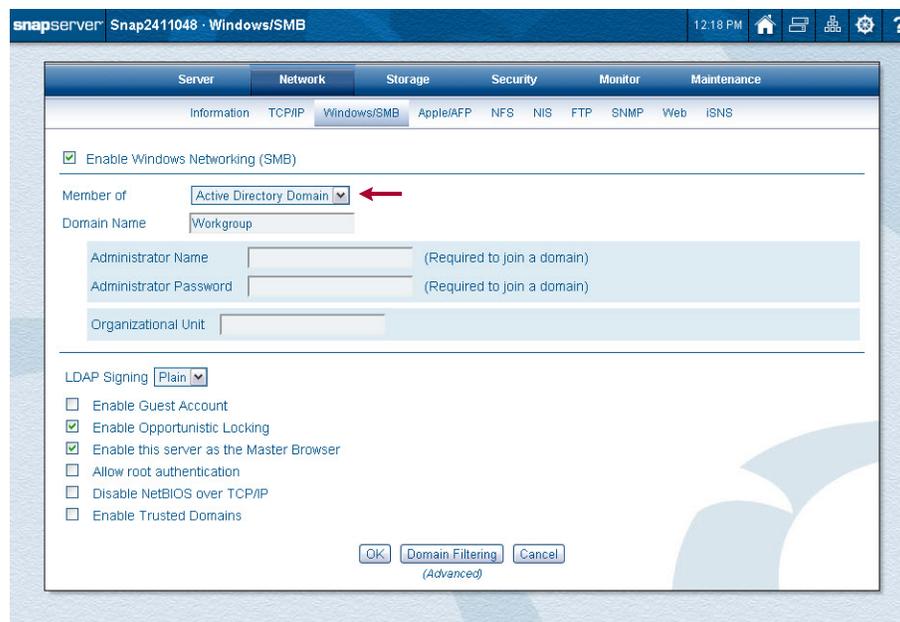
Option	Settings
Enable Windows SMB	Check the Enable Windows Networking (SMB) checkbox to enable SMB. Clear the checkbox to disable SMB.
Member Of	Verify it is set to Workgroup . NOTE: For the Active Directory Domain option, see “To Join an Active Directory Domain”.
Workgroup Name	The default settings make the SnapServer available in the workgroup named <i>Workgroup</i> . Enter the workgroup name to which the server belongs.
Enable Guest Account	Check the Enable Guest Account checkbox to allow unknown users or users explicitly logging in as “guest” to access the SnapServer using the guest account. Clear the option to disable this feature.
Enable Opportunistic Locking	Enabled by default. Opportunistic locking can help performance if the current user has exclusive access to a file. Clear the checkbox to disable opportunistic locking.
Enable this Server as the Master Browser	Enabled by default. The SnapServer can maintain the master list of all computers belonging to a specific workgroup. (At least one Master Browser must be active per workgroup.) Check the checkbox if you plan to install this server in a Windows environment and you want this server to be able to serve as the Master Browser for a workgroup. Clear the checkbox to disable this feature.

Option	Settings
Allow root authentication	<p>Check the Allow root authentication checkbox to allow a root login on the selected server.</p> <p>NOTE: The root password is synchronized with the admin password.</p>
Disable NetBIOS over TCP/IP (Active Dir Domain only)	<p>Some administrators may wish to disable NetBIOS over TCP/IP. Select the checkbox to disable NetBIOS; clear the checkbox to leave NetBIOS enabled.</p> <p>NOTE: If you disable NetBIOS and you are joining a domain, you must enter the domain name as a fully qualified domain name (for example, actdirdomainname.companyname.com). A short form such as ActDirDomName will not work.</p>
Enable Trusted Domains (Active Dir Domain only)	<p>SnapServers recognize trust relationships established between the domain to which the SnapServer is joined and other domains in a Windows environment by default. Select the checkbox to toggle this feature.</p> <p>NOTE: SnapServers remember trusted domains. That is, if this feature is disabled and then activated at a later time, the previously downloaded user and group lists, as well as any security permissions assigned to them, will be retained.</p>

4. Click **OK** to update Windows network settings immediately.

To Join an Active Directory Domain

1. Go to **Network > Windows/SMB**.
2. From the drop-down Member list, select **Active Directory Domain** to view the configuration page.



3. Edit the **fields** shown in the following table:

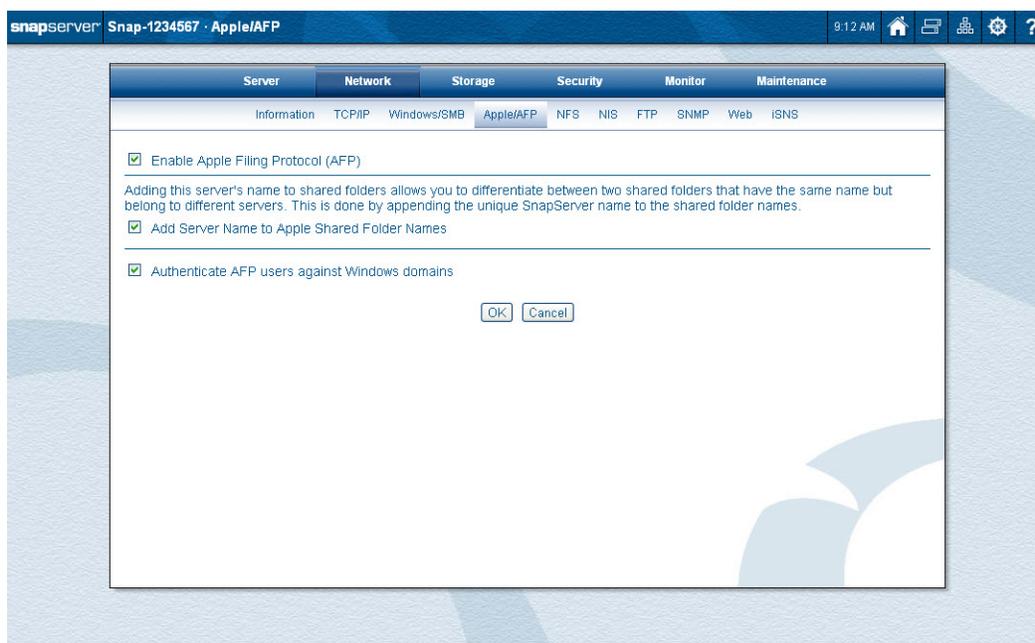
Option	Description
Enable Windows SMB	Check the Enable Windows Networking (SMB) checkbox to enable SMB. Clear the checkbox to disable SMB.
Member Of	Verify it shows <i>Active Directory Domain</i> .
Domain Name	<p>The default settings make the SnapServer available in the workgroup named <i>Workgroup</i>. Enter the domain name to which the server belongs.</p> <p>If you join a Windows domain through Advanced Security (Security > Security Guides > appropriate option), the domain name you entered displays here and can be changed on the next page or the Advanced Security page.</p> <p>NOTE: Windows 2000 domain controllers must run SP2 or later.</p>
Administrator Name and Password	If joining a domain, enter the user name and password of a user with domain join privileges (typically an administrative user).
Organizational Unit	<p>To create a machine account at a different location than the default, enter a name in the Organizational Unit field. By default, this field is blank, signaling the domain controller to use a default defined within the controller.</p> <p>NOTE: Sub-organizational units can be specified using Full Distinguished Name LDAP syntax or a simple path ([organizationalunit]/[sub-unit1]/[sub-unit1a])</p>
LDAP Signing	Set ADS domain LDAP signing to Plain (no signing), Sign, or Seal, as appropriate for your domain. Default setting is Plain.
Enable Guest Account	Check the Enable Guest Account checkbox to allow unknown users or users explicitly logging in as "guest" to access the SnapServer using the guest account. Clear the option to disable this feature.
Enable Opportunistic Locking	Enabled by default. Opportunistic locking can help performance if the current user has exclusive access to a file. Clear the checkbox to disable opportunistic locking.
Enable this Server as the Master Browser	Enabled by default. The SnapServer can maintain the master list of all computers belonging to a specific workgroup. (At least one Master Browser must be active per workgroup.) Check the checkbox if you plan to install this server in a Windows environment and you want this server to be able to serve as the Master Browser for a workgroup. Clear the checkbox to disable this feature.
Allow Root Authentication	<p>Check the Allow Root Authentication checkbox to allow a root login on the selected server.</p> <p>NOTE: The root password is synchronized with the admin password.</p>

Option	Description
Disable NetBIOS over TCP/IP (Active Dir Domain only)	Some administrators may wish to disable NetBIOS over TCP/IP. Select the checkbox to disable NetBIOS; clear the checkbox to leave NetBIOS enabled. NOTE: If you disable NetBIOS and you are joining a domain, you must enter the domain name as a fully qualified domain name (for example, actdirdomainname.companyname.com). A short form such as ActDirDomName will not work.
Enable Trusted Domains (Active Dir Domain only)	SnapServers recognize trust relationships established between the domain to which the SnapServer is joined and other domains in a Windows environment by default. Select the checkbox to toggle this feature. NOTE: SnapServers remember trusted domains. That is, if this feature is disabled and then activated at a later time, the previously downloaded user and group lists, as well as any security permissions assigned to them, will be retained.

4. Click **OK** to update Windows network settings immediately.

Apple Networking (AFP)

Apple File Protocol (AFP) settings are configured on the **Network > Apple/AFP** page of the Web Management Interface.



The default settings provide access to AFP clients over a TCP/IP network. Mac clients connecting over AFP can log in to the server either as local users on the SnapServer or as Active Directory domain users (if the server belongs to a domain). For more granular control over client access for Mac users who do not belong to a recognized Windows domain, create local user accounts.

NOTE: Mac OS X users can also connect to the SnapServer using Windows networking (SMB). See “Connect a Mac OS X Client Using SMB” on page 3-10.

AFP Configuration Guidelines

Consider the following when configuring access for your AFP clients.

Some SnapServer terms may cause confusion for those familiar with Apple terminology:

Term	Definitions
Share	<p>A SnapServer share appears as a Mac volume that can be accessed through the Finder.</p> <p>NOTE: Unlike standard AppleShare servers, SnapServers allow nested shares (folders within folders). As a result, it is possible for some files or directories to appear in more than one share.</p>
Volume	<p>A volume on a SnapServer is a logical partition of a RAID's storage space that contains a filesystem.</p>
Right-click	<p>This document uses the Windows convention in describing keyboard/mouse access to context-sensitive menus. For example, “To rename a group, right-click a group and then choose Rename.”</p> <p>NOTE: Mac users with a single-button should substitute control-click to achieve the same result.</p>

Authenticating Clients Against a Configured Windows Domain

You can authenticate AFP clients against a Windows domain by navigating to **Network > Apple/AFP** and checking the *Authenticate AFP users against Windows domains* box. When domain authentication is enabled, user names will first be authenticated against the Windows domain and then authenticated against the local database. Local and domain users with the same name will connect as the domain user. To force either local or domain authentication, prefix the user name with the name of the domain to authenticate against or the name of the SnapServer. For example:

`mydomain\username` (domain authentication)

`snap12345\username` (local authentication)

Distinguishing Share Names on the Desktop and Finder

By default, the Finder identifies SnapServer shares using only the share name. To display both the share name and the server name, the **Add Server Name To Apple Shared Folder Names** checkbox on the **Network > Apple/AFP** page of the Web Management Interface is enabled by default. This option makes it easier to differentiate between shared folders with the same share name on multiple servers. For example, SHARE1 on SNAP61009 refers to the share named SHARE1 on the SnapServer named SNAP61009.

AFP Procedures

Use these procedures with AFP.

To Edit AFP Access

1. Go to **Network > Apple/AFP**.

2. Edit settings as described in the following table:

Options	Usage
Enable Apple Filing Protocol (AFP)	Check the Enable Apple Filesharing (AFP) checkbox to enable AFP; leave the checkbox blank to disable AFP access.
Add Server Name to Apple Shared Folder Names	Select this option to identify shares to AFP clients using both the server name and share name. Clear the checkbox to display only the share name.
Authenticate AFP users against Windows domains	Select this option to automatically authenticate AFP users against a Windows domain, if configured. NOTE: By default, users are authenticated against the domain first, then against the local database, so if the same user name exists on both the domain and the SnapServer, the domain user will take precedence. To force an AFP client to log in as either user, prefix the user name with either the Windows domain name or the SnapServer's servername. For example: <i>windowsdomain\username</i> or <i>snap12345\username</i>

3. Click **OK** to update network AFP settings immediately.

To Connect from a Mac (Running OS X)

From the Connect to Server window (from **Finder** press **Cmd + K**, or select **Finder > Go > Connect to Server**), enter one of the following:

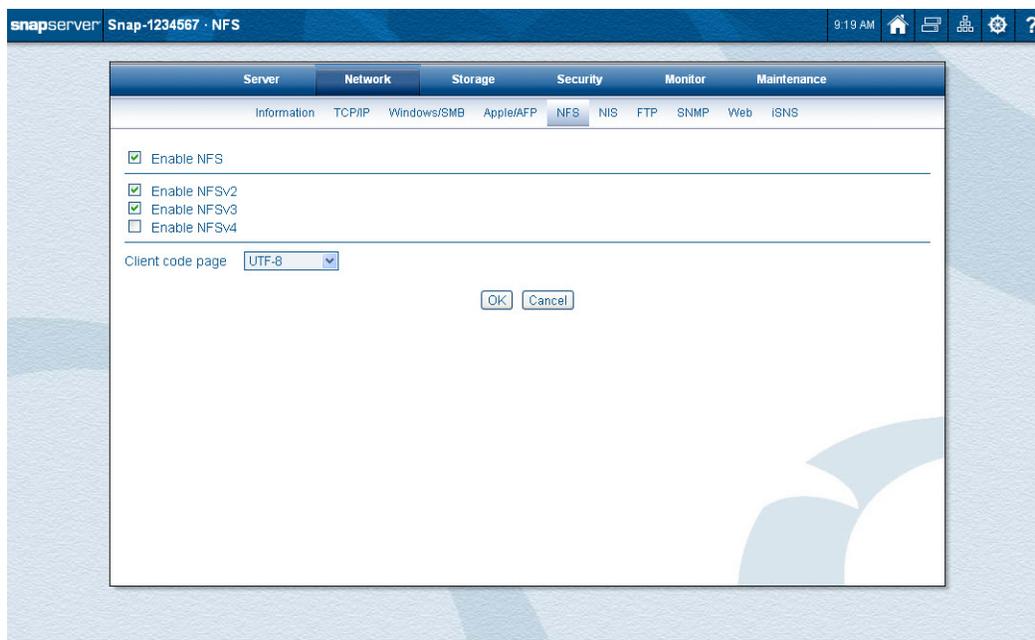
- `afp://servername`
- `afp://ipaddress`

Tip: To disconnect from the SnapServer, drag its icon into the trash.

NFS Access

NFS access to the server is enabled on the **Network > NFS** page of the Web Management Interface. By default, NFS access is enabled and any NFS client can access the SnapServer through the guest account.

NOTE: Only NFSv2 and v3 are enabled by default. If you wish to enable NFSv4, select the Enable NFSv4 checkbox on the Network > NFS page.



NFS client access to shares can be specified by navigating to the **Security > Shares** page and clicking the **NFS Access** link next to the share. Because you are in Unicode mode, you must configure the SnapServer's protocol for the code page being used.

Support for NFS

Consider the following technical information when configuring access for your NFS clients.

Assigning Share Access to NFS Users

The NFS protocol does not support user-level access control, but rather supports host- and subnet-based access control. On a standard UNIX server, this is configured in an “exports” file. On SnapServers, the exports for each share are configured on the NFS Access page independently of user-based share access for other protocols.

Supported Protocols

SnapServers support these versions of the NFS protocol:

Protocol	Version	Source
NFS	2.0, 3.0, 4.0*	RFC 1094, RFC 1813, RFC 3530
Mount	1.0, 2.0, 3.0	RFC 1094 Appendix A, RFC 1813, RFC 3530
Lockd	1.0, 4.0	RFC 1094, RFC1813, RFC 3530

*NFSv4 ACLs are not supported.

NFS Procedures

Use these procedures with NFS.

To Enable NFS Access to the Server

1. Go to **Network > Windows/SMB**.

2. Check the **Enable NFS** box, then check the versions you want to enable (NFSv2, NFSv3, and NFSv4).
Leave the checkbox blank to disable access to this server via the NFS protocol.
3. Choose the desired **Unicode Client** code page from the drop-down list.
4. Click **OK** to update NFS network settings.

To Configure NFSv4 Access to the Server

1. Check the **Enable NFS** and **Enable NFSv4** checkboxes.
A new set of security options are shown below the Enable NFSv4 option.
2. Use this table to select the security you want to apply:

Option	Description
Domain Name	The default setting "localdomain" is shown in the field. If necessary, you can change it.
	 CAUTION: This setting is used by the NFSv4 IDMAP daemon and must be set to the same value on all NFSv4 clients and servers for proper functionality. If set incorrectly, UID and GID resolution will not work properly.
Standard NFS Security	Click this option button if you want to use standard NFS security.
RPSEC GSS Security (UNIX Kerberos)	Click this option button and then complete the fields that appear if you want to use UNIX Kerberos security to authenticate NFSv4 connections. NOTE: Kerberos security can only be configured for UNIX-based Kerberos implementations. Windows ADS Kerberos is not supported for NFSv4 authentication.

3. If you select **UNIX Kerberos** security, complete the options in the table below. Note the following:
- The service will not start unless the TCP/IP domain name (page 3-3) is set up exactly the same as the keytab.
 - You must create the NFS and host service entries in the keytab with the fully qualified domain name of the SnapServer.
 - The SnapServer assumes the domain name from the **primary** Ethernet interface. For more information, see “TCP/IP Options”.

Option	Description
KDC Host Name	Enter the host name of the Kerberos server (for example, <code>kerberos-2000.mit.edu</code>).
Realm Name	Enter the Kerberos realm name (For example, <code>ATHENA.MIT.EDU</code>). NOTE: Realm names are conventionally specified in all CAPITAL letters, but this is not required to function correctly.
Key Tab File	Click Browse to locate and upload the Kerberos key tab file (for example, <code>zeus.keytab</code>) . NOTE: This file can have any name the administrator wishes to give it. If you do not have a keytab file for the SnapServer: <ul style="list-style-type: none"> - create a host and NFS principle for the SnapServer on the KDC - generate a keytab file - save it to a location the client administering the SnapServer can access.

4. Click **OK** to save the configuration.

NOTE: After enabling NFSv4 with Kerberos security, read-write host entries for `gss/krb5`, `gss/krb5i`, and `gss/krb5p` are automatically added to the NFS access entries for each NFS-enabled share.

To Mount Shares from NFS Clients

A share on a SnapServer is equivalent to an exported filesystem on an NFS server. NFS users can mount SnapServer shares and access content directly, or mount a subdirectory of a share, using the following procedure:

1. To mount an NFS client, enter one of the following commands:
 - a. To mount via NFSv2 or NFSv3:


```
mount server_name:/share_name /local_mount
```

 where `server_name` is the name or IP address of the server, `share_name` is the name of the share you want to mount, and `local_mount` is the name of the mount target directory.
 - b. To mount via NFSv4 with standard NFS host-based security, modify the above command as follows:


```
mount -t nfs4 server_name:/share_name /local_mount
```
 - c. To mount via NFSv4 with Kerberos security, enter one of the following commands, depending on the level of security desired:
 - Authentication only (client is authenticated via Kerberos):


```
mount -t nfs4 -o sec=krb5 server_name:/share_name /local_mount
```

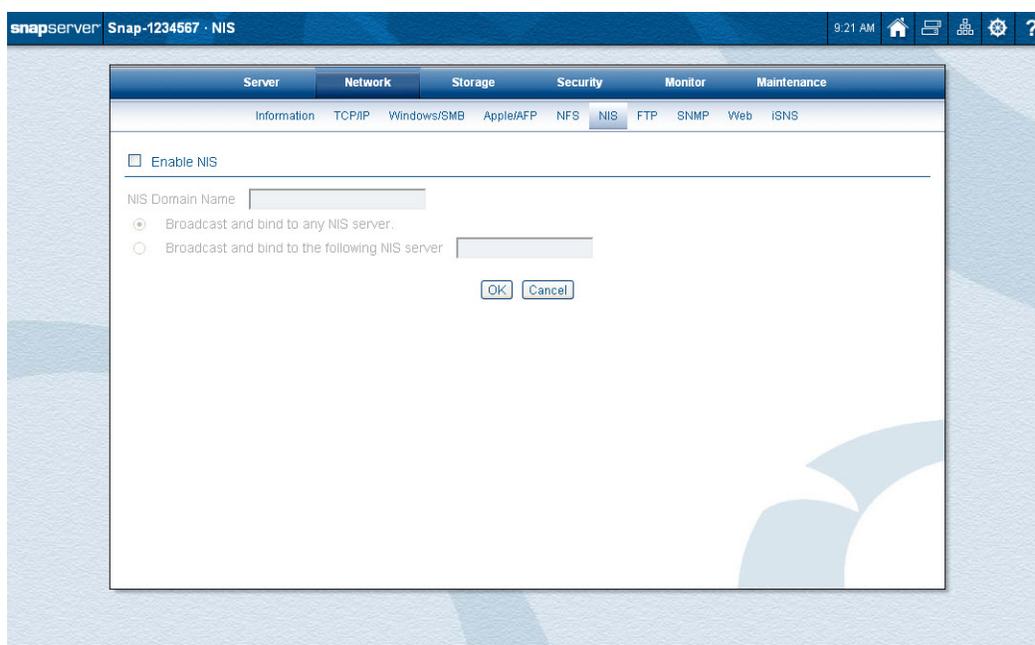
- Authentication and integrity (client is authenticated and data contains a cryptographic checksum that guarantees the data has not changed):
`mount -t nfs4 -o sec=krb5i server_name:/share_name /local_mount`
- Authentication, integrity, and privacy (client is authenticated, data integrity is guaranteed, and data itself is encrypted):
`mount -t nfs4 -o sec=krb5p server_name:/share_name /local_mount`

NOTE: Syntax can vary depending upon the operating system.

2. Press **Enter** to connect to the specified share on the server.

NIS Domain

NIS domains are configured on the **Network > NIS** page of the Web Management Interface.



The SnapServer can join an NIS domain and function as an NIS client. It can then read the users and groups maintained by the NIS domain. Thus, you must use the NIS server to make modifications. Changes you make on the NIS server do not immediately appear on the SnapServer; it may take up to 10 minutes for changes to be replicated.

Guidelines for Configuring NIS

Unless UID/GID assignments are properly handled, NIS users and groups may fail to display properly. For guidelines on integrating compatible SnapServer UIDs, see [“User and Group ID Assignments”](#) on page 7-3.

NIS identifies users by UID, not user name, and although it is possible to have duplicate user names, Overland Storage does not support this configuration.

To Join an NIS Domain

1. Go to **Network > NIS**.
2. Edit the **settings** shown in the following table:

Options	Description
Enable NIS	Check the Enable NIS checkbox to enable NIS, leave the checkbox blank to disable NIS.
NIS Domain Name	Enter the NIS domain name.
NIS Server	To bind to an NIS server, select either: <ul style="list-style-type: none"> • Broadcast and Bind to Any NIS server to bind to any available NIS servers. • Broadcast and Bind to the following NIS server enter the NIS server IP address in the field provided.

3. Click **OK** to update the settings immediately.

FTP/FTPS Access

FTP and FTPS settings are configured on the **Network > FTP** page of the Web Management Interface. FTPS adds encryption to FTP for increased security. By default, FTP and FTPS clients can access the server using the anonymous user account, which is mapped to the SnapServer's *guest* user account and *AllUsers* group account. You can set share access and file access for anonymous FTP users by modifying permissions for these accounts. For more granular control over FTP access, you must create local user accounts for FTP users.

SnapServer also supports explicit FTPS (such as, FTPES or Auth TLS).

NOTE: If standard FTP is enabled, only the data channel is encrypted for FTPS connections – the control channel (including user password) is not encrypted. To force FTPS to encrypt the control channel as well, disable standard FTP.

Supported FTP Clients

SnapServers have been tested with the most common FTP clients and work as expected based on the commands required by RFC 959. SnapServers have been proven to work with these products for standard FTP: Internet Explorer 6.0 and later, Safari 2.0 and later, and Firefox 2.0 and later, and Chrome 1.0 and later.

NOTE: Most standard FTP clients do not support FTPS. A client designed to support FTPS is required for FTPS connections.

To Configure FTP/FTPS Access

1. Go to **Network > FTP**.
2. Edit the **settings** shown in the following table:

Option	Settings
Enable FTP	Check the Enable FTP checkbox to enable standard FTP services; leave the checkbox blank to disable access to this server via standard FTP.

Option	Settings
Enable FTPS	Check the Enable FTPS checkbox to enable FTPS services; leave the checkbox blank to disable access to this server via FTPS.
Allow Anonymous User Access	<p>When you allow anonymous login, FTP/FTPS users employ an email address as the password. When you disallow anonymous login, only FTP/FTPS users who are configured as local SnapServer users can access the server. Select one of the following access options:</p> <ul style="list-style-type: none"> • <i>Checking the checkbox</i> allows users to connect to the server using the anonymous user account. The anonymous user is mapped to the SnapServer's local guest user account. You can set share access for anonymous FTP/FTPS users by granting either read-write (the default access) or read-only access to the guest account on a share-by-share basis. • <i>Leaving the checkbox blank</i> means users cannot log in anonymously but must instead log in via a locally created user name and password.

3. Click **OK** to update the settings immediately.

To Connect via FTP/FTPS

1. To connect to the server through standard FTP, enter the server's name or IP address in the FTP Location or Address box of a web browser or FTP client application.
 - To connect via a command line, enter:
`ftp server_name`
 - To connect via a Web browser, enter:
`ftp://server_name`
(where *server_name* is the name or IP address of the server)
2. To connect to the server through FTPS:
 - Configure your FTPS client application to use explicit FTPS (such as, FTPES or "Auth TLS").
 - Enter the SnapServer's server name or IP address.

NOTE: With anonymous login enabled, access to folders is determined by the share access settings for the guest account. With anonymous login disabled, log into the server using a valid local user name and password.

3. Press **Enter** to connect to the FTP root directory.
All shares and subdirectories appear as folders.

NOTE: FTP users cannot manage files or folders in the FTP root directory.

SNMP Configuration

The SnapServer can act as an SNMP agent. SNMP managers collect data from agents and generate statistics and other monitoring information for administrators. Agents respond to managers and may also send traps, which are alerts that indicate error conditions. The server communicates with SNMP managers in the same community. A community name is a password that authorizes managers and agents to interact. The server only responds to managers that belong to the same public or private community.

Default Traps

A *trap* is a signal from the SnapServer informing an SNMP manager program that an event has occurred. The SnapServer supports the following default traps:

Trap	Initiating Action
coldStart	Whenever SNMP is enabled and the server boots.
linkDown	An Ethernet interface has gone offline.
linkUp	An Ethernet interface has come online.
authenticationFailure	An attempt to query the SNMP agent using an incorrect public or private community string was made, and resulted in a failure.
enterpriseSpecific	<p>SnapServer-generated traps that correspond to the error-level, warning-level, and fatal-error-level traps of GuardianOS. These traps contain a descriptive message that helps to diagnose a problem using the following OID's:</p> <ul style="list-style-type: none"> 1.3.6.1.4.1.6411.2000.1000.1:loglevel 0 syslog messages ("emergency") 1.3.6.1.4.1.6411.2000.1001.1:loglevel 1 syslog messages ("alert") 1.3.6.1.4.1.6411.2000.1002.1:loglevel 2 syslog messages ("critical") 1.3.6.1.4.1.6411.2000.1003.1:loglevel 3 syslog messages ("error")

NOTE: There is no Snap-specific MIB that defines traps sent by SnapServers.

Supported Network Manager Applications and MIBs

SnapServers respond to requests for information in MIB-II (RFC 1213) and the Host Resources MIB (RFC 2790 or 1514). You can use any network manager application that adheres to the SNMP V2 protocol with the SnapServer. The following products have been successfully tested with SnapServers: CA Unicenter TNg, HP Open View, and Tivoli NetView.

Configure SNMP

Edit settings as described in the following table, and then click **OK**. Once enabled, SNMP managers can access MIB-II and Host Resources MIBs management data on the server.

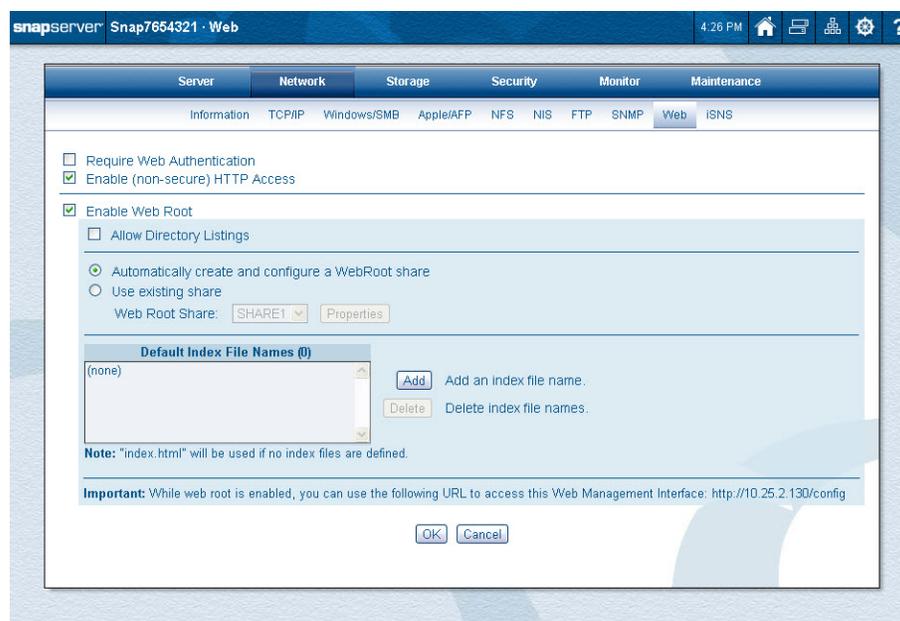
Option	Description
Enable SNMP	To enable SNMP, check the Enable SNMP checkbox. Leave the check box blank to disable SNMP.
Public Community	To enable SNMP managers to read data from this server, enter the name of one or more public communities, or accept the default <i>snap_public</i> .
Private Community	To enable SNMP managers to remotely configure this server, enter the name of one or more private communities, or accept the default <i>snap_private</i> . Create unique public and private names. As a precaution against unauthorized access, Overland Storage recommends that you create your own public and private community names.

Option	Description
Server Location	Enter information that helps a user identify the physical location of the server. For example, you might include a street address for a small business, a room location such as <i>Floor 37, Room 308</i> , or a position in a rack, such as <i>rack slot 12</i> .
Contact Person	Enter information that helps a user report problems with the server. For example, you might include the name and title of the system administrator, a telephone number, pager number, or email address.
Enable SNMP Traps	Check the Enable SNMP Traps check box to enable traps. Clear the check box to disable SNMP traps.
IP Address 1-4	Enter the IP address of at least one SNMP manager in the first field as a trap destination. You can enter up to three additional IP addresses.
Send a Test Trap	To verify your settings, check the Send a test trap box, then click OK.

Web Access

HTTP and HTTPS are used for browser-based access to the server via Web View, Web Root, or the Web Management Interface. HTTPS enhances security by encrypting communications between client and server, and cannot be disabled. You can, however, disable HTTP access on the **Network > Web** page of the Web Management Interface. Additionally, you can require browser-based clients to authenticate to the server.

NOTE: To access the CA Antivirus configuration interface (on the SnapExtensions page), HTTP must be enabled.



GuardianOS supports the following browsers: Internet Explorer 7 and higher, Firefox 3 and higher, Apple Safari 5, and Google Chrome 9 and higher.

Configuring HTTP/HTTPS

You can require web authentication, disable HTTP (non-secure) access, and enable the Web Root feature.

To Require Web Authentication

Edit the following option and click **OK**.

Option	Description
Require Web Authentication	Check the Require Web Authentication checkbox to require clients to enter a valid user name and password in order to access the server via HTTP/HTTPS. Leave the checkbox blank to allow all HTTP/HTTPS clients access to the server without authentication. NOTE: This option applies to both Web View and Web Root modes.

To Enable HTTP Access to the Server

Edit the following option and click **OK**.

Option	Description
Enable (non-secure) HTTP Access	Check the Enable HTTP Access checkbox to enable non-secure HTTP access. Leave the checkbox blank to disable access to the server via HTTP. NOTE: This option applies to both Web View and Web Root modes. To access the CA Antivirus configuration interface, HTTP must be enabled.

To Connect via HTTPS or HTTP

1. Enter the **server name or IP address** in a Web browser.

Web access is case-sensitive. Capitalization must match exactly for a Web user to gain access. To access a specific share directly, Internet users can append the full path to the SnapServer name or URL, as shown in the following examples:

```
https://SNAP61009/Share1/my_files
https://10.10.5.23/Share1/my_files
```

2. Press **Enter**.

The Web View page opens.

Using Web Root to Configure the SnapServer as a Simple Web Server

When you enable the Web Root feature from the **Network > Web** page, you can configure your SnapServer to open automatically to an html page of your choice when a user enters the following in the browser field:

```
http://[servername] or http://[IP address]
```

In addition, files and directories underneath the directory you specify as the Web Root can be accessed by reference relative to `http://[servername]` without having to reference a specific share. For example, if the Web Root points to the directory *WebRoot* on share *SHARE1*, the file *SHARE1/WebRoot/photos/slideshow.html* can be accessed from a web browser:

```
http://[servername]/photos/slideshow.html
```

The Web Root can also be configured to support directory browsing independent of Web View (access through shares).

NOTE: The SnapServer supports direct read-only web access to files. It is not intended for use as an all-purpose Web Server, as it does not support PERL or Java scripting, animations, streaming video, or anything that would require a special application or service running on the server.

Configuring Web Root

1. Complete the following information, then click **OK**.

Option	Description
Enable Web Root	Check the Enable Web Root checkbox to configure the SnapServer to serve the Web Root directory as the top level web access to the server, and optionally, automatically serve an HTML file inside. When the box is checked, the options described below will appear.
Allow Directory Listings	If Allow Directory Listings is checked and no user-defined index pages are configured or present, the browser will open to a page allowing browsing of all directories underneath the Web Root. NOTE: Checking or unchecking this option only affects directory browsing in Web Root. It does not affect access to Web View directory browsing.
Create and configure a Web Root share	Select one of the following: <ul style="list-style-type: none"> • Automatically create and configure a Web Root share: A share named "WebRoot" will automatically be created. By default, the share will be hidden from network browsing and will have all network access protocols except HTTP/HTTPS enabled (as such, it can be accessed from a browser as the Web Root but can not be accessed via Web View). You can change these settings from the Security > Shares page. • Use existing share: Click the Browse button to locate an existing share you want to use as the Web Root share.
Default Index File Names	Files found underneath the Web Root with names matching those in this list will be automatically served to the web browser when present, according to their order in the list. To add a filename, click the Add button, enter the name of one or more index HTML files, then click OK . The file you entered will be shown in the Index Files box. NOTE: If no files are specified, <code>index.html</code> will be automatically loaded if found.

2. Map a drive to the share you have designated as the Web Root share and upload your HTML files to the root of the directory, making sure the file names are listed in the Index Files box.

Accessing the Web Management Interface when Web Root is Enabled

By default, when you connect to a SnapServer with Web Root enabled, the browser will load the user-defined HTML page or present a directory listing of the Web Root. To access the Web Management Interface (for example, to perform administrative functions, change a password, etc.), enter the following in the browser address field:

```
http://[servername or ip address]/config
```

You will be prompted for your User ID and password, then you will be placed into the Web Management Interface.

If you need to access the Web View page to browse shares on the server independent of Web Root, enter this in the browser address:

```
http://[servername or ip address]/sadmin/GetWebHome.event
```

Web View

Web View opens when the user accesses a SnapServer using a Web browser, unless the administrator has enabled the Web Root feature (see [“Using Web Root to Configure the SnapServer as a Simple Web Server”](#) on page 3-25). This page displays a list of all shares to which the user has access. Users can navigate the share structure to locate and view or download files, but they cannot modify or upload files.

For users with admin rights, a key icon () appears next to the file/folder in the share. Clicking this icon displays a popup box with security information about the file/folder.

From this page, the user can also change a password, switch to another user, or log in to perform Administrative functions (if the user has Administrator permissions).

To Switch to a Different User

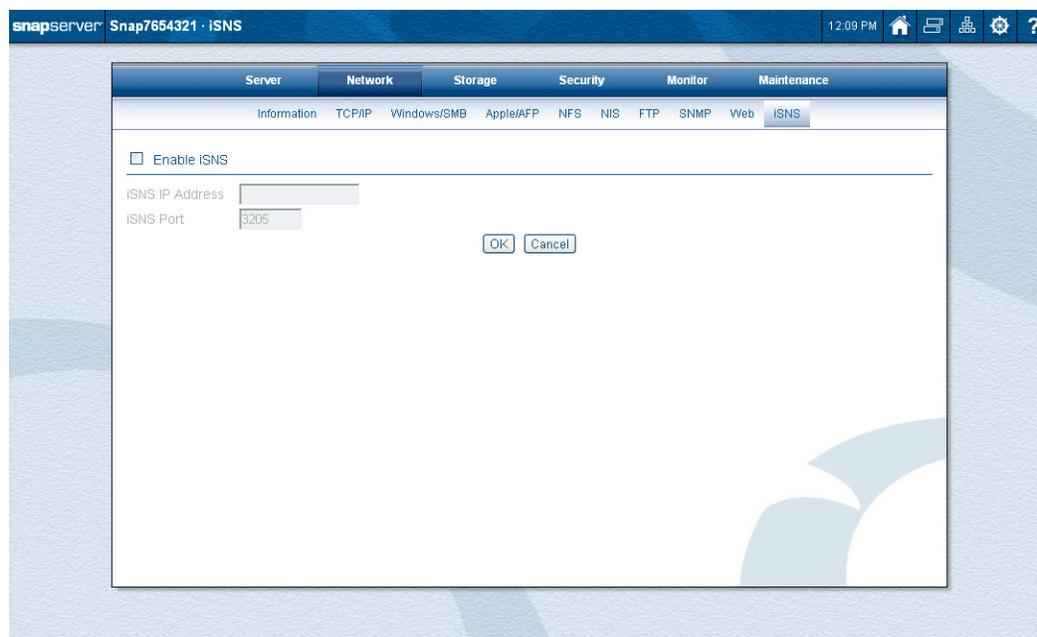
Users can switch to a different user name from the opening Web View page by clicking the **Switch User** link and entering the new user name and password.

To Change a User Password

Users can change their passwords from the opening Web View page by clicking the **Change Password** link, and then completing the user name, old password, and new password information.

iSNS Configuration

Microsoft iSNS Server can be used for the discovery of targets on an iSCSI network. The iSNS software package installs a *readme* file that contains extensive release notes on bug fixes and current iSNS limitations.

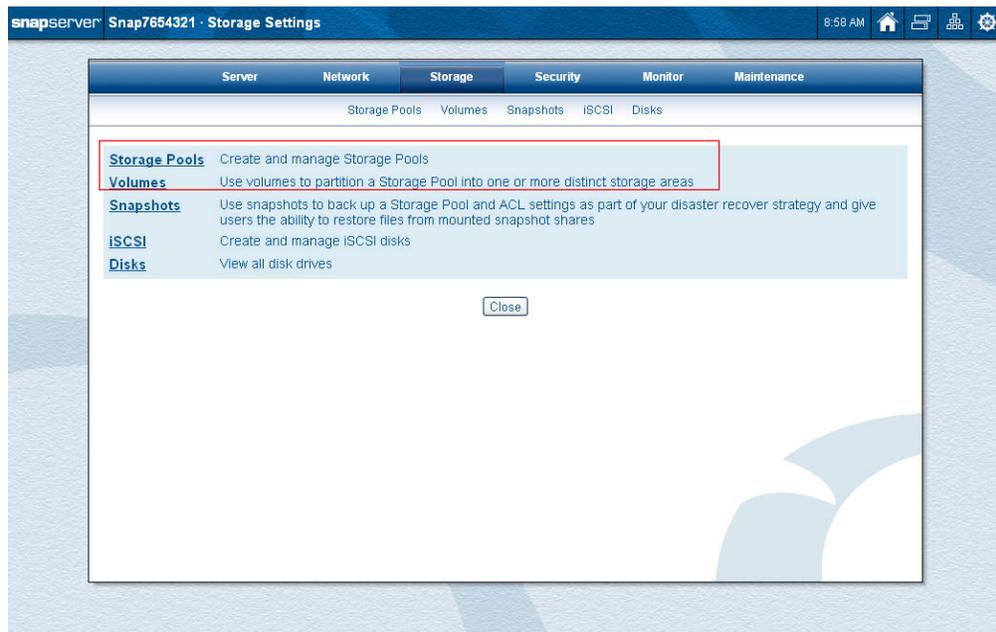


To configure the iSNS settings:

1. Install the iSNS service on a **Windows server**.
Follow the instructions provided in the iSNS *readme* file. Note the IP address of the server or workstation on which the iSNS service is installed.
2. Configure iSNS on the SnapServer.
On the **Network > iSNS** page, check to select the **Enable iSNS** box, enter the IP address of the iSNS workstation, and then click **OK**. The iSNS port default value of 3205 can be changed on this page as well (if changing the port is supported).
3. Configure iSNS in the **iSCSI initiator**.
Run the initiator software and configure the iSNS service from the **iSNS Servers** tab. For example, from a Windows client:
 - When using the Microsoft initiator, run the Microsoft initiator software, select the iSNS Servers tab, and click **Add**. Enter the name or address of the iSNS server, and then click **OK**.
 - When using the QLogic4010/4050 initiator, right-click the QLogic adapter and select **Properties**. Select the Discovery Configuration tab, and check **Perform Discovery**. Check **Use iSNS Server**, enter the server name or IP Address, and click **OK**.

NOTE: After you have completed this procedure, all the iSCSI targets on the SnapServer automatically appear in the Microsoft Initiators target list.

This chapter covers the key options of a DynamicRAID configuration used to manage your SnapServer storage pools and volumes with a maximum of flexibility.



To determine which RAID configuration is appropriate for your needs, see “Should I use DynamicRAID or Traditional RAID?”

For information on the Traditional RAID configuration option, see Chapter 5, “Traditional RAID Storage.” For other storage features, see Chapter 6, “Other Storage Options.”

Topics in DynamicRAID Storage:

- [Storage Pools](#)
- [Volumes](#)

Storage Pools

If you chose the DynamicRAID option during the initial setup of your SnapServer, the wizard creates a separate storage pool on the head unit and on each attached expansion unit. When you navigate to **Storage > Storage Pools**, an overview of the configured storage pool is shown.



IMPORTANT: A SnapServer head unit or expansion unit supports only one storage pool created from its drives and contained within that enclosure. Multiple volumes can be created on that storage pool.

All disk drives in a chassis (head unit or expansion unit) are part of a single storage pool. If disk drives are added to fill empty slots, they become part of the same storage pool. If a new expansion unit is added, a new storage pool is created on it.

Disk drives that have been previously configured can be added to a head or expansion unit and to a chassis and incorporated into its storage pool. These drives are indicated in the list by the  icon and a message stating that the disk has previously been used in a different system. This includes a drive that has any kind of storage configuration on it (from any machine, including the current one) that is not recognized by the server.

Examples include a drive that was installed, a drive that was removed, a drive that was in the server and then added to a RAID, or a drive was part of a RAID that was deleted. This applies to drives that are current RAID members and may have been removed inadvertently. Upon re-insertion, they will not be automatically incorporated, regardless of whether automatic incorporation of unassigned drives is turned on.

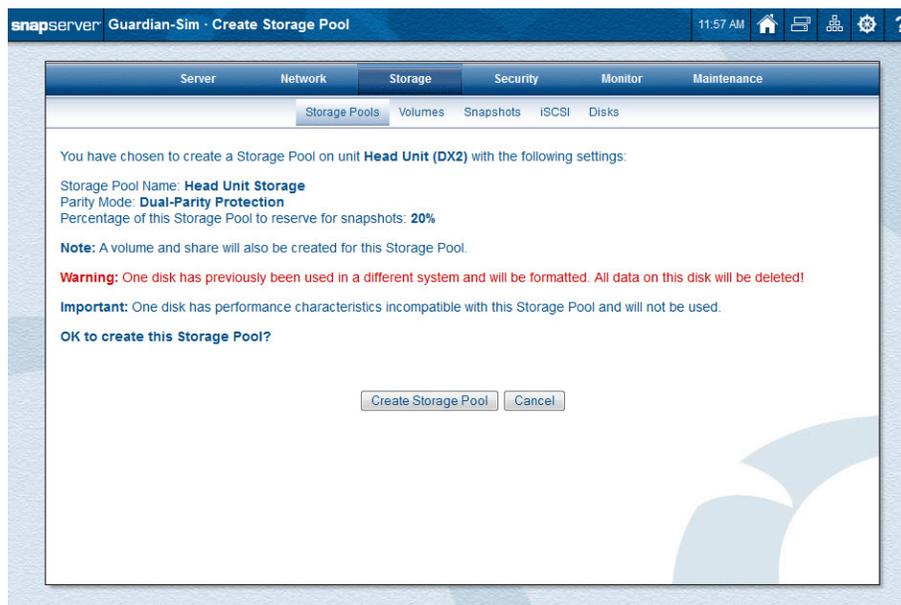
Storage Pool Creation

During the setup process, storage pools are created on the SnapServer using all disk drives available on the head unit and all attached expansion units. DynamicRAID always maximizes the space available based on the type of parity mode and size requested for the snapshot pool. Should it become necessary to create a new storage pool, click the link in the **Status** column of the Storage Pool table to open the **Create Storage Pool** page.

NOTE: You can create storage pools on the head and expansion units with each pool completely contained within the unit, one storage pool per unit. When a storage pool is created and a disk in the storage pool has previously been used in a different system, it will be reformatted and all data on the disk will be deleted.

The first detected drive is used as the basis for the drive-size characteristics of the storage pool. All other drives in the storage pool must conform to the first drive's size characteristics. Otherwise, they are not used. See "Configure Storage – Head Unit" for more details.

The example below demonstrates the notification of both a disk that has previously been used in a different system, and an incompatible drive.



At the **Create Storage Pool** page, you can set these options:

Option	Description
Storage Pool Name	Use this field to enter the name of the storage pool. It can be up to 32 alphanumeric characters and spaces.
Parity Mode	Based on the total number of disks that are available for a storage pool, you can set the parity mode of the storage pool: <ul style="list-style-type: none"> • 1 disk drive – No parity available. • 2 or 3 disk drives – Single-parity protection only. • 4 or more disk drives – Single- or dual-parity protection available. <p>NOTE: Increasing the parity level may require additional disks. You will need to install disks if none are currently available.</p>
Snapshot Pool	Use the drop-down list to choose a percentage of the storage pool that is reserved for snapshots. For more details about snapshots, refer to “Snapshots”. NOTE: Once snapshot space is set up, it can be decreased at any time. To increase the size of the snapshot pool, either the storage pool must be deleted and re-created, or you must add more storage capacity to your storage pool. Default: 20%

To Create a Storage Pool

To create a new storage pool (on a unit that doesn't have a storage pool already present):

1. At the **Storage Pools** page (**Storage > Storage Pools**), click **No Storage Pool** in the **Status** column.
2. At the **Create Storage Pool** page:
 - Select the desired parity mode from the options provided.
 - From the drop-down list, choose the percentage of storage pool space reserved for the snapshot.
3. Click **Create Storage Pool**.
4. At the confirmation page, verify your selections and, if everything is correct, click **Create Storage Pool** again. (To make changes, click **Cancel**.)

If any disk in a storage pool has previously been used in a different system, a warning appears that it will be reformatted and all data on the disk will be deleted.

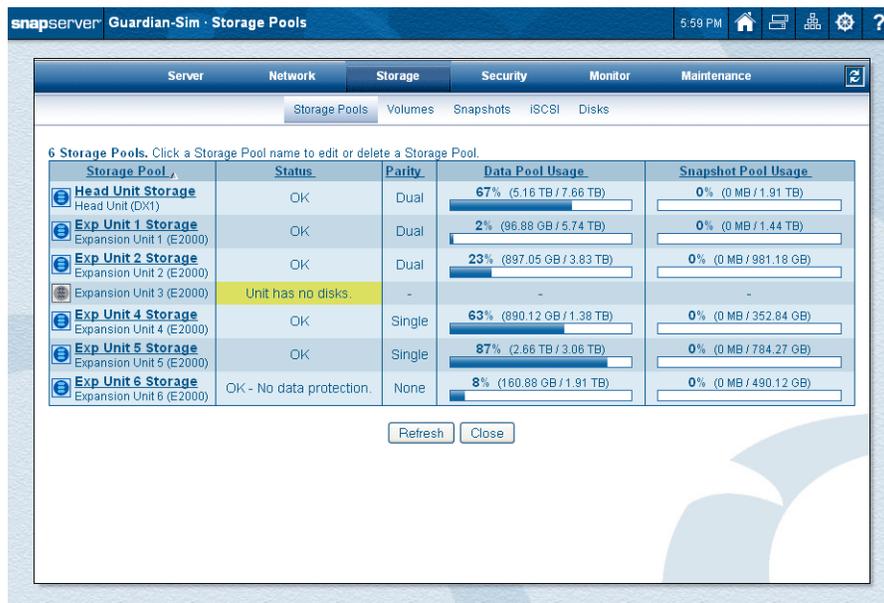
If any disk has size characteristics that are incompatible with the other disks in a storage pool, a message appears that this disk will not be used.

While a storage pool is being created, progress is shown in the **Status** column.
5. When a storage pool has been created, click **OK** to continue.
6. You are returned to the **Storage Pools** page where the **Status** shows **Resync**. Click **Refresh** to see the current **Status** and determine when the resync is complete.

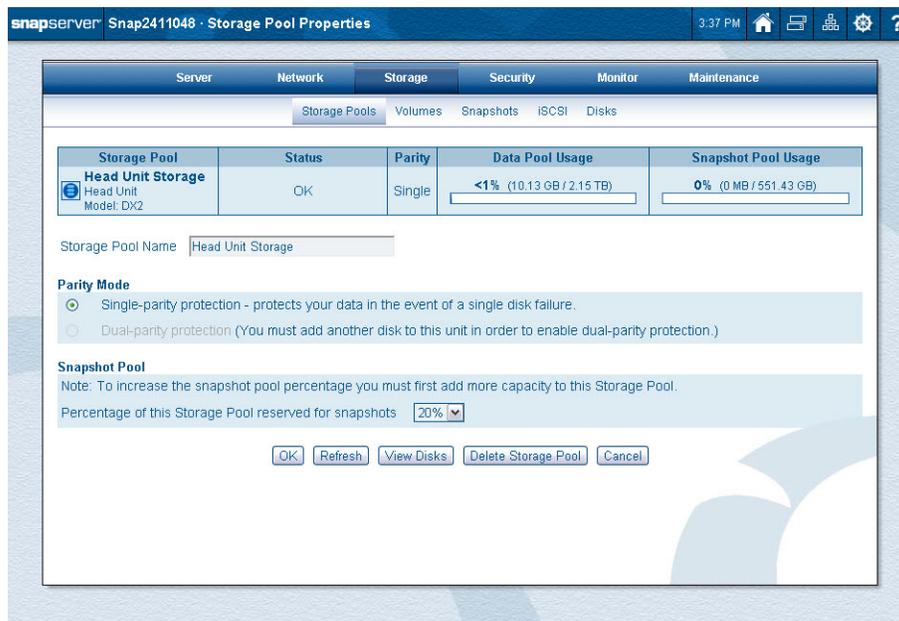
NOTE: The new storage pool is currently being synchronized in the background. Do not apply a heavy load to this storage pool until the synchronization operation is complete. Also, unless you are performing necessary tasks using the Web Management Interface, it is recommended to log out of the interface during the synchronization to give the synchronization operation the full system resources necessary to complete as quickly as possible.

Storage Pool Properties

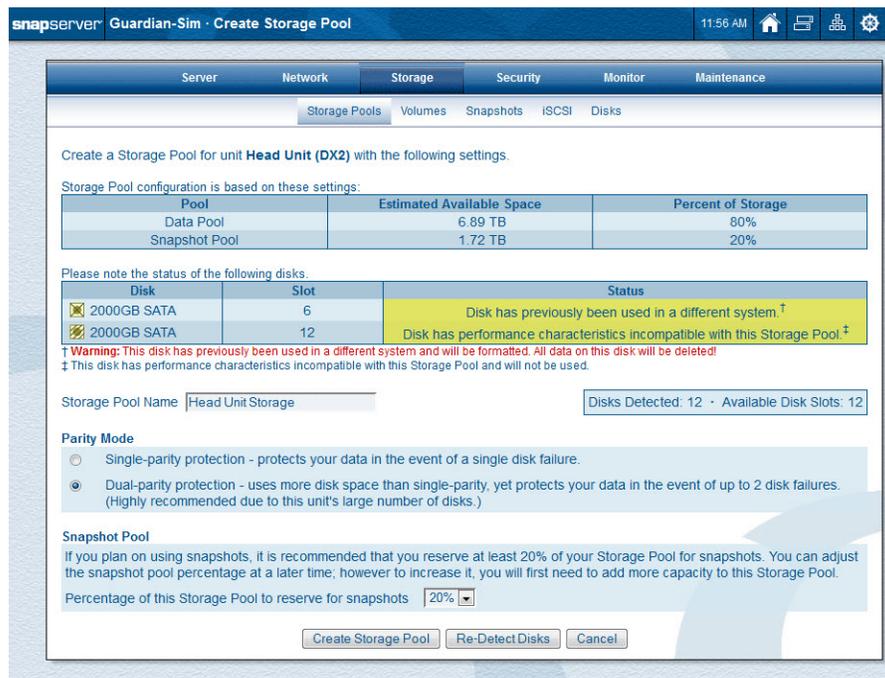
To access the **Storage Pool Properties** page for a storage pool, click the storage pool's name.



When you click the Head Unit Storage name in the Storage Pool list, the properties page for the head unit is shown:



When you click an Expansion Unit Storage name in the Storage Pool list, the properties page for that specific expansion unit is shown:



At the **Storage Pool Properties** page, you can edit these options:

Option	Description
Storage Pool Name	Use this field to change the name of the storage pool. It can be up to 32 alphanumeric characters and spaces.
Parity Mode	<p>You can change the Parity Mode. Your options are based on the current setting and available disk drives. For more information on Parity management, see “Parity Management”.</p> <p>Possible options include:</p> <ul style="list-style-type: none"> • No parity available. • Single-parity protection only. • Single- or dual-parity protection available. <p>NOTE: Increasing the parity level always requires the addition of an unassigned disk to the storage pool. In addition, it may require the installation of additional disks if none are currently available.</p>
Snapshot Pool	<p>Use the drop-down list to choose a percentage of the storage pool that you want reserved for snapshots. You can only decrease the current reserved space from the Properties page.</p> <p>NOTE: If you grow the storage pool by adding a drive and not changing the parity mode, you can allocate the new space to increase snapshot space.</p> <p>For more details about snapshots, refer to “Snapshots”.</p>

If changes are made to the storage pool, a confirmation page is shown. Click **OK** to accept the changes.

View Disks from Storage Pool Properties Page

To view all of the disks in a storage pool, from the **Storage > Storage Pools** page, select a storage pool to open the properties page, then click **View Disks**.

Any disk that is incompatible with the head unit/storage pool is shown with a highlighted message, as demonstrated in the example below.

Storage Pool Disks

Server Network **Storage** Security Monitor Maintenance

Storage Pools Volumes Snapshots iSCSI Disks

Disks for Storage Pool **Head Unit Storage**

Disk	Slot	Status
<input type="checkbox"/> (No Disk.)	1	Slot is empty. (No Disk.)
<input checked="" type="checkbox"/> 932GB SATA	2	OK
<input type="checkbox"/> (No Disk.)	3	Slot is empty. (No Disk.)
<input type="checkbox"/> (No Disk.)	4	Slot is empty. (No Disk.)
<input checked="" type="checkbox"/> 932GB SATA	5	OK
<input checked="" type="checkbox"/> 932GB SATA	6	OK
<input type="checkbox"/> (No Disk.)	7	Slot is empty. (No Disk.)
<input type="checkbox"/> (No Disk.)	8	Slot is empty. (No Disk.)
<input type="checkbox"/> (No Disk.)	9	Slot is empty. (No Disk.)
<input checked="" type="checkbox"/> 932GB SATA	10	OK
<input type="checkbox"/> (No Disk.)	11	Slot is empty. (No Disk.)
<input type="checkbox"/> (No Disk.)	12	Slot is empty. (No Disk.)

Refresh Close

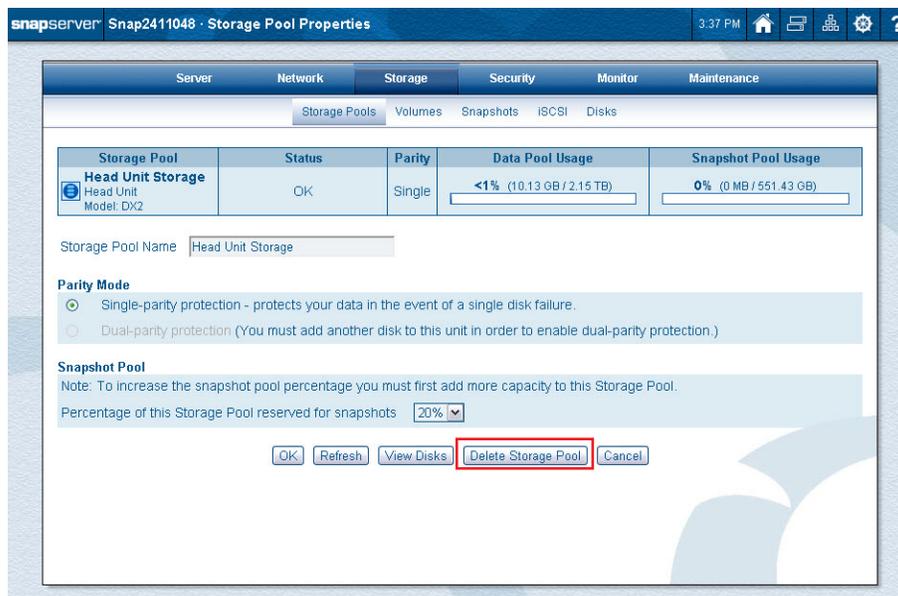
Storage Pool Deletion



CAUTION: Deleting a storage pool deletes all volumes and their data on the storage pool. The data cannot be recovered.

To Delete a Storage Pool

1. Go to the **Storage > Storage Pools** page.



2. Click the **name** of the storage pool being deleted.
3. At the **Storage Pool Properties** page, click **Delete Storage Pool**.
4. At the confirmation page, click **Delete Storage Pool** again.

You are returned to the **Storage Pools** page. The Status for the unit should show **No Storage Pool**. To create a new storage pool, click that link and follow the steps in “[To Create a Storage Pool](#)”.

Parity Management

Parity is set either when creating a new storage pool, or when modifying an existing storage pool to either increase parity by adding a new drive, or decrease parity to expand storage space (and sacrifice redundancy). Parity and snapshot space are selected by the user according to the best estimate of necessary storage requirements.

A move from dual parity to single parity is allowed at any time, provided the storage pool is healthy. A move from single parity to dual parity is only allowed when a new disk drive is added that is large enough to support the new parity mode. See “[Additional Information on DynamicRAID Sizing](#)”

NOTE: A storage pool that was converted from dual parity to single parity cannot be converted back to dual parity until a new disk drive is added. This is due to the extra dual-parity drive that was rolled into the single-parity RAID set.

To Add a Disk Drive to Upgrade Parity

To increase the parity protection of the storage pool, new disk drives can be added to empty slots in the unit containing the storage pool. The DynamicRAID will then obtain user input on how you want to use the new, additional space.

Current Number of Disks	Impact of Adding One More Disk
1	Parity is upgraded from no parity to single parity.
2	Dual-parity option is activated: <ul style="list-style-type: none"> • If dual parity selected, system migrates to it. • If single parity is kept, filesystem space is expanded.
3 or more	The filesystem space is expanded and, if dual parity has been selected to replace single parity, migration commences. See “Adding Drives” .

When a new disk drive is added, the **Administration Home** page displays the message **New Disks Detected**. Going to the **Storage > Storage Pools** page displays a message, **New disks detected (click to use)**. When you click the link, the **Storage Pool Properties** page is shown. You can then change the parity and snapshot settings for the storage pool to take advantage of the additional space.

NOTE: Disk drives that have been previously configured can be added; they are indicated in the list by the  icon and a message stating that the disk has previously been used in a different system.

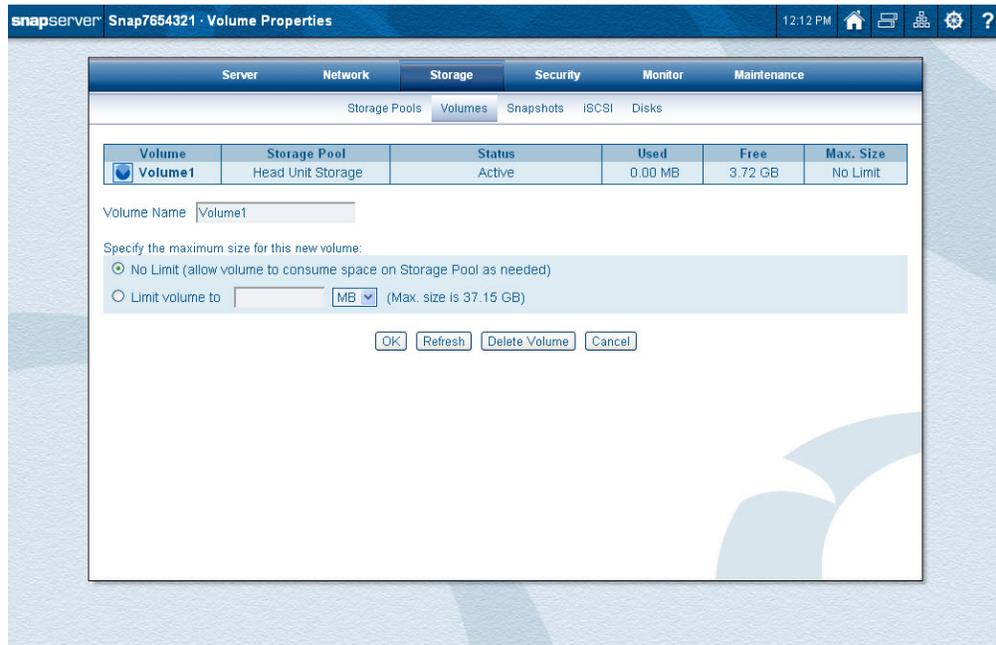
There are no separate spare or global spare disk drives when using the DynamicRAID option. With single parity, if a disk drive fails, a warning is issued and the system reverts to degraded mode with no protection, so that a second drive failure will cause the system to fail. With dual parity, if two disk drives fail, a warning is issued and the system reverts to degraded mode with no protection, so that a third drive failure will cause the system to fail.

Adding Drives. Adding new disks to a storage pool sometimes requires the SnapServer to perform multiple queued operations. During this multi-step resynchronization process, the estimated data pool size will be displayed and may be different than the size currently displayed in the **Data Pool Usage** column. The actual pool size won't be known until the process is complete.

NOTE: New drives added to a storage pool must at least be the same size or larger than the smallest drive in that pool.

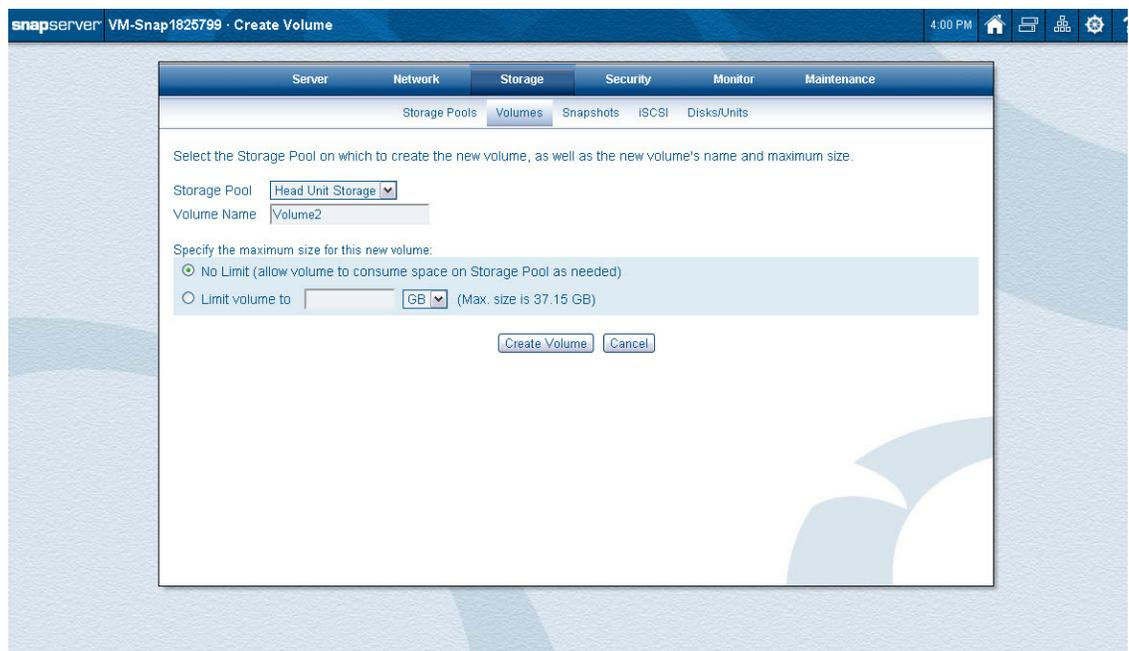
Volumes

GuardianOS supports multiple volumes in a storage pool. During the initial creation of your DynamicRAID storage pool, an initial volume was also created. To view that volume (and create other volumes if needed), navigate to **Storage > Volumes**. To access the **Properties** page for a volume, click the volume name.



Volume Creation

If a storage pool exists, at the **Volumes** page, you can click the **Create Volume** button to set up a new volume.



To Create a New Volume

1. Navigate to **Storage > Volumes**.
2. Click the **Create Volume** button.
3. Choose the **options** for the new volume:
 - Select a **storage pool name** from the drop-down list.
 - Enter a unique **volume name** of 32 alphanumeric characters and spaces.
 - Specify the **maximum size** of the volume:
 - **No Limit** – this allows the volume to expand as needed by making available as much (or as little) of the remaining unused space on the storage pool.
 - **Maximum size** – Establish a maximum volume size limit by entering the amount and selecting a unit of measure. The volume then expands in size as needed until it reaches its maximum. If email notification has been enabled, alerts are sent as the maximum is approached. (To enable email notification, see “[Email Notification](#)”.)

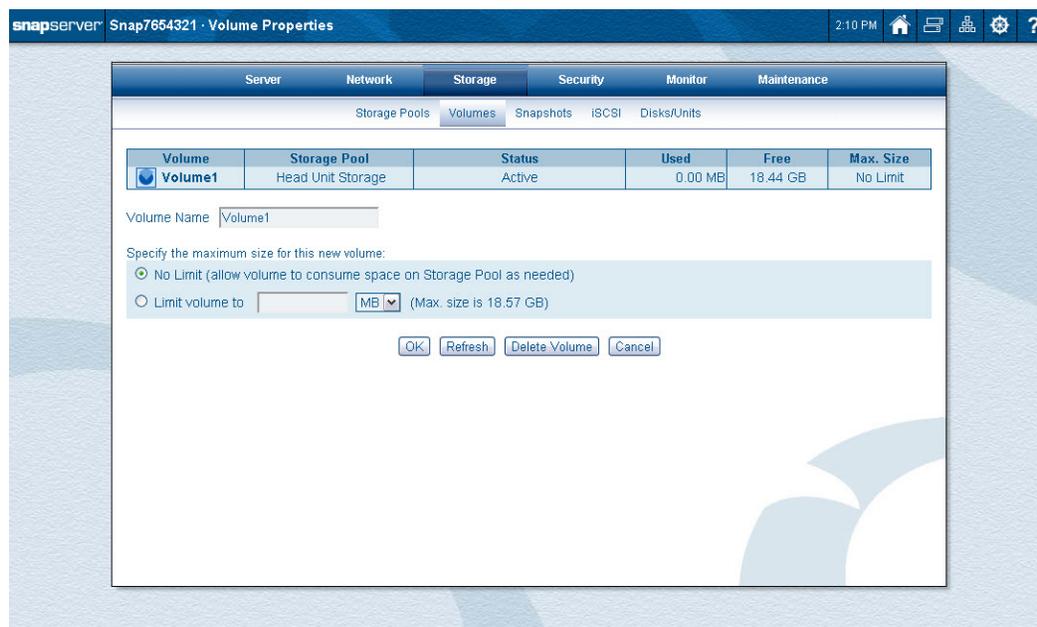
NOTE: If you set the maximum size to less than the current size, the volume is treated as full and no more data can be written to it until the actual space consumed is below the maximum size again.

4. Click the **Create Volume** button on this page to create the volume.

A message appears that the volume has been created. If desired, you can now create a share by clicking the **Create Share** button. See “[Shares](#)” for more information about creating shares.

Volume Properties

At the **Volume Properties** page, you can edit several items.



To Edit Volume Properties

1. Navigate to **Storage > Volumes**.

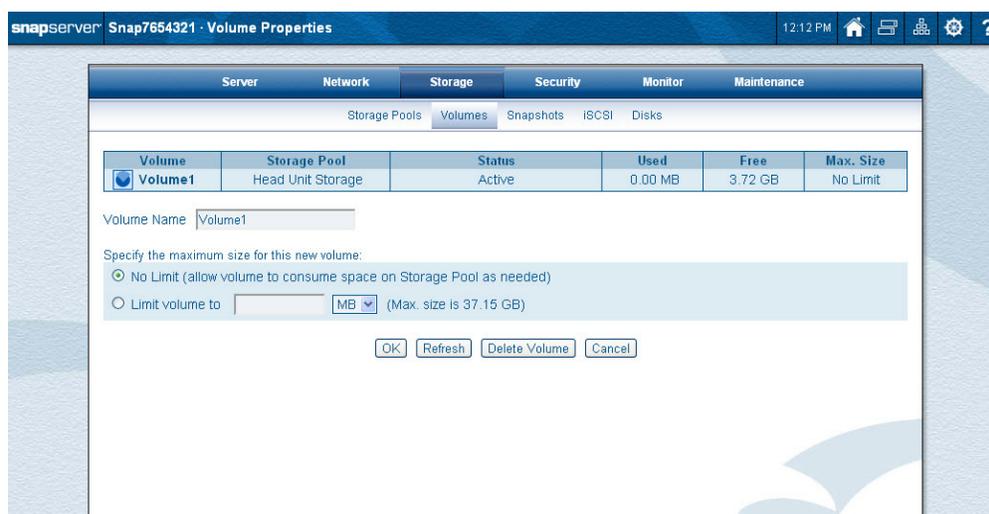
2. Click the **volume name** in the table.
3. At the **Volume Properties** page, change the **options** desired:
 - Edit the **volume name** using up to 32 alphanumeric characters and spaces.
 - Specify the **maximum size** of the volume:
 - **No Limit** – this allows the volume to expand as needed incorporating the remaining unused space on the storage pool.
 - **Maximum size** – Establish a maximum volume size limit by entering the amount and selecting a unit of measure. The volume then grows in size until it reaches its maximum. If email notification has been enabled, alerts are sent as the maximum is approached. (To enable email notification, see “[Email Notification](#)”.)

NOTE: If you set the maximum size to less than the current size, the volume is treated as full and no more data can written to it until the actual space consumed is below the maximum size again.

4. When you are done, click **OK**.

Volume Deletion

To delete a volume, click the **Delete** button on the **Storage > Volumes** page.



To Delete a Volume

1. Navigate to **Storage > Volumes**.
2. Click the **volume name** in the table.
3. At the **Volume Properties** page, click **Delete Volume**.

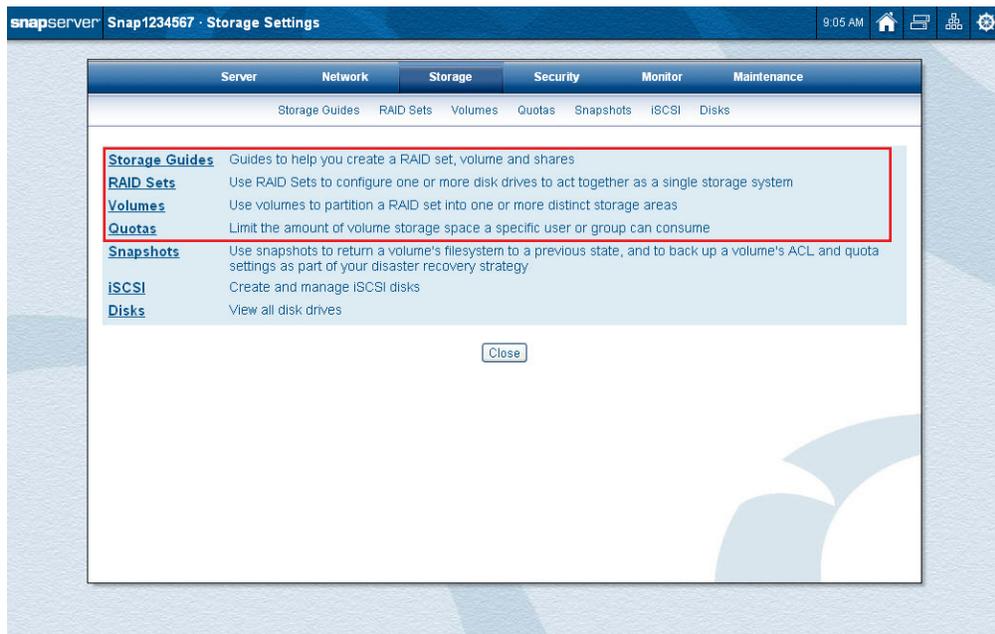


CAUTION: Deleting a volume deletes all data on the volume.

4. At the confirmation page, click **Delete Volume** again.
You are returned to the **Volumes** page. The volume is deleted in the background.

Traditional RAID Storage

This chapter covers the key options of a Traditional RAID configuration. It explains how best to use the Storage Guides and manage your RAID sets, volumes, and quotas.



IMPORTANT: To simplify the management of your SnapServer RAID sets, it is recommended that you use the DynamicRAID option on your server and expansion units.

Using the Traditional RAID option requires you to manually configure and manage RAID sets to meet your specific needs. For simplified storage management and additional configuration options not available in Traditional RAID, use the DynamicRAID option instead. For information on the DynamicRAID configuration option, see [Chapter 4, “DynamicRAID Storage.”](#) For other storage features, see [Chapter 6, “Other Storage Options.”](#)

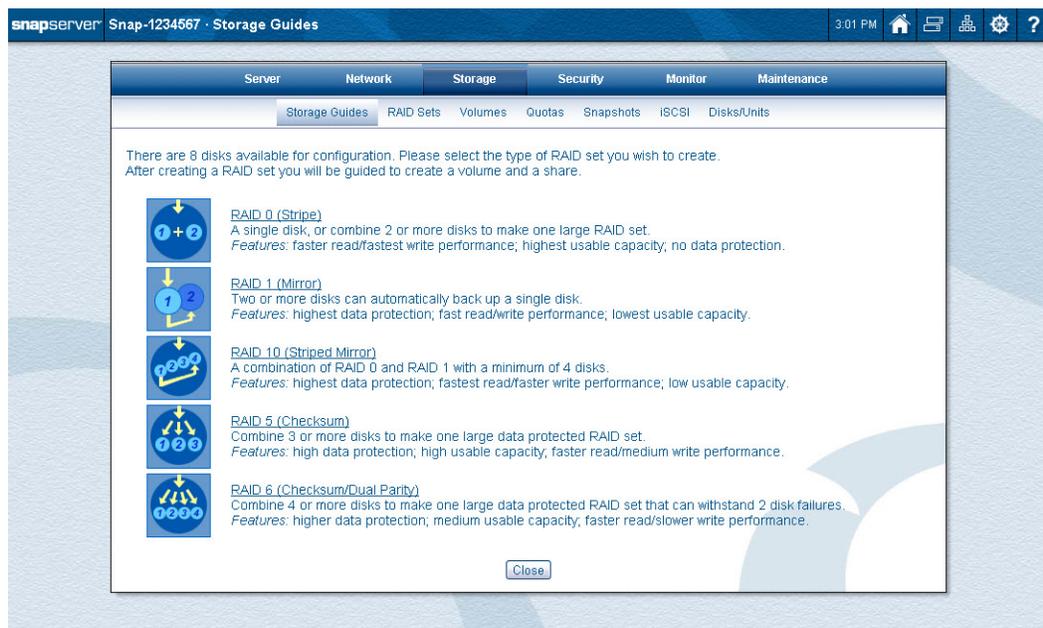
Topics in Traditional RAID Storage

- [Storage Guides](#)
- [RAID Sets](#)
- [Volumes](#)
- [Quotas](#)

Storage Guides

Five different storage guides (wizards) are available for creating a RAID set, volume, and share.

NOTE: If you do not have enough disk drives for the more advanced RAID set configurations, they will be grayed out and unavailable.



The basic steps for storage configuration are:

1. Create a **RAID set**.
2. Create a **volume** on the new RAID set.
3. Create a **share** to access files on the new volume.

Factors in Choosing a RAID Type

The type of RAID configuration you choose depends on a number of factors:

- The importance of the data
- Performance requirements
- Drive utilization
- The number of available drives

For example, in configuring the disk drives of a four-drive SnapServer, the decision whether to include a spare in the RAID depends on the value you place on capacity vs. high availability. If capacity is paramount, you would use all drives for storage; if high availability were more important, you would configure one of the drives as a spare.

The following table summarizes the advantages and disadvantages of each type of RAID.

Features	RAID 0	RAID 1	RAID 5	RAID 6	RAID 10
Data Loss Risk	Highest	Lowest	Low	Lower	Very Low
Write Access Speeds	Fastest	Fast	Medium	Slower	Faster
Usable Capacity	Highest	Lowest	High	Medium	Low
Disks Required	1 or more	2 or more	3 or more	4 or more	4 or more
Supports Spares	No	Yes	Yes	Yes	Yes



CAUTION: To reduce exposure to double-drive disk failures on RAID 5, use no more than eight drives in a single RAID set and group smaller RAID sets together. RAID 6 is recommended for RAID sets with more than four drives.

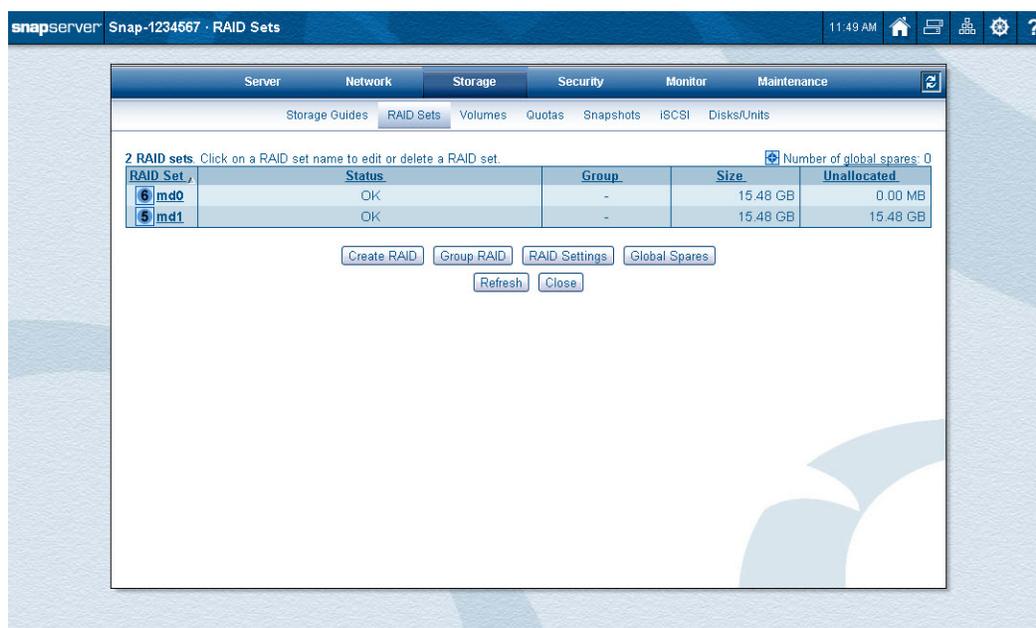
Local and Global Spares

A *spare* is a disk drive that can automatically replace a failed drive in a RAID 1, 5, 6, or 10 set. Designating a disk drive as a spare helps ensure that data is available at all times. If one disk drive in a RAID fails or is not operating properly, the RAID automatically uses the spare to rebuild itself without administrator intervention. SnapServers offer two kinds of spares: local and global.

Item	Description
Definitions	<p>Local (hot) spare – A local (or dedicated) spare is associated with and is available only to a single RAID. Administrators typically create a local spare for RAID sets containing mission-critical data that must always be available.</p> <p>Global (hot) spare – A spare that may be used for any RAID 1, 5, 6, or 10 in the system on any unit (assuming sufficient capacity) as it becomes needed.</p>
Identifying	<p>Spares are identified on the Storage > Disks page using the following icons:</p> <p> Global Spare (GS)</p> <p> Local Spare</p> <p>Each icon will be associated with a disk in the RAID, identifying that disk as either a local spare or a global spare.</p>
Interaction	<p>When a drive in a RAID fails, the system looks for a spare in the following order:</p> <ol style="list-style-type: none"> 1. If a local spare dedicated to the RAID exists, use the local spare. 2. If no local spare is available, and there is a single global spare of sufficient capacity, use the global spare. 3. If no local spare is available, and two global spares of different capacity are available, use the smaller global spare with sufficient capacity.

RAID Sets

Use the **Storage > RAID Sets** page to manage RAID sets and their options.



From the **RAID Sets** main page, you can do the following:

- **Create RAID Sets** – If unassigned drives exist, a new RAID set can be created by launching a wizard.
- **Group RAID Sets** – If more than one RAID set exists, they can be grouped together.
- **Change RAID Settings** – Change two RAID options:
 - Enable/disable automatic incorporation of unused disks into degraded RAID sets.
 - Enable/disable background disk scans during idle I/O system time.
- **Manage Global Spares** – Add, delete, and configure global spares.
- **RAID Set Properties** – Edit the RAID set properties (click the name to access).

Create RAID Sets

If you choose not to use the Storage Guide wizards to expedite the configuration of your RAID sets, you can manually configure them using these steps:

1. Under **Storage > RAID Sets**, click **Create RAID**.

The following page is displayed. Based on the disk drives available, only the supported RAID options have active links. The other options and icons are grayed out.



2. Click the **RAID option** you want.

The following table summarizes the advantages and disadvantages of each type of RAID:

Features	RAID 0	RAID 1	RAID 5	RAID 6	RAID 10
Data Loss Risk	Highest	Lowest	Low	Lower	Very Low
Write Access Speeds	Fastest	Fast	Medium	Slower	Faster
Usable Capacity	Highest	Lowest	High	Medium	Low
Disks Required	1 or more	2 or more	3 or more	4 or more	4 or more
Supports Hot Spares	No	Yes	Yes	Yes	Yes

CAUTION: To reduce exposure to double-drive disk failures on RAID 5, use no more than eight drives in a single RAID set and group smaller RAID sets together. RAID 6 is recommended for RAIDs with more than four drives.

3. Place a check mark next to the **disks** you want to include in the RAID set.

NOTE: Disks can be from the head unit or any attached expansion unit. However, creating a RAID with disks from different units increases the chance of a multiple-disk RAID failure due to communication issues that may arise between units.

CAUTION: Do not mix drives of different capacity in a RAID 1, 5, 6, or 10 set. Because all drives within a RAID must be the same capacity, using mixed-capacity drives in the same RAID will result in wasted capacity. Also, do not mix drives of different rotational speeds in the same slot column. See “Adding Disk Drives” on page 6-28 for illustrations of supported and unsupported drive configurations.

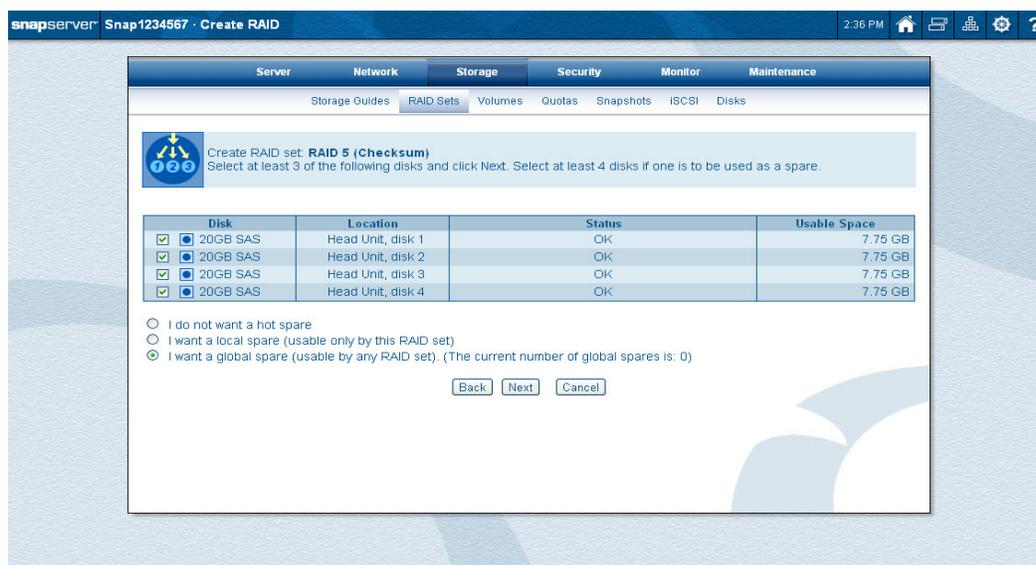
For example, if a RAID is configured with the drives listed in the following table, some capacity of the larger drives will go unused.

Drive	Raw Capacity	Actual Used Capacity	Usage
Drive1	750 GB	750 GB	100%
Drive 2	750 GB	750 GB	100%
Drive 4	1 TB	750 GB	75% of 1 TB
Drive 6	2 TB	750 GB	38% of 2 TB

4. Select an **option** for spares:

- **I do not want a hot spare** – No spare will be created.
- **I want a local spare** – A local spare will be usable only by this RAID set.
- **I want a global spare** – A global spare is usable by any RAID set.

For more information about spares, see “Manage Global Spares”.



5. Click **Next**.

6. Verify your configuration, then click **Next**.



A message appears confirming the creation of the RAID set, and how much storage space it provides. Before you can place any data on this RAID set, you must create a volume. Choose to create the volume now or later:

- To create the volume now, click **Create Volume Now**, configure your volume, click **Create Volume** when you are finished, and click **Create Volume** again to confirm. For specifics on creating volumes, see [“Volume Creation”](#).
- To create the volume later, select **Create Volume Later** and you will be returned to the RAID Sets page. You will need to remember to go to the Volumes option to create the volume at a later time.

You can cancel the process at any time.

Group RAID Sets

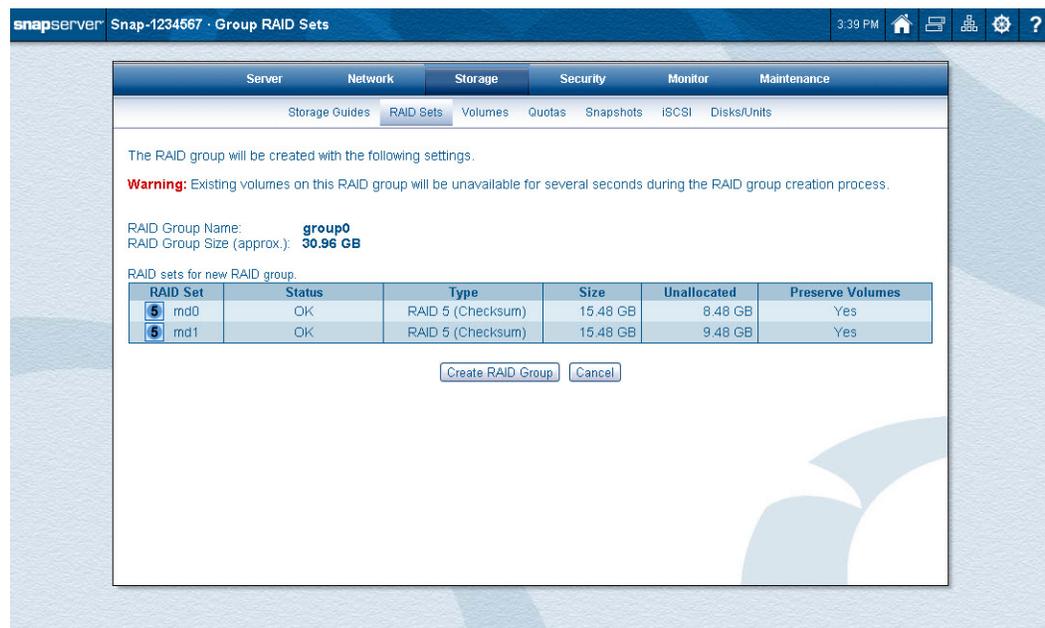
RAIDs can be grouped together to neatly resolve a number of capacity issues. For example, a volume on one RAID nearing full utilization can be expanded using spare capacity on another RAID. The ability to grow volumes beyond the capacity of a single RAID allows administrators to expand a volume without reconfiguring RAID sets and allows users to continue working as usual with no interruption.

Grouped RAID sets must be the same type. For example, you can group two RAID 1 sets or two RAID 5 sets but you cannot group a RAID 1 set and a RAID 5 set.

Click the **Group RAID** button to show the Group RAID Sets page.



1. Select (check) the RAID sets you want to include in the group and click **Next**. You will see a confirmation page.



2. Click the **Create RAID Group** button to complete the process.
3. At the primary **RAID Sets** page, click the group name to see the details of the group.



From this page you can view the status, add another RAID set of the same type to the group, or delete the entire group. The status shows the following information:

Label	Description
<i>Group Table</i>	
RAID Group	The name of the RAID Group to which the RAID belongs
Status	The current condition of the Group: <ul style="list-style-type: none"> • <i>Active</i> – The group and all its RAID sets is functioning properly. • <i>Resync</i> – A device repair operation is in progress. • <i>Failure</i> – The RAID is offline. • <i>Degraded</i> – A drive has failed or been removed.
Type	Type of RAID configured on members of the group.
Size	The total capacity of the group.
Unallocated	The total storage space in the group not allocated to a volume or snapshot pool.
<i>RAID Set Table</i>	
RAID Set	The name of each RAID set. A symbol of the RAID type is shown to the left of the name.
Status	The current condition of the RAID: <ul style="list-style-type: none"> • <i>OK</i> – The RAID is functioning properly. • <i>Resync</i> – A device repair operation is in progress. • <i>Failure</i> – The RAID set is offline. • <i>Degraded</i> – A drive has failed or been removed.
Type	Type of RAID configured on the RAID set.
Size	The total capacity of the RAID set.
Unallocated	The total storage space not allocated to a volume or the RAID set's snapshot pool.

Adding an Expansion Unit

In a common scenario, a SnapServer configured as a RAID 5 is nearing full utilization. The administrator decides to add an expansion unit. The administrator creates a RAID 5 on the expansion unit, groups it with the existing RAID set on the SnapServer, and then expands volume capacity using the new storage from the expansion unit. Clicking the **Expand Volume** button that appears does this automatically (see “[Expand Volume Capacity](#)”).

Expansion unit RAID sets are created in the same way as head unit RAID Sets. See “[Create RAID Sets](#)”.

Grouping RAIDs with other Grouped RAIDs

Just as RAID sets can be grouped, individual RAID groups can be brought together to form an even larger group.

For example, a SnapServer is running out of capacity. Two 12-drive expansion units are attached to the SnapServer to provide increased capacity. You can configure a RAID 5 on each of the expansion units, then group the two of them together. The resulting RAID set group can then be grouped with the RAID set on the SnapServer, allowing network users to take advantage of the full capacity of the head and expansion units with no loss of capacity.

See “[Group RAID Sets](#)”.

Deleting Grouped RAIDs



CAUTION: Deleting a RAID group deletes all the RAID sets, volumes, and shares. Any data on those volumes will be lost.

If one RAID set becomes inaccessible for any reason, the entire RAID group containing that RAID set will also become inaccessible. Depending on the cause, the RAID group may or may not be recoverable.

For example, if a RAID group spans a SnapServer and an expansion unit and one of the RAIDs goes down because of a disconnected cable, the RAID group is fully recoverable by reconnecting the cable and rebooting the system. On the other hand, if one of the RAIDs becomes corrupted and remains unrecoverable, the data in the other RAID will also be lost.

See “[Delete a RAID Set](#)”.

Snapshot Pools are Combined

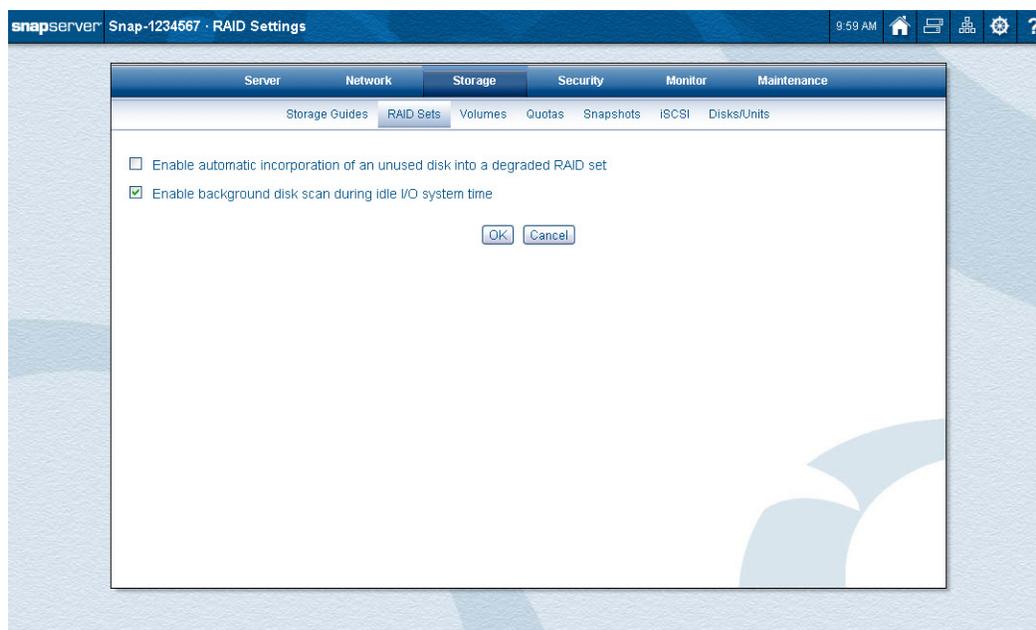
When two RAID sets are grouped, the size of the resulting snapshot pool is the sum of each RAID set's formerly separate snapshot pools.

Two RAIDs at a Time Grouping Rule

To group more than two RAID sets, create a RAID group with two RAID sets, then add each remaining RAID set to the RAID group, one at a time.

Change RAID Settings

Click the **RAID Settings** button on the **RAID Sets** page to allow you to enable or disable the automatic incorporation of a disk into a degraded RAID set or a background scan.



Automatic Incorporation of Hot-Swapped Drives

If a RAID (except RAID 0) is running in degraded mode and a raw drive, a non-GuardianOS drive, or an unassigned GuardianOS-partitioned drive is “hot-inserted” into a SnapServer, it can be automatically assigned as a local spare and used to rebuild the degraded RAID. If there are no degraded RAIDs, a hot-inserted non-GuardianOS or unassigned drive will be automatically configured as a global spare. To enable automatic incorporation of unassigned drives, go to the **Storage > RAID Sets** page and click the **RAID Settings** button.

NOTE: Drives that have previously been configured for use in a different RAID set on any SnapServer are not automatically incorporated, regardless of whether automatic incorporation of unassigned drives is turned on. You must manually incorporate and configure these previously used drives.

Background Disk Scan

The background disk scan checks the integrity of RAID data by continuously scanning the disk drives for errors. Each RAID (except RAID 0) has its own background disk scan that is set to run when disk I/O drops to a low level of activity. Once the activity rises above the *idle threshold*, the background scan stops and waits for the activity to fall to the idle threshold again before resuming. As a result, there should be minimal to no impact on performance. Once the disk scan has completed a pass on a given RAID set, it waits a designated period of time before starting again.

The background disk scan is enabled by default. To disable the background disk scan, go to the **Storage > RAID Sets** page and click the **RAID Settings** button. Note the following:

- If the background disk scan is disabled, the SnapServer will still initiate a scan on a RAID if problems are detected on one of the RAID drives.
- The background scan will not run on RAIDs that are degraded, syncing, or rebuilding.

Manage Global Spares

A **spare** is a disk drive that can automatically replace a damaged drive in a RAID 1, 5, 6, or 10. Designating a disk drive as a spare helps ensure that data is available at all times. If one disk drive in a RAID fails or is not operating properly, the RAID automatically uses the spare to rebuild itself without administrator intervention. SnapServers offer two kinds of spares: local and global (see “[Local and Global Spares](#)”).

Click the **Global Spares** button to view the **Global Spares** page showing all the disks available for use, or that are in use, as global spares. To enable a disk as a global spare, check the checkbox next to the desired disk and click **OK**. More than one disk can be checked at a time. To disable or delete a disk assigned as a global spare, clear the checkbox next to the disk and click **OK**.

The following disks are available for use as (or are currently assigned as) global spares. Global spares are automatically used to replace failed members of RAID sets. Check the disks you want to use as global spares. Uncheck the disks you do not want to use as global spares.

Legend: Disk available for use as global spare Disk currently assigned as global spare

Disk	Location	Status	Usable Space
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 6	OK	7.74 GB
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 7	OK	7.74 GB
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 8	OK	7.74 GB

OK Cancel

RAID Set Properties

By clicking a RAID set name on the **RAID Sets** main page, details of that particular RAID set are shown on a **RAID Set Properties** page.

RAID Set	Status	Group	Size	Unallocated
md0	Resync: 40% complete Members (active / configured): 4 / 4 Spares (active / configured): 1 / 1	-	15.48 GB	0.00 MB

Disks for RAID set md0. (*Note: Usable Space specifies the amount of disk space usable by the RAID set.)

Disk	Location	Status	Usable Space*	Action
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 1	OK	7.74 GB	-
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 2	OK	7.74 GB	-
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 3	OK	7.74 GB	-
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 4	OK	7.74 GB	-
<input checked="" type="checkbox"/> 20GB SAS	Head_Unit_disk 5	Local Spare - OK	7.74 GB	Remove

Number of global spares: 0

Refresh Delete RAID Add Disk Close

The following table shows details about member drives of that specific RAID:

Label	Description
RAID Set	The name of each RAID.
Status	<p>The current condition of the RAID:</p> <ul style="list-style-type: none"> • <i>OK</i> – The RAID is functioning properly. • <i>OK-Spare Missing</i> – The RAID is functioning properly after a repair and rebuild. Because the local spare was consumed to repair the RAID, it is no longer available as a spare. <p>It is recommended that the original drive that failed be replaced to restore the RAID to its proper configuration and provide the full protection by one or more local spares. Alternately, you can click the link to reset the RAID spare count; however, the RAID will not be able to automatically recover from a drive failure.</p> <ul style="list-style-type: none"> • <i>Resync</i> – A device repair operation is in progress. • <i>Failed</i> – The RAID is offline. • <i>Degraded</i> – A drive has failed or been removed. <p>Number of members in the RAID:</p> <ul style="list-style-type: none"> • <i>Active</i> – Number of non-spare disks in the RAID that have a status of OK. • <i>Configured</i> – Number of non-spare disks with which the RAID was configured.
Group	The name of the RAID Group to which the RAID belongs.
Size	The total capacity of the RAID.
Unallocated	The total storage space not allocated to a volume.



CAUTION: Actions on this page can result in a loss of data. Be sure you have backed up your data before making changes to RAID sets.

From this secondary page, you can:

- Remove an individual RAID disk drive or local spare.
- Add a disk drive.
- Delete the entire RAID set (if not part of a group).

Remove a RAID Drive

From the **RAID Set Properties** page, you can remove a RAID disk drive or local spare by clicking the **Action** link on the far right of disk table. If you are removing a primary RAID disk, you will see a message page warning of RAID running in a degraded mode (no or reduced parity).

NOTE: The only types of drives that can be removed are local spares, failed drives, or members of a RAID 1, 5, 6, or 10.

Add a Disk Drive as Local Spare to RAID

Clicking the **Add Disk** button at the bottom of the page displays a table of available disk drives. Check one or more of the boxes to add them to the RAID set and click **Next** for the confirmation page.

Adding disks is limited based on the type of RAID it is being made a member of:

- Disks cannot be added to a RAID 0.
- Disks can only be added to a RAID 1 as full members.
- Disks can only be added to all other RAID types (5, 6, or 10) as local spares.

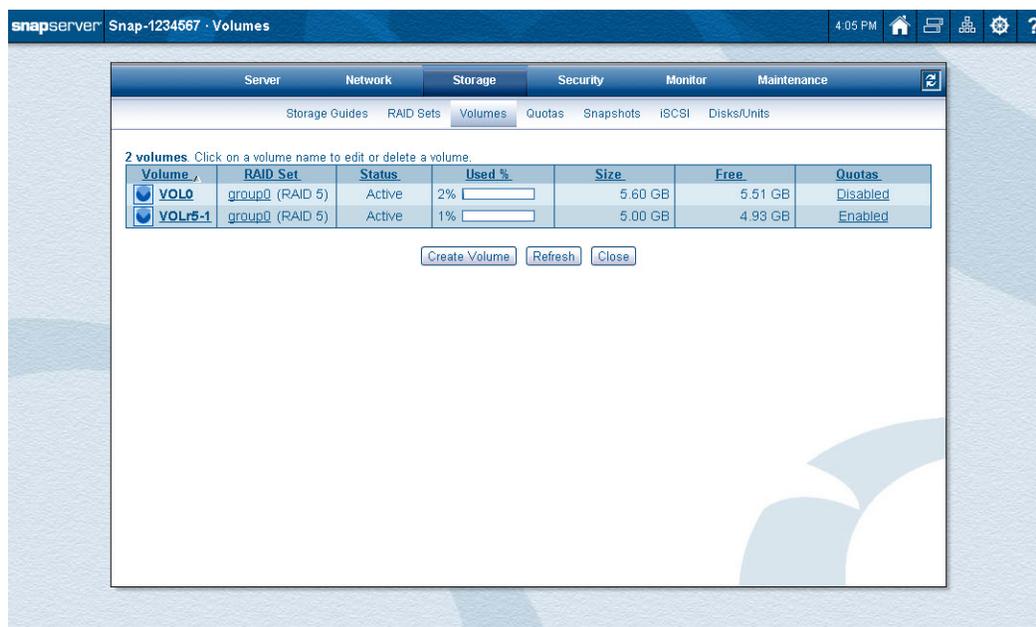
NOTE: Disk drives that have been previously configured can be added; they are indicated in the **Storage > Disks** list by the  icon and a message stating that the disk has previously been used in a different system. If you want to use the drive, add it to the RAID as you would any other drive.

Delete a RAID Set

Click the **Delete RAID** button at the bottom to completely delete the RAID set. This also deletes the Volume and Shares, including any data on them. Click the **Delete RAID** button again to complete the deletion.

Volumes

Use the **Storage > Volumes** page to manage the volumes that have been created on the RAID set.



From this page, you can:

- Create a new volume.
- Edit or delete the volume (click the name).
- Enable/disable quotas on the volume (click the **Quotas** link on far right).

Volumes and the Snapshot Pool

The default capacity settings for the filesystem and future snapshot use are 80% for the filesystem and the remaining 20% for snapshots. You may need to adjust this figure depending on your snapshot strategy or expand the volume to all available space if you plan

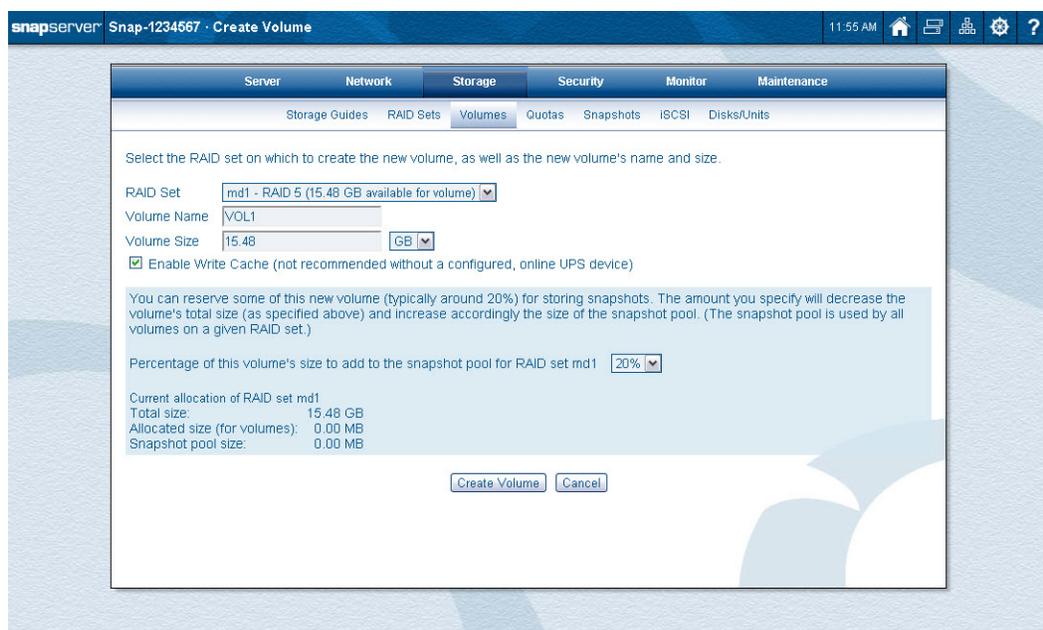
never to use snapshots. Keep in mind that you can increase or decrease snapshot pool size at any time, but volume space can only be increased. For more information, see “[Estimating Snapshot Space Requirements](#)”.

NOTE: GuardianOS snapshots should not be used on volumes that contain iSCSI disks. If a volume will contain one or more iSCSI disks, decrease the Snapshot pool size to zero. For information about creating snapshots of iSCSI disks, see “[Configuring VSS/VDS for iSCSI Disks](#)” on page 6-22.

Volume Creation

When you click the **Create Volume** button (or the **Create Volume Now** or **Create Volume Later** buttons when creating a RAID set), the **Create Volume** page is displayed. Make your selections and click the **Create Volume** button.

For more information, refer to “[Volume Properties](#)”.

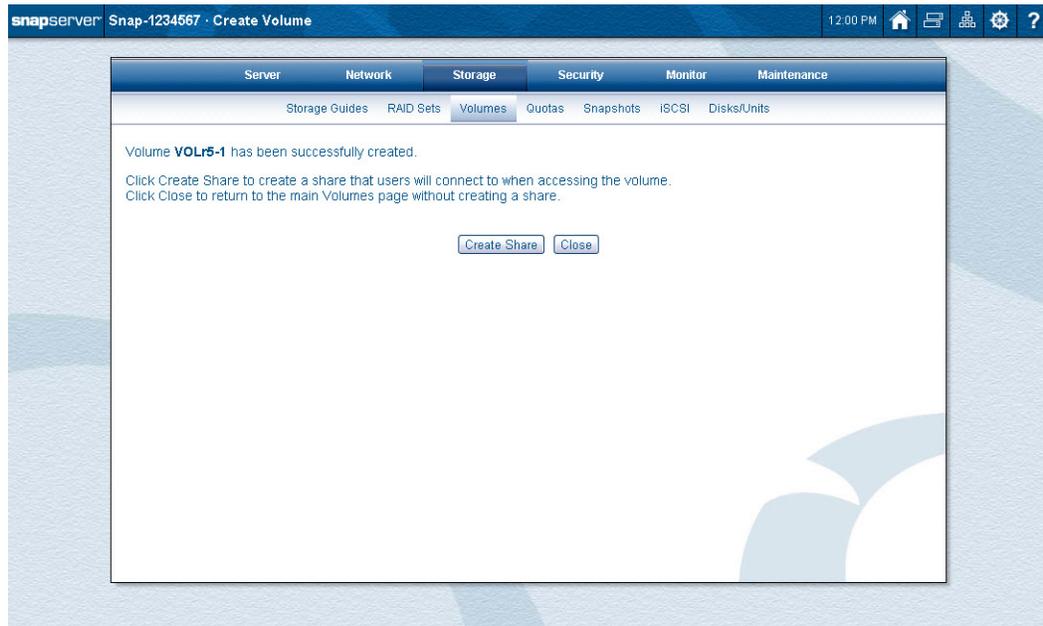


The screenshot displays the 'Create Volume' page in the SnapServer GuardianOS 7.2 web interface. The page title is 'SnapServer Snap-1234567 · Create Volume'. The interface includes a navigation menu with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. The 'Storage' tab is active, and the 'Volumes' sub-tab is selected. The main content area contains the following form elements:

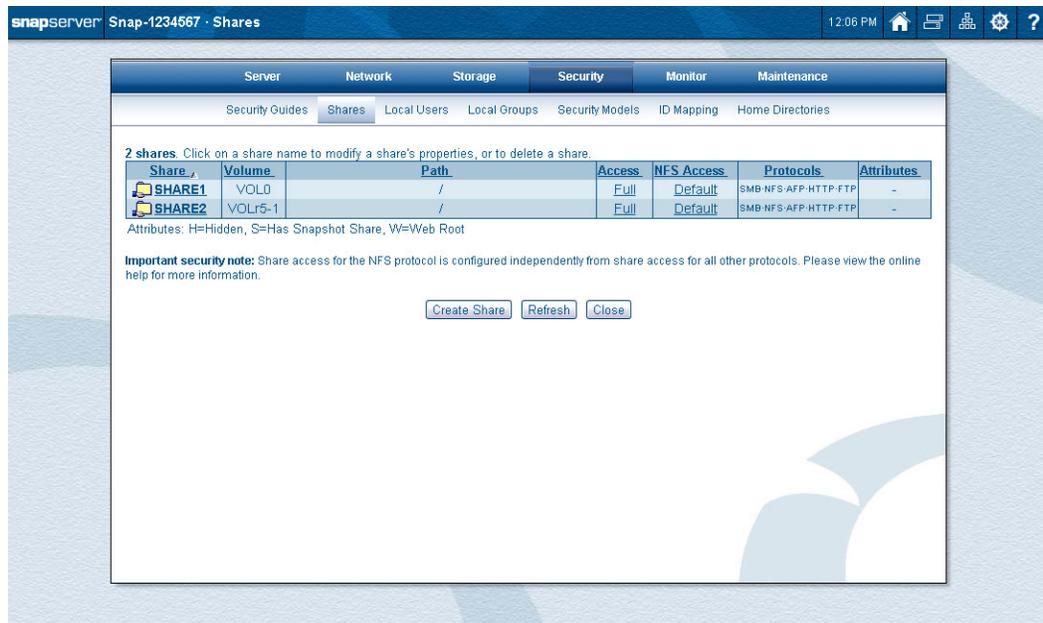
- A heading: "Select the RAID set on which to create the new volume, as well as the new volume's name and size."
- A dropdown menu for "RAID Set" with the selected value "md1 - RAID 5 (15.48 GB available for volume)".
- A text input field for "Volume Name" containing "VOL1".
- A text input field for "Volume Size" containing "15.48" and a unit dropdown menu set to "GB".
- A checked checkbox labeled "Enable Write Cache (not recommended without a configured, online UPS device)".
- A text box with a warning: "You can reserve some of this new volume (typically around 20%) for storing snapshots. The amount you specify will decrease the volume's total size (as specified above) and increase accordingly the size of the snapshot pool. (The snapshot pool is used by all volumes on a given RAID set.)"
- A dropdown menu for "Percentage of this volume's size to add to the snapshot pool for RAID set md1" set to "20%".
- A summary table for "Current allocation of RAID set md1":

Total size:	15.48 GB
Allocated size (for volumes):	0.00 MB
Snapshot pool size:	0.00 MB
- Buttons for "Create Volume" and "Cancel".

At the confirmation page, click **Create Volume** a third time to create the volume. At the successful creation page, click **Create Share** to provide access to this new volume.



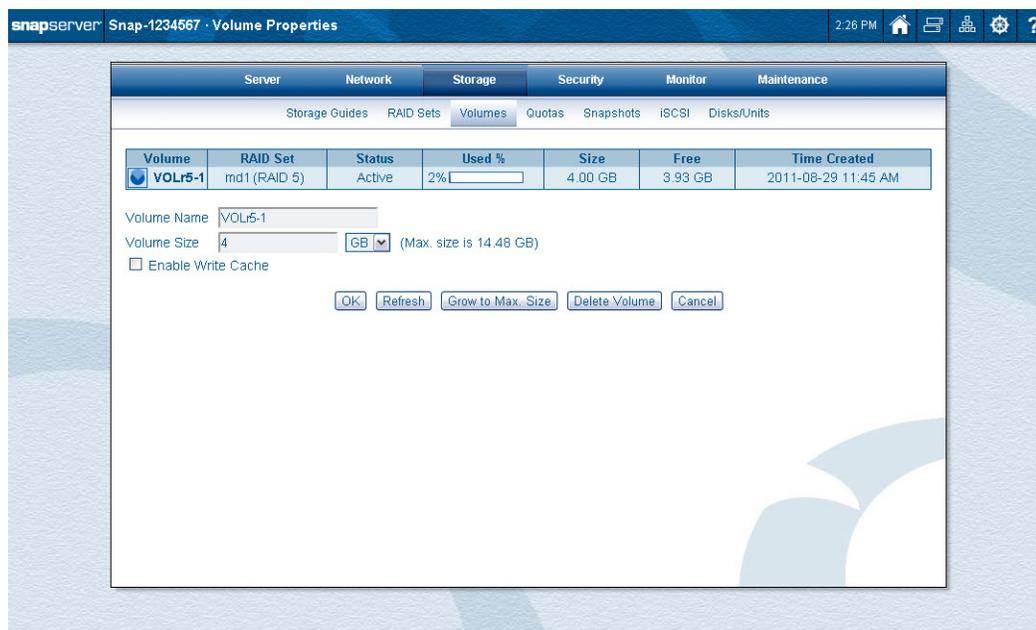
This opens the **Security > Shares** option page. Enter the options you want and click the **Create Share** button again. The share is automatically created and shown in the Share table. See “Shares” for complete details.



Volume Properties

By clicking a volume's name on the main page, details of that particular volume are shown on a **Volume Properties** page. From this secondary page, you can:

- Change the volume name.
- Increase the volume size.
- Enable the write cache (only recommended if a UPS system is attached).
- Delete the entire volume.



Rename a Volume

On the **Volume Properties** page, enter the new name starting with an alphanumeric character and using up to 20 alphanumeric characters or hyphens (but not spaces). Then click **OK**.

Expand Volume Capacity

A volume's capacity can be expanded by navigating to the **Storage > Volumes** page and clicking the name of a volume. There are two ways to expand the size of a volume:

- **Adding Unallocated Capacity** – If there is unallocated capacity remaining on the RAID, you can add this capacity to the volume by either (1) editing the **Volume size field** to a size less than or equal to the maximum size of the volume, or (2) clicking the **Grow to Max. Size** button and then clicking **OK**.

NOTE: You cannot decrease the size of an existing volume. You can, however, delete the volume and recreate it as a smaller size.

- **Creating a New RAID Set** – If all capacity on the existing RAID set is allocated, and either (1) a sufficient number of drives to create a new RAID set exists, or (2) a RAID set of the same type with excess capacity exists, then the **Expand Volume** button appears. Click this button to create an additional RAID set, group the new RAID set with the existing RAID, and then expand the volume into the space on the new RAID group.

This is usually accomplished through the addition of an expansion unit.

NOTE: If you expand the volume onto an existing RAID set with existing volumes, those volumes will be preserved and the expanded volume will only consume the free space on the RAID set.

Configure Write Caching

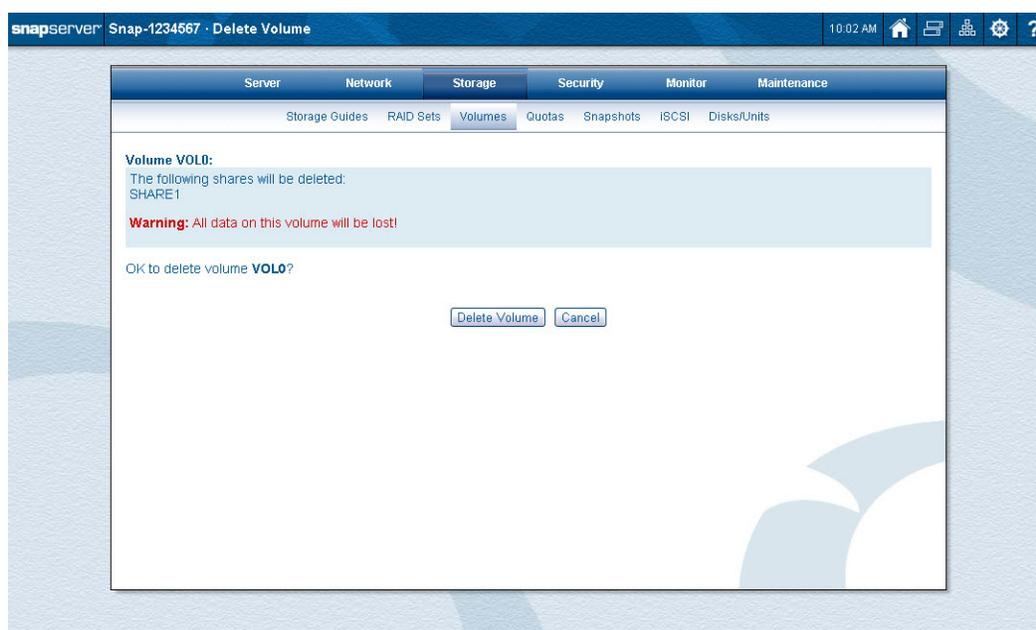
NOTE: This is not related to write caching on iSCSI disks. For information about configuring write caching on iSCSI disks, see “Write Cache Options with iSCSI Disks” on page 6-18.

By default, write cache is enabled on all volumes. For systems that do not use a UPS device to help protect data during a power outage, or for applications that require synchronous writes to disk, write cache can be disabled on a volume-by-volume basis. When a volume's write cache is disabled, all data written to the volume bypasses memory buffers and writes directly to disk, helping to protect the data when writes are occurring during a power outage. While disabling write cache does help protect data, it also significantly impacts disk write performance.

NOTE: When write cache is disabled on a volume, disk cache is also disabled on all disk drives that are members of the RAID or RAID group hosting the volume. This can impact performance on other volumes with write cache enabled that are hosted by the same RAID or RAID group.

Delete a Volume

To delete a volume, go to the **Volume Properties** page and click the **Delete Volume** button. At the confirmation page, click the **Delete Volume** button again.



The volume and all its shares and data are deleted.

Third-Party Applications on Deleted Volumes

Deleting volumes may move or disable certain third-party applications that are installed on the user volume space.

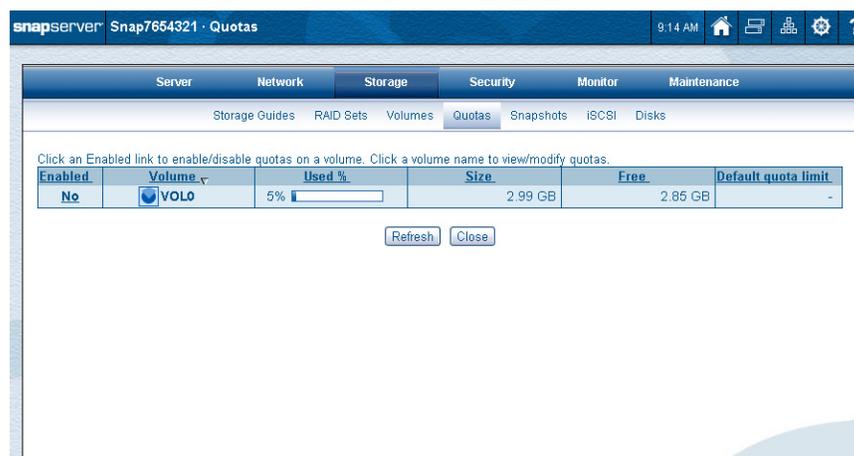
The CA Antivirus software and Snap EDR can reside on one or more volumes. If you delete a volume containing one of these applications, these components will be automatically moved to another volume, or deleted if no other volume or volumes of sufficient space are available. If deleted, CS Antivirus will need to be re-enabled and Snap EDR will need to be reinstalled when a new volume with sufficient space exists.

After creating your new storage configuration, you can reenabling the antivirus software by navigating to the SnapExtensions page and selecting CA Antivirus. On the next page, check the **Enable** checkbox and click **OK**. The SnapServer reinstalls the antivirus software (using default settings) on the volume with the most available space. However, the installation process does not preserve custom antivirus configuration settings, so make a note of any such settings before deleting a RAID or volume. To reconfigure the antivirus software, click **Configure Antivirus**.

To reactivate Snap EDR functionality after creating a new volume, download the Snap EDR package from the SnapServer website and install it on the server using the OS Update feature. Then go to the **Misc. > SnapExtensions** page using the Site Map and enable it.

Quotas

Quotas, which are only available in Traditional RAID, are configured in the **Storage > Quotas** screen of the Web Management Interface.



Assigning quotas ensures that no one user or group consumes a disproportionate amount of volume capacity. Quotas also keep tabs on how much space each user or NIS group is currently consuming on the volume, allowing for precise tracking of usage patterns. You can set individual quotas for any NIS, Windows domain, or local user known to the SnapServer. Group quotas are available only for NIS groups.

For users and groups, there are no pre-assigned default quotas on the SnapServer. When quotas are enabled on the SnapServer, you can assign a default quota for all users, or allow all users to have unlimited space on the volume. Unless you assign individual user or group quotas, all users and groups will receive the default quota when it is enabled.

In calculating usage, the SnapServer looks at all the files on the server that are owned by a particular user and adds up the file sizes. Every file is owned by the user who created the file and by the primary group to which the user belongs. When a file is copied to the server, its size is applied against both the applicable user and NIS group quotas.

Quotas Page

From the Quotas default page, you can enable/disable quotas on the volume by clicking the Yes/No link in the Enabled column on the far left.

NOTE: Quotas can also be accessed from Storage > Volumes by clicking the link in the Quotas column on the far right.

The screenshot shows the SnapServer Quotas page. The page title is "snapserver Snap7654321 - Quotas". The navigation menu includes Server, Network, Storage, Security, Monitor, and Maintenance. Under the Storage menu, there are links for Storage Guides, RAID Sets, Volumes, Quotas, Snapshots, iSCSI, and Disks. The main content area contains a table with the following data:

Enabled	Volume	Used %	Size	Free	Default quota limit
Yes	VOL0	2%	6.19 GB	6.09 GB	2.00 GB
No	VOL1a	2%	2.00 GB	1.96 GB	-
Yes	VOL1b	2%	2.00 GB	1.96 GB	no limit

Below the table are "Refresh" and "Close" buttons.

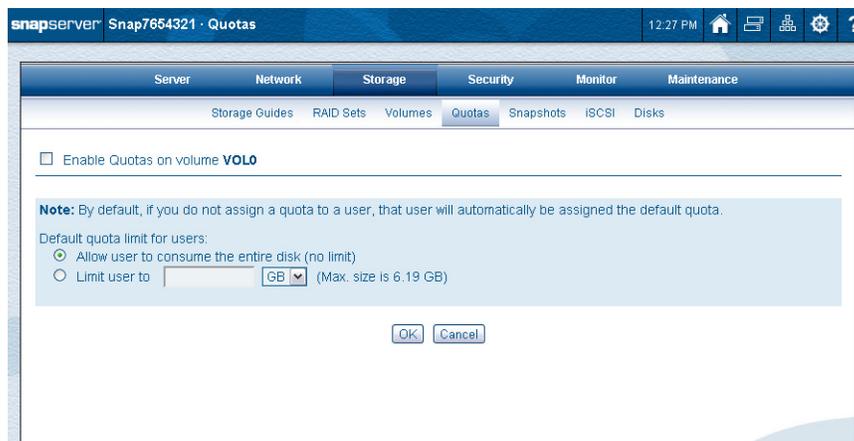
The Default Quota Limit column at the far right shows one of the following:

- **Any amount** – the default quota size assigned to users in that volume who do not have a specific quota assigned to them.
- **“no limit”** – text displayed when quotas are enabled but no default quota size is configured for users in that volume (users can consume the entire disk).
- **“-” (dash)** – a dash indicates quotas are disabled for that volume.

The Default Quota Limit amount on the far right doubles as a link to access the Quota Enable page. Once quotas are set for a volume, the volume name on this page doubles as a link to the view/modify quotas search page.

Enable/Disable Quotas

When you click the Enabled status link (left-most column in the Quota table), a secondary page is shown for managing the quota properties.



1. Check/uncheck the **Enable Quotas** box to enable/disable quotas.

NOTE: The Enable Quotas check box must be checked before changes are accepted.

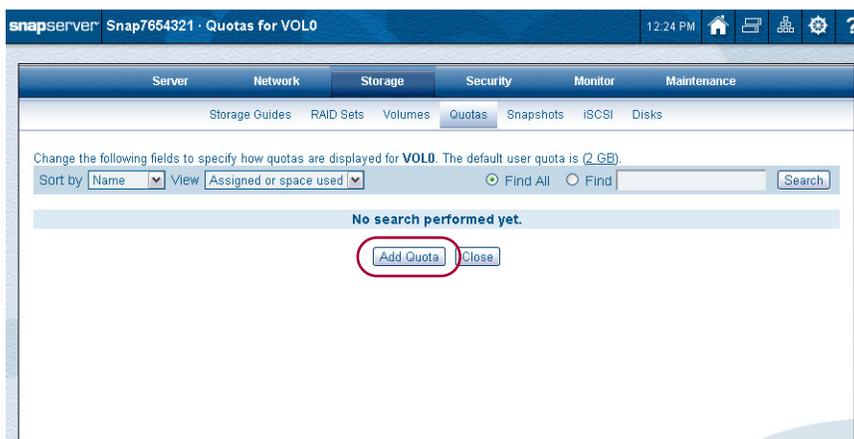
2. Select one of the two **default quota options** to set the quota applied to users who do not have individual quotas assigned to them, and click **OK**.

NOTE: The server may require a restart. If so, a warning message is displayed. Click OK to proceed.

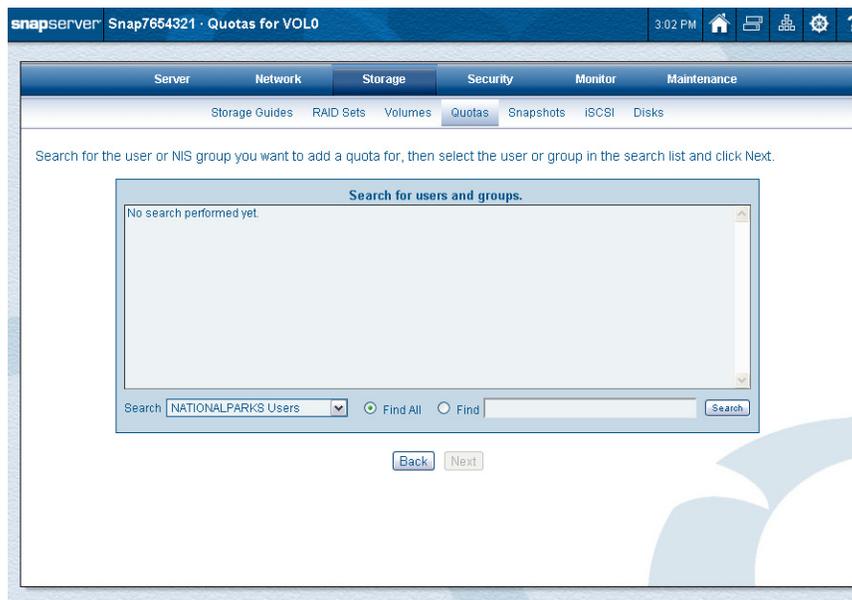
You are returned to the default Quotas page.

Adding Quotas Wizard

1. Click the **Volume name** link on the Quotas default page to open the quota search and configuration page for that volume.

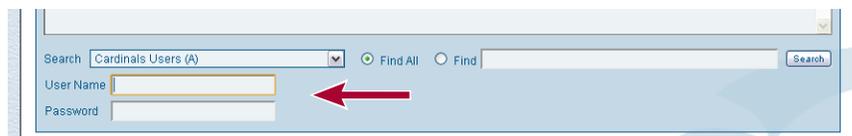


2. Click **Add Quota** to launch the search wizard.



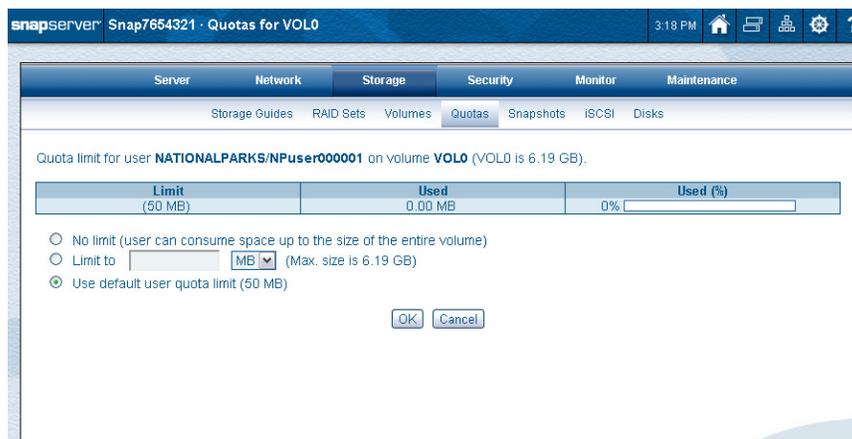
3. To search for a user or NIS group, select the local or domain option from the Search drop-down list, enter the search string (or select Find All), and click **Search**.

NOTE: For domains that require authentication (showing an "(A)" after the name), after selecting the domain name, enter the User Name and Password for that domain.



- Returned results will include all users and NIS groups whose name **begins** with the string entered in the Search field.
- The search results returned may be limited. Fine tune your search by using a more specific string to return the names desired.
- On the rare occasion you need to search for a Windows domain that's not listed ("remote domain"), select a Windows domain from the Search drop-down list through which to search, then enter in the Find box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, `remote_domain\user_name`). After you click Search, another authentication prompt may be presented to authenticate with the remote domain.

- From the search results, select the appropriate **user or NIS group**, and click **Next**.

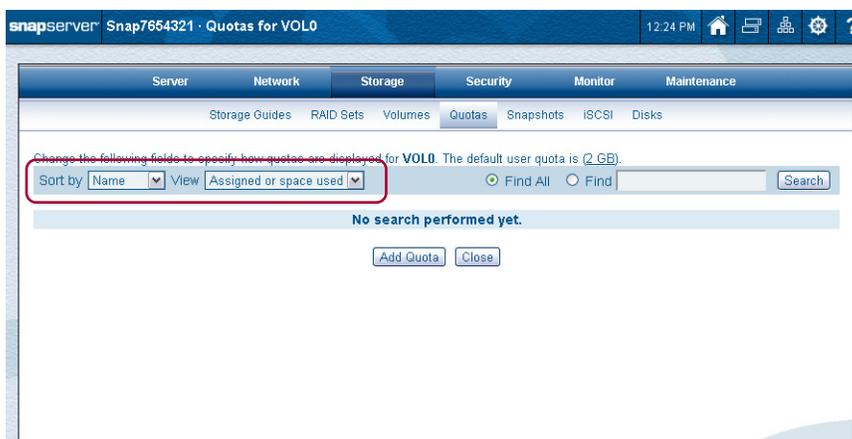


- Select or enter the **quota** desired, and click **OK**.
NIS groups do not display the third option for a default user quota.

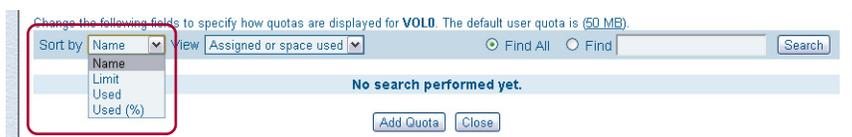
Displaying/Changing Quotas

To display and configure quotas of users or groups that have used space on this volume or have had specific quotas assigned to them from the volume:

- Click the **Volume name** link on the Quotas default page to open the quota search and configuration page for that volume.
- Select the sort and view **parameters**.



- From the **Sort** drop-down list, choose Name, Limit, Used, or Used (%).



- From the **View** drop-down list, choose Only assigned quotas, Only with space used, Assigned or space used, or > 95% used.



3. Enter the **search string** (or select Find All).

When entering a search string:

- Returned results will include all users and groups whose name **begins** with the string entered.
- To search a specific Windows or NIS domain, enter the domain name, followed by a slash (/) or backslash (\) before the search string.
- To search only local user and groups, enter “local” followed by a backslash (\) before the search string.
- The search results returned may be limited. Fine tune your search by using a more specific string to return the names desired.

4. Click **Search**.

A detailed list of users or NIS groups that match the parameters:

User or Group (click to edit)	Domain	Limit	Used	Used (%)
nisGroup000	atlanta Groups	no limit	1.67 TB	24%
nisGroup001	atlanta Groups	982.18 GB	442.65 GB	50%
nisGroup002	atlanta Groups	441.09 GB	147.72 GB	33%
nisGroup003	atlanta Groups	641.58 GB	342.43 GB	53%
nisGroup004	atlanta Groups	no limit	1.73 TB	25%
nisGroup005	atlanta Groups	542.88 GB	295.71 GB	54%
nisGroup006	atlanta Groups	705.74 GB	398.36 GB	56%
nisGroup007	atlanta Groups	371.44 GB	334.92 GB	90%
nisGroup008	atlanta Groups	542.88 GB	475.08 GB	87%
nisGroup009	atlanta Groups	no limit	732.20 GB	10%
nisGroup010	atlanta Groups	1,008.20 GB	920.75 GB	91%

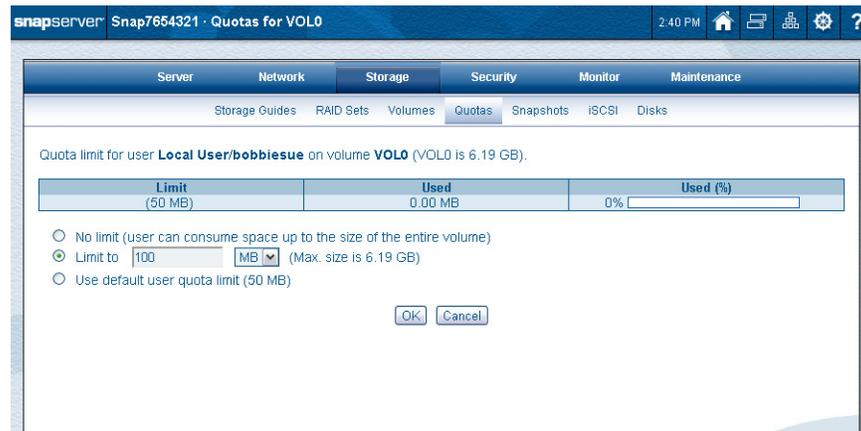
NOTE: The search results returned may be limited. Fine tune your search by using a more specific string to return the names desired.

Parenthesis around a quota limit amount indicates the volume default quota is being used. If the volume's default quota limit is set to “no limit,” then “(no limit)” is displayed. If the volume's default quota limit is set to an actual value, such as 500 GB, then “(500 GB)” is displayed.

No parenthesis around the limit amount indicates a specific quota has been assigned that is different from the default value. If the default quota limit is set to “no limit” but a particular user's or group's quota is set to 750 GB, then “750 GB” is shown instead of the default “(no limit).”

The one exception to this is NIS groups. They don't use a volume default quota, so “no limit” (without parenthesis) is shown.

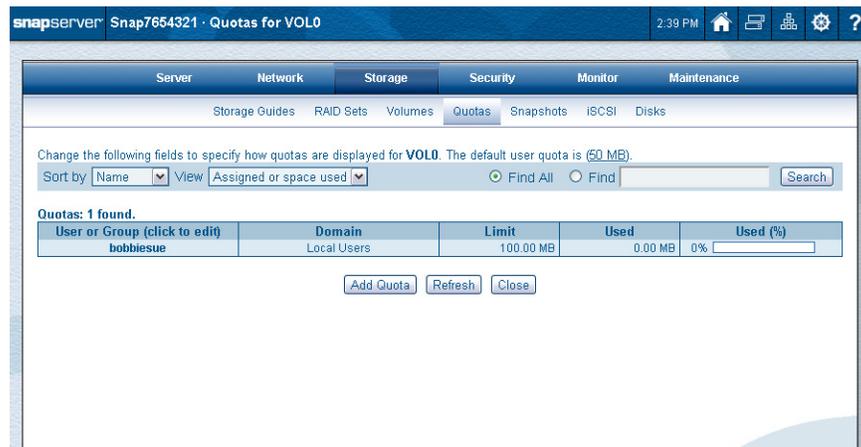
- From the search results, select the appropriate **user or NIS group** from the left column to open the Quotas settings page.



- Select or enter the **quota** desired, and click **OK**.
NIS groups do not display the third option for a default user quota.

NOTE: Any changes override the default volume quota for this user or NIS group.

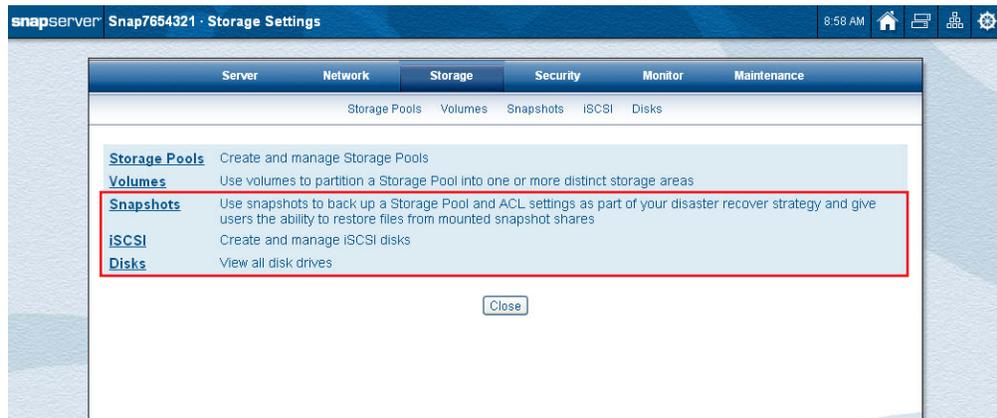
The main Quotas search page is displayed and your changes are reflected here if allowed by your search criteria.



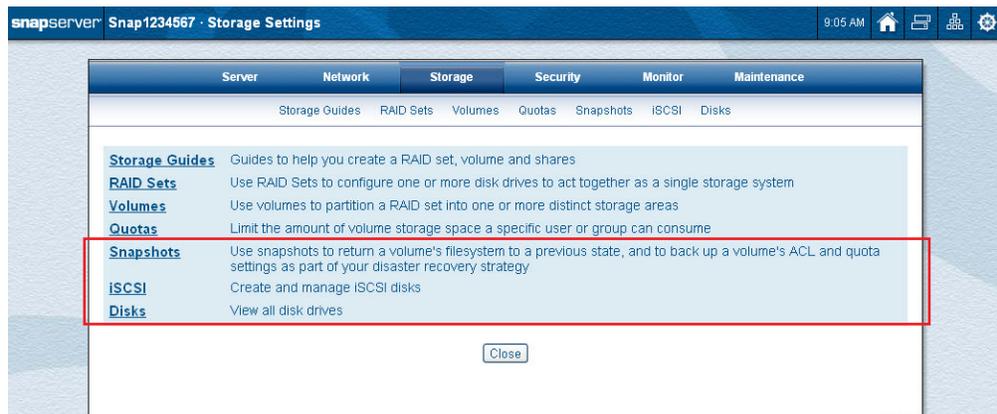
Other Storage Options

Once the RAID sets have been configured using either the DynamicRAID or Traditional RAID options, you can configure the other three storage options for your SnapServer.

DynamicRAID Configuration



Traditional RAID Configuration



For information on the DynamicRAID configuration option, see [Chapter 4, “DynamicRAID Storage.”](#) For information on the Traditional RAID configuration option, see [Chapter 5, “Traditional RAID Storage.”](#)

Topics in Other Storage Options:

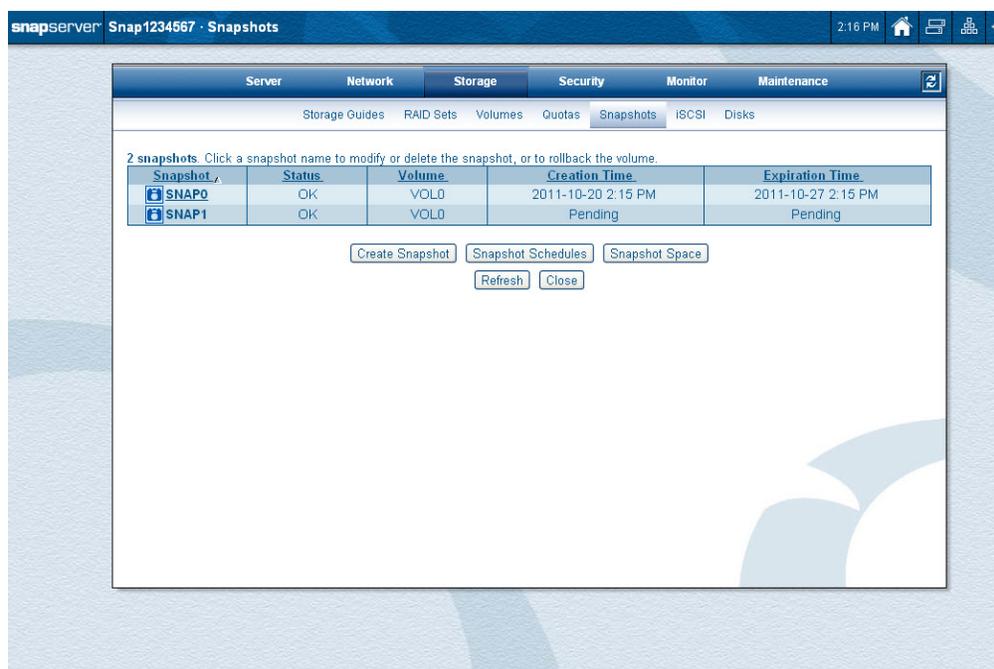
- [Snapshots](#)
- [iSCSI Disks](#)
- [Disks](#)
- [Expansion Units](#)

Snapshots

A *snapshot* is a consistent, stable, point-in-time image of a Traditional RAID volume or DynamicRAID storage pool that can be backed up independent of activity on the live volume or storage pool. Snapshots can also satisfy short-term backup situations such as recovering a file deleted in error, or even restoring an entire filesystem, without resorting to tape. Perhaps more importantly, snapshots can be incorporated as a central component of your backup strategy to ensure that all data in every backup operation is internally consistent and that no data is overlooked or skipped.

NOTE: The Snapshot feature described here does not apply to snapshots for iSCSI disks. Supported Windows servers can create native snapshots of iSCSI disks using VSS. For more information, see “Configuring VSS/VDS for iSCSI Disks” on page 6-22.

To manage the snapshot options using the SnapServer Web Management Interface, go to **Storage > Snapshots**.



These options are available in the Snapshots section of the Web Management Interface:

Action	Procedure
Create a New Snapshot	Click Create Snapshot . The process involves first defining snapshot parameters, and then scheduling when and how often to run the snapshot. Do not take more snapshots than your system can store, or more than 250 snapshots. Under normal circumstances, between nine and ten snapshots are sufficient to safely back up any system.
Edit a Snapshot Schedule	Click the Snapshot Schedules button, and then click the snapshot name. You can modify all snapshot parameters.

Action	Procedure
Adjust Snapshot Space	NOTE: Traditional RAID only. Click the Snapshot Space button, then click the RAID set name for the snapshot space you want to adjust. You can adjust the amount of space allotted for snapshots on each RAID set or RAID group.
Edit and Delete	Click the snapshot's name to open the Snapshot Properties page. You can edit the snapshot's name and duration, or delete the snapshot.
Roll Back a Snapshot	NOTE: Traditional RAID only. Click the snapshot's name to open the Snapshot Properties page. You can roll back the snapshot to a volume.

Clicking the **Refresh** button updates the data shown. This is helpful when waiting for a snapshot to complete.

When single snapshots are originally created or while recurring snapshots are active, a refresh icon () is displayed to the right on the tab bar. It indicates that the snapshot data in the table is being refreshed every 5 minutes.

Clicking the **Close** button returns you to the **Storage** home page.

NOTE: The presence of one or more snapshots on a volume (Traditional RAID) or storage pool (DynamicRAID) can impact write performance. Additional snapshots taken of the same volume or storage pool do not have additional impact; in other words, the write performance impact of one snapshot on a volume is the same as the impact of 100 snapshots on the same volume.

Creating Snapshots

Creating a snapshot involves first defining the snapshot and then scheduling the snapshot. For regular data backup purposes, create a recurring snapshot. A recurring snapshot schedule works like a log file rotation, where a certain number of recent snapshots are automatically generated and retained as long as possible, after which the oldest snapshot is discarded. You can also create individual, one-time-only snapshots as needed.

NOTE: If you have created a new volume or have numerous existing snapshots, make sure you have enough space allocated in the snapshot space; otherwise, you will not be able to create the snapshot.

Scheduling Snapshots

Snapshots should ideally be taken when your system is idle. It is recommended that snapshots be taken before a backup is performed. For example, if your backup is scheduled at 4 a.m., schedule the snapshot to be taken at 2 a.m., thereby avoiding system activity and ensuring the snapshot is backed up. See “[Schedule Snapshots](#)” on page 6-5 for more information.

Snapshots and Backup Optimization

When you back up a live volume directly, files that reference other files in the system may become out-of sync in relation to each other. The more data you have to back up, the more time is required for the backup operation, and the more likely these events are to occur. By backing up the snapshot rather than the volume itself, you greatly reduce the risk of archiving inconsistent data. For instructions, see “[Schedule Snapshots](#)” on page 6-5.

Snapshots and iSCSI Disks

Running a GuardianOS snapshot on a volume containing an iSCSI Disk will abruptly disconnect any clients attempting to write to the iSCSI Disk and the resulting snapshot may contain inconsistent data. Do not use GuardianOS snapshots on a volume containing an iSCSI Disk.

To create a native snapshot of an iSCSI disk on Windows systems, use the VSS feature described in [“Configuring VSS/VDS for iSCSI Disks” on page 6-22](#).

To Create a Snapshot

Follow these steps to create a snapshot:

Step 1: Create the snapshot definition.

Complete the following to define the snapshot:

- Name the snapshot (20 character maximum).
- Identify the source volume/storage pool.

Step 2: Specify when to create the snapshot.

Click either the **Create Snapshot Now** button to run the snapshot immediately or the **Create Snapshot Later** button to schedule the Snapshot for a later time.

When you select the **Create Snapshot Later** button, a new input section appears below the option. Complete the following:

- Schedule a start date and start time to run the snapshot.
- Select either to create the snapshot only once (**One Time**) or to have it recurring. To repeat a snapshot periodically using the **Recurring** option, specify the repeat interval in hours, days, weeks, or months.

Step 3: Specify the duration of the snapshot.

In the **Duration** field, specify how long the snapshot is to be active in hours, days, weeks, or months. The SnapServer automatically deletes the snapshot after this period expires, as long as no older unexpired snapshots exist that depend on it. If any such snapshot exists, its termination date is displayed at the bottom of the page. You must set the duration to a date and time after the displayed date.

Step 4: Specify whether to create a recovery file.

If you plan to create a backup from the snapshot and want to save filesystem security configuration in the backup, check the **Create Recovery File** box. (See the [“Schedule Snapshots” on page 6-5](#) for information on coordinating snapshots and backup operations.)

Step 5: Create the snapshot.

Click **Create Snapshot**. If you elected to run the snapshot immediately, it appears in the Current Snapshots table. If you scheduled the snapshot to run at a later time, it appears in the Scheduled Snapshots table.

Accessing Snapshots

After snapshots are created, they can be accessed via a snapshot share. Just as a share provides access to a portion of a live volume (or filesystem), a snapshot share provides access to the same portion of the filesystem on all current snapshots of the volume. The snapshot share's path into snapshots mimics the original share's path into the live volume. The

snapshot share is created in the **Shares** section under the **Security** tab. See “[Shares](#)” and “[Snapshot Access](#)” for details.

Schedule Snapshots

Like backups, snapshots can be scheduled to recur at a designated time and interval. Part of the initial creation process is to set the time and date when the snapshot will occur or recur.

In addition to synchronizing the backup and snapshot schedules, you must create a share (and snapshot share) to the appropriate directory so that the backup software can access the snapshot. For most backup purposes, the directory specified should be one that points to the root of the volume so that all of the volume's data is backed up and available from the snapshot share.

Step 1: Create a snapshot for each Traditional RAID volume or DynamicRAID storage pool you want to back up.

In the Web Management Interface, navigate to **Storage > Snapshots**, and click **Create Snapshot**. When defining and scheduling the snapshot, consider the following:

- Check the **Create Recovery File** checkbox to ensure that the ACL, extended attributes, and quota information are captured and appended to the snapshot. This step is needed because many backup packages do not back up native ACLs and quotas. Placing this information in a recovery file allows all backup packages to include this information. If the volume needs to be restored from tape, or the entire system needs to be recreated from scratch on a different server, this information may be required to restore all rights and quota information.
- Offset the snapshot and backup schedules such that the backup does not occur until you are sure the snapshot has been created. The snapshot itself does not require much time, but creating the recovery file may take up to 30 minutes, depending on the number of files in the volume.

For example, assuming you schedule nightly backups for a heavily used volume at 3:00 a.m., you might schedule the snapshot of the volume to run every day at 2:30 a.m., allowing half an hour for the snapshot to run to completion.

Step 2: If you have not already done so, create a share for each volume with snapshot share enabled.

In the Web Management Interface, navigate to the **Security > Shares** page, and click **Create Share**. Select the volume you want the share to point to (if you want to create a share to the root of the volume, simply accept the default path). Click **Advanced Share Properties**, then select **Create Snapshot Share**.

Step 3: Set the backup software to archive the latest version of the snapshot.

The SnapServer makes it easy to configure your backup software to automatically archive the most recent snapshot. Simply configure your backup software to copy the contents of the `latest` directory within the snapshot share you created.

For example, assume the snapshot share named `SHARE1_SNAP` contains the following four directories:

```
latest
2011-09-25.120000
2011-10-01.000100
2011-10-07.020200
```

Each directory inside the snapshot share represents a different snapshot. The directory names reflect the date and time the snapshot was created. However, the `latest` directory always points to the latest snapshot (in this case, `2011-10-07.020200`, or October 7th, 2011, at 2:02 a.m.). In this case, configuring the backup software to copy from:

```
\SHARE1_SNAP\latest
```

ensures that the most recently created snapshot is always archived.

Snapshot Space

Snapshots are stored in a RAID set or storage pool in snapshot space reserved within the RAID set for this purpose. Each RAID set on the system contains its own independent snapshot space. This space contains all snapshot data for all the volumes on the RAID set or storage pool.

Estimating Snapshot Space Requirements

Snapshot data grows dynamically for as long as a snapshot is active and as long as there is enough space available in the snapshot space to store them. When the snapshot space approaches its capacity (at about 95 percent), the SnapServer deletes the oldest snapshot's data to create space for more recent snapshot data.

By default, 80 percent of RAID set or storage pool capacity is allocated to volumes and 20 percent to snapshot space. You can adjust the amount of snapshot space on the RAID set or storage pool up (assuming unallocated space exists) or down according to your needs. If you find that your snapshot strategy does not require all of the space allocated to the snapshot space by default, consider decreasing snapshot space capacity and reallocating the capacity to the Traditional RAID volumes or data storage in the DynamicRAID storage pool.

Adjusting Snapshot Space Size

The size of the snapshot space can be adjusted at any time. However, under DynamicRAID, to increase the size of the space a new disk drive must be added to the Storage Pool.

To adjust the size of the snapshot space:

- For DynamicRAID, navigate to the **Storage > Storage Pools** page, and then click the **Storage Pool name** for the snapshot space you want to adjust. Using the drop-down list, select the percentage of space you want to reserve on this pool.
- For Traditional RAID, navigate to the **Storage > Snapshots** page, click the **Snapshot Space** button, and then click the **RAID set name** for the snapshot space you want to adjust. Enter the new amount in the **Snapshot Space** field.

The number of snapshots that a RAID set can support is a function of these factors:

- The space reserved for the snapshot data.
- The duration of the snapshots you create.
- The amount and type of write activity to the volume since the snapshot was created.

The following table describes minimum and maximum allocation cases.

Allocate about 10% of RAID set if	Allocate about 25% of RAID set if
<ul style="list-style-type: none"> • Activity is write-light. • Write access patterns are concentrated in a few places. • A small number of Snapshots must be available at any point in time. 	<ul style="list-style-type: none"> • Activity is write-heavy. • Write access patterns are randomized across the volume. • A large number of Snapshots must be available at any point in time.

There are two other processes that may affect the size of the snapshot space:

- **Creating a Traditional RAID Volume** – In the course of creating a new volume, a drop-down list allows you to add a percentage of the capacity being allocated to the new volume to the snapshot space. This feature defaults to 20 percent, the recommended amount of space to reserve for snapshots. If you do not plan to use snapshots with this volume, maximize volume capacity by reducing this percentage to zero; if you do plan to use snapshots, adjust this percentage in accordance with the guidelines discussed in the previous section, “[Estimating Snapshot Space Requirements](#).”
- **Creating a Traditional RAID Group** – When two or more RAID sets are grouped together, their snapshot spaces are added together. For example, if RAID set A with a snapshot space of 50 GB is grouped with RAID set B with a snapshot space of 25 GB, the resulting RAID set group will have a snapshot space of 75 GB. Depending on the purpose you had in mind when grouping the RAID sets, the result of combining the two snapshot spaces may or may not be desirable, and you will need to readjust the size as described previously.

Snapshot Properties

From the **Snapshot** primary page table, you can click a snapshot name to access the **Snapshot Properties** page. There you can edit the name and duration, delete the snapshot, or, for Traditional RAID configurations, roll back to a previous state.

Edit a Snapshot

You can edit the name and duration by changing the data in the detail fields.

Delete a Snapshot

Click the **Delete Snapshot** button and then click it again on the confirmation page. The snapshot is deleted and all its associated data.

Rollback to a Previous State

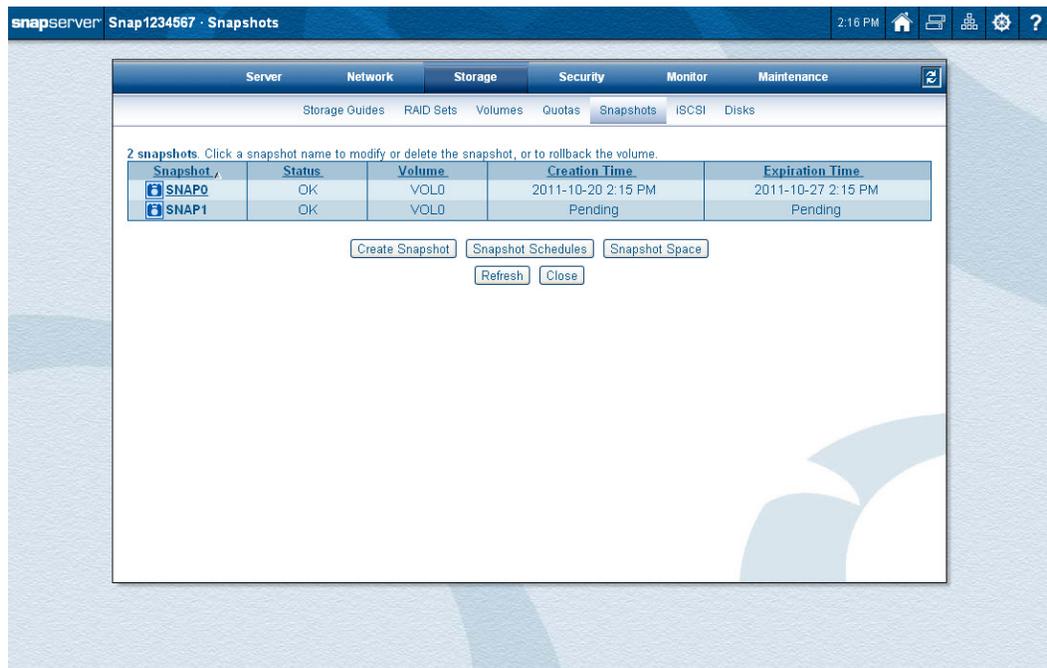
NOTE: This is only available on a Traditional RAID configuration.

If you need to restore an entire filesystem to a previous state, you can do so without resorting to tape. The snapshot rollback feature allows you to use any archived snapshot to restore an entire filesystem to a previous state simply by selecting the snapshot and clicking the **Rollback** button. During the rollback operation, data on the volume will be inaccessible and changes blocked.

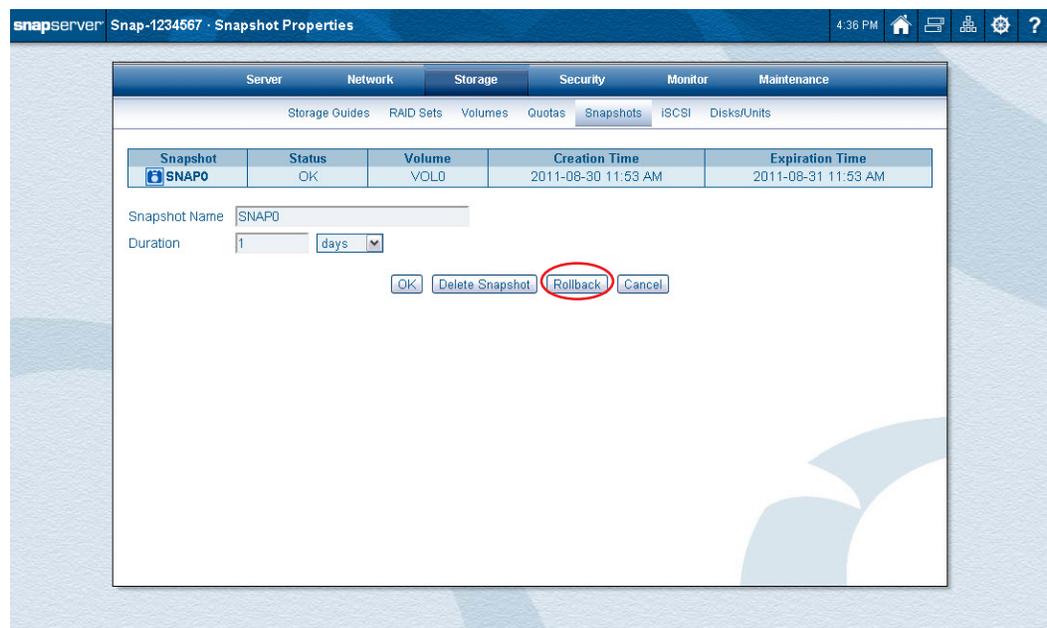


CAUTION: Rolling back a volume cannot be undone and should only be used as a last resort after attempts to restore selected directories or files have failed. Performing a rollback on a volume may disable the antivirus software. If you are using the antivirus software, take the necessary precautions described in [Chapter 11, “CA Antivirus Software.”](#)

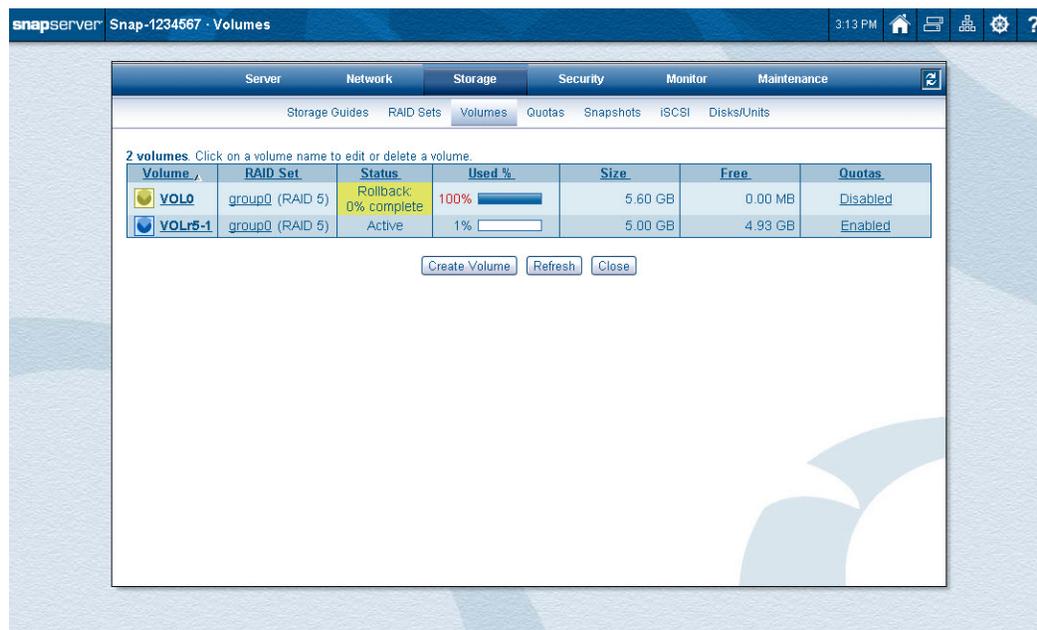
To access the Rollback option, navigate to the **Storage > Snapshots** page.



Click the name of the Snapshot you want to use that is shown in the left-most column. At the displayed **Snapshot Properties** page, click the **Rollback** button.



At the confirmation page, click the **Rollback** button again. The **Storage > Volumes** page is displayed showing the rollback progress.



IMPORTANT: A rollback can disable Snap EDR and result in its removal. If this occurs, download Snap EDR from the SnapServer website, reinstall it using the OS Update feature, then re-enable and configure it from the SnapExtensions page.

iSCSI Disks

Internet SCSI (iSCSI) is a standard that defines the encapsulation of SCSI packets in Transmission Control Protocol (TCP) and their transmission via IP. On SnapServers, an iSCSI disk is based on an expandable, RAID set-protected volume, but appears to a client machine as a local SCSI drive. This storage virtualization frees the administrator from the physical limitations of direct-attached storage media and allows capacity to be expanded easily as needed. Unlike standard SnapServer volumes, SnapServer iSCSI disks can be formatted by the iSCSI client to accommodate different application requirements.

Connectivity to the iSCSI disk is established using a software package or PCI card, known as an initiator, that must be installed on a client machine. The initiator sees the SnapServer as a “target portal” and an iSCSI disk as a “target.”

To use the SnapServer as an iSCSI target, you need to configure iSCSI on both the client initiating the iSCSI connection, and on the SnapServer. Use the information presented here in conjunction with the documentation supplied with your initiator to install, configure, and connect the iSCSI initiators to the SnapServer.

iSCSI Disk Limitations:

- The the size of any iSCSI disk is limited to 2 TB.
- GuardianOS can maintain up to 256 iSCSI disks.

For Additional Information:

The following resources provide further information you may need to plan and complete your iSCSI implementation.

- **SnapServer Online Help** – Available from the **Storage > iSCSI page**, the online help provides details on creating and managing iSCSI disks on SnapServers.
- **RFC3720: Internet Small Computer System Interface (iSCSI)** – Detailed specification for the iSCSI protocol, available from <http://www.ietf.org>.
- **RFC4171: Internet Storage Name Service (iSNS)** – Detailed specification for the iSNS protocol, available from <http://www.ietf.org>.
- **The Microsoft iSCSI Software Initiator User's Guide** (*uguide.doc*) – This document is packaged with the initiator download and installs to the default location, usually: `C:\Windows\iscsi\uguide.doc`. It can also be downloaded from the [Microsoft website](#).
- **The SANSurfer iSCSI HBA CLI Application Users Guide** – This document is available for download on the QLogic website at http://support.qlogic.com/support/drivers_software.asp.
- **The RedHat or Novell (SuSE Linux) websites** – Information on configuring the Linux in-box initiators can be found by searching for *iSCSI* on the RedHat (<http://www.redhat.com>) or Novell (<http://www.novell.com/home/>) websites.
- **The Novell NetWare Administrator's Guide** – This document is available for download on the [Novell website](#).
- **The VMware Server Configuration Guide** – This document is available for download on the [VMware website](#).
- **ReadMe files and Help menus** – For Solaris 10 and operating systems using Open iSCSI (SuSE 10, RedHat 4/5, and CentOS 5), the readme files and help menus provide information on installing and configuring iSCSI.
- **Specifications, Briefs, and White Papers** – The Overland Storage website offers a wide array of informational guides regarding iSCSI and its uses, from product overviews and problem solving for iSCSI, to product specifications and knowledge base articles. For more information about iSCSI and its uses, please browse the Overland Storage website.

Configuring iSCSI Initiators

Overland Storage has qualified a number of software initiators, PCI cards, and drivers to interoperate with SnapServers.

The following sections briefly describe the initiators supported by GuardianOS and some of the more common configuration options.

- [iSCSI Configuration for Microsoft Windows using MS Initiator](#)
- [iSCSI Configuration for Linux and UNIX](#)
- [iSCSI Configuration for Novell NetWare](#)
- [iSCSI Configuration for VMware](#)
- [iSCSI Configuration for Mac](#)

iSCSI Configuration for Microsoft Windows using MS Initiator

Installation and configuration information is included with the MS Initiator download (*uguide.doc*). It can also be downloaded from the Microsoft website.

Before implementing iSCSI using MS Initiator, please consider the following:

- On pre-Vista operating systems, Microsoft does not support dynamic disks for use with the Microsoft iSCSI initiator. Overland Storage recommends:
 - Using a QLogic QLA4010 or QLA4050 HBA which supports dynamic disks.

- Using only “basic” disks with the Microsoft initiator to avoid unexpected behavior and possible data loss when connecting to iSCSI targets in a SnapServer.
- To extend the size of a basic disk on pre-Vista operating systems, use the diskpart.exe utility as described in “[To Use the Microsoft Diskpart Utility to Grow iSCSI Basic Disks](#)” on page 6-13 or refer to Microsoft KB article [325590](#). The Microsoft knowledge base can be found at <http://support.microsoft.com>. On Vista, Windows 2008, and Windows 7 systems, use the disk management tool to resize the disks.

To Configure Microsoft Services Installed on iSCSI Disks to Start Automatically. iSCSI technology allows SnapServers to host the data files for applications that otherwise require local disk storage, such as MS SQL Server 2000 and Exchange Server 2003. If you use the Microsoft initiator on Windows XP, Windows 2003, Vista, Windows 7, or Windows 2008 server, services installed on iSCSI disks will start up automatically by default once you have configured them to persistently reconnect. On the Windows 2000 server, however, you must edit the Windows registry to make the service dependent on the iSCSI Initiator Service.



CAUTION: Use the Registry Editor with caution. Changes suggested by SnapServer should be evaluated by qualified technical staff to ensure that they do not affect the proper functionality of the Windows implementation, installed applications, or other components on the Windows system whose registry is being modified. The result of any modifications to the Windows registry can vary. Implied outcomes of any modification suggested by SnapServer are NOT guaranteed, and may not be supported.

Overland Storage strongly recommends backing up your registry before making any modifications. Please see Microsoft Knowledge Base article [322755](#) (Windows 2000) for details on backing up and restoring the Windows registry.

To Configure the Server to Persistently Connect. 1. Create an iSCSI disk on the SnapServer (see “[Create iSCSI Disks](#)”).

2. From the Target tab of the Initiator's Property dialog box, select the Target and click the **Logon** button
3. Check the **Automatically restore this connection when the system reboots** box to make this a persistent target, and click **OK** to log in to the SnapServer target.
4. Use the Disk Administrator to configure all **volumes** on top of the disks.
5. From the Bound Volumes/Devices tab on the Property dialog box, click **Bind All** to allow the iSCSI service to configure the list of persistent volumes.

If you are running Windows XP, Windows 2003 Server, Vista, Windows 7, or Windows 2008 Server, your iSCSI disks will now start automatically on reboot.



IMPORTANT: If you are running Windows 2000 Server, you must continue to the following procedure and edit the registry to make services dependent on the iSCSI Initiator service.

To Edit the Windows Registry for MS Exchange Server or MS SQL Server.

NOTE: This applies to Windows 2000 only.

1. Install **Exchange Server 2003** and configure it to use the iSCSI disk as the location to store database files.
2. On a Windows workstation running Windows 2000, enter on the command line **regedt32**.
3. Navigate to the **registry key**:

- For Exchange Server:
`HKey_Local_Machine > System > Current Control Set > Services > lanmanserver`
- For SQL Server:
`HKey_Local_Machine > System > Current Control Set > Services > MSSQLServer`

4. If **DependOnService** does not exist, create it:
 - a. Select **Add Value** from the Edit menu.
 - b. In the **Name** field of the Add Value dialog box, enter **DependOnService**.
 - c. Click **OK**.
5. Double-click **DependOnService**.
6. In the Data box that opens, enter **msiscsi**, click **OK**, and then close the registry.
7. **Reboot** the Windows server.

To Configure Shares to iSCSI Disks. When using the Microsoft initiator, shares to iSCSI disks may not automatically reconnect when the Windows system hosting the shares is rebooted. There are two methods to resolve this issue:

- Share an iSCSI target that has an assigned drive letter. This method requires changes to the Windows registry and is described in [Microsoft Knowledge Base article #870964](#).
- Mount the iSCSI disk to a folder on an existing NTFS volume as described in “[To Mount an iSCSI Disk Without a Drive Letter](#)”. This method does not require changes to the Windows registry and is described below.

To Mount an iSCSI Disk Without a Drive Letter. To complete this procedure, you must create and format an iSCSI target on the SnapServer and connect to this iSCSI disk using the Microsoft initiator. You must also have an existing NTFS volume on a local disk within the Windows server, initiating the connection.

1. Right-click **My Computer** and select **Manage**.
The new formatted volume will appear in the Disk Management window.
2. Right-click the **New Volume** and select **Change Drive Letter and Paths**.
3. Click **Remove** in the Change Drive Letter and Paths for (New Volume) dialog, and click **Yes** to confirm drive letter removal.
4. Right-click the **New Volume** again and select **Change Drive Letter and Paths**.
5. Select **Add** in the Change Drive Letter and Paths for (New Volume) dialog.
6. In the Add Drive Letter or Path dialog, select **Mount in the following empty NTFS folder**.
7. Create a folder or enter the path to the one that will be shared from the Windows server and select **OK**.
8. Select **OK** in the Add Drive Letter or Path dialog. This returns you to the Disk Management window.
You will see the icon of a disk in place of the folder icon in the File Management window.
9. Create a **share** to the iSCSI disk in the standard method, then reboot the Windows machine and verify that the share is persistent.

To Configure Dynamic Disks to Persistently Reconnect. On pre-Vista operating systems, when iSCSI targets are configured as dynamic disks, the Microsoft iSCSI initiator connecting to the dynamic disk may fail to connect properly during system boot. Using dynamic disks for iSCSI targets on pre-Vista operating systems is not supported by Microsoft. For more information, see the *Microsoft iSCSI Software Initiator User's Guide*, available on the Microsoft website (*uguide.doc*).

To Use the Microsoft Diskpart Utility to Grow iSCSI Basic Disks. In a Microsoft environment, *basic disk* is the simplest configuration method for an iSCSI disk. Basic disks are given the highest priority at both system and application services startup to ensure proper initialization.

For Vista, Windows 7, and Windows 2008 Server, use the Disk Management utility. For Windows 2003 Server, Windows 2000 Server, and Windows XP, Microsoft offers a command line utility called Diskpart that allows you to expand basic disks. This utility ships with Windows 2003 Server, and is available for download for Windows 2000 Server and XP. Additional details on the Diskpart utility can be found in [Microsoft Knowledge Base article Q300415](http://support.microsoft.com/kb/300415) (<http://support.microsoft.com/kb/300415>).

Step 1: Preparing to Expand a Microsoft Basic iSCSI Disk

The following steps must be taken to prepare for the expansion of a basic iSCSI disk from a Windows host:

1. Using the Microsoft Services GUI, stop all application services that are using the volume you intend to expand.
2. If it is not already installed, load the Diskpart utility on the host machine that is running the iSCSI initiator

NOTE: If Diskpart is already installed, you will get the appropriate response when entering `diskpart` at the command line. If the command returns `command not found`, locate [Diskpart on the Microsoft website](#), download the utility, and install it on the local host.

3. Log off the iSCSI volume that is to be expanded:
 - a. Open the Microsoft initiator tool.
 - b. Under Connected Targets, highlight the specific iSCSI disks you want to expand.
 - c. Click **LogOff**. This will log you off the specific target.
4. Verify that you have additional space available on the SnapServer to expand an existing volume:
 - a. Open the Web Management Interface from a client on the network.
 - b. Navigate to **Storage > iSCSI**.
 - c. Select the iSCSI disk you intend to expand.

If you have not disconnected from the iSCSI disk at the host, you will be unable to proceed to the configuration page.
 - d. From the configuration page, ensure that you have additional space on the volume to expand the selected iSCSI disk.
 - e. Make changes to the iSCSI disk size as desired.
 - f. Click **OK**. The disk should now reflect the larger size.

Step 2: To Expand the Basic Disk on the Microsoft Host

1. Open the Microsoft initiator tool.

2. Under **Available Targets**, highlight the specific iSCSI disks you expanded in the previous procedure.
3. Click **LogOn**. This will connect the initiator to the selected iSCSI target.
4. Close the Microsoft initiator tool.
5. Open the Disk Management tool by right-clicking **My Computer** and selecting **Manage**. In the Computer Management GUI, select **Disk Management**.
The disk will automatically reattach, and the additional expanded space in the iSCSI disk will appear as unallocated space on the same disk.

Step 3: To Expand an iSCSI Volume using the Microsoft Diskpart Utility

1. In the Start menu, select **Run** and enter **cmd** in the Run dialog to open a command-line window.
2. Enter the command **diskpart**.
3. To show all the available disks on the host, enter
list disk
4. Identify the specific disk you are expanding.
5. To show all the available volumes on the host, enter:
list volume
6. Identify the specific **volume** you are expanding.
7. Enter the necessary data:
 - a. **select disk n**
where *n* is the disk number that Diskpart indicated from the list command.
 - b. **select volume n**
where *n* is the volume number that Diskpart indicated from the list command.
 - c. **extend size=n**
where *n* is the number of megabytes you want to expand the disk.
For example, if you are adding 10 GBs to an existing disk of 100 GBs, use the following command:
 - d. **extend size=10240** (the number is in megabytes, 1024MBs = 1GB)
The Disk Management GUI will show the newly expanded disk size.
8. Exit the **Computer Management** tool.
9. **Restart** the necessary application services.

To Configure the QLogic iSCSI Initiators for Microsoft Windows. The Overland Storage-recommended QLogic QLA4010 and QLA4050/52c HBAs are iSCSI adapters that appear as a SCSI adapter instead of a network adapter in Windows Device Manager. Before a QLA4010 or QLA4050/52c can successfully connect to iSCSI targets, you must:

- Set initiator parameters (for example, initiator name, alias, IP address).
- Enter target information (for example, target portal information and target iSCSI name).

You can use either the SANSurfer Management application that came with the QLA4010/4050/4052c or Microsoft's iSCSI initiator applet to set initiator parameters and enter target information. Follow the instructions in the documentation to install and configure the adapter.

iSCSI Configuration for Linux and UNIX

Before implementing iSCSI on Linux or UNIX systems, consider the following:

- The QLogic QLA4010/4050/4052c hardware initiator supports Red Hat Enterprise Linux 3, QU5; Red Hat Enterprise Linux 4, QU1; and SuSE Linux Enterprise Server 9, SP3. This initiator provides CHAP authentication and can connect to multiple targets simultaneously. The SANSurfer utility is included with the HBA to initiate, monitor, and change iSCSI targets using its text-based user interface.
- The Cisco-based in-box iSCSI software initiators for Linux support Red Hat Enterprise Linux 3, QU6, Red Hat Enterprise Linux 4, QU2, and SuSE Linux Enterprise Server 9, SP3.
- The Open iSCSI-based in-box iSCSI software initiators for Linux support Red Hat Linux 5 QU1 and higher, SuSE Linux Enterprise Server 10, SP1 and higher and CentOS 5.0 and higher.
- The Open iSCSI-based in-box iSCSI software initiator for UNIX supports Solaris 10 U4.

Installation and configuration information for the QLogic QLA4010/4050/4052c HBA is included with the adapter and is also available for download from the QLogic website. Information about the in-box iSCSI initiators is available from the Red Hat, Novell (SuSE Linux), and Sun Microsystems websites.

To Use CHAP Authentication to Enable Multiple Linux Systems to Share iSCSI Disks

Securely on a SnapServer. You can use CHAP authentication to enable multiple Linux systems with in-box initiators to share different iSCSI disks on a SnapServer or SnapServers. To do this, you would set up different user names and passwords for a Discovery Address.

For example, on a SnapServer (IP address:192.3.2.193), iSCSI disks can be configured for System A and System B. With CHAP enabled, set the System A User name to *a*, and set the Password to *PasswordForA*. Then, for system B, set the user name *b*, and set the password to *PasswordForB*. The configuration will look like the following:

In System A's `/etc/iscsi.conf`, enter the following:

```
DiscoveryAddress=192.3.2.193
Username=a
Password=PasswordForA
```

In System B's `/etc/iscsi.conf`, enter the following:

```
DiscoveryAddress=192.3.2.193
Username=b
Password=PasswordForB
```

System A and B can connect to their own iSCSI disks on the same SnapServer (IP address 192.3.2.193) without the possibility of data corruption caused by sharing the same iSCSI disk.

iSCSI Configuration for Novell NetWare

Consider the following information before implementing iSCSI on NetWare servers:

- NetWare 6.5 with SP1 for NetWare is required, and the iSCSI packages must also have been installed using the Custom Install method to utilize the NetWare iSCSI initiator.

- The server initiating the connection should be a P-III or higher with a minimum of 512MB of RAM and a GbE adapter. To validate the NetWare server's ability to communicate with the SnapServer, ping the SnapServer from the NetWare server.
- With GuardianOS 5.0 or later, CHAP authentication is supported on NetWare 6.5, SP7.

NOTE: CHAP authentication is not supported on versions of NetWare 6.5 earlier than SP7, nor is it supported on pre-GuardianOS 5.0 systems.

- iSCSI implementation requires configuration using the NetWare Remote Manager or the command line in the Server Console.

For more information regarding installation and configuration of required NetWare components, refer to the documentation included with the Novell initiator distribution.

iSCSI Configuration for VMware

When you install VMware ESX Server or vSphere Server, the iSCSI Initiator is automatically installed.

On connecting to the SnapServer targets, the VMware ESX 3.5 Server initiator will find all iSCSI disks and automatically log into them. If iSCSI disks are shared across multiple servers, you can use CHAP authentication to restrict the number of iSCSI disks the VMware initiator can access. See "Create iSCSI Disks" for more information. The VMware vSphere 4.0 Server initiator provides the option for Static Discovery, allowing you to enter the IP addresses of only those targets you want the VMware initiator to access.

For more information regarding installation and configuration of required VMware components, refer to the documentation included with the VMware Server installation.

To Use the VI Client to Configure iSCSI Services. Follow the instructions in the *VMware Server Configuration Guide*, available from <http://www.vmware.com>, to configure your iSCSI service. Use the VI Client to:

1. Configure the Service Console that connects to the VMware host.
2. Create the VMKernel on the NIC used for the iSCSI connection.
3. Enable the iSCSI software initiator, set up target IP addresses, and configure CHAP authentication (if desired). Rescan if necessary to see the new iSCSI service.

On pre-VMware ESX 3i systems, you must open a port in your security profile to enable the iSCSI port. From the Configuration tab, select **Security Profile**, click **Properties**, and check the port for the iSCSI Initiator.

4. Use the **Add Storage** option to configure your storage.

iSCSI Configuration for Mac

GuardianOS supports globalSAN 4.1 and the SmallTree abcSAN iSCSI initiator for use with Mac OS 10.5 and later. Download the initiator software from the [SmallTree website](#), and follow the installation instructions.

NOTE: If iSCSI is used on a SnapServer with more than one Ethernet port, Mac OS X iSCSI clients can encounter connectivity issues if multiple ports are connected to one or more networks. To avoid these issues, configure the server from Network > TCP/IP to enable and connect only one standalone interface or one bonded pair (Load Balance, Failover, etc.) to a single network.

iSCSI Configuration on the SnapServer

iSCSI disks are created on the **Storage > iSCSI** page of the Web Management Interface. Before setting up iSCSI disks on your SnapServer, carefully review the following information.

Basic Components of an iSCSI Network

iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. A basic iSCSI network has two types of devices:

- iSCSI initiators, either software or hardware, resident on hosts (usually servers), that start communications by issuing commands; and
- SCSI Targets, resident on storage devices, that respond to the initiators' requests for data.

The interaction between the initiator and target mandates a server-client model where the initiator and the target communicate with each other using the SCSI command and data set encapsulated over TCP/IP. Overland Storage is one of the first to embed iSCSI target support in its SnapServers.

Isolate iSCSI Disks from Other Resources for Backup Purposes

It is important to isolate iSCSI disks from other resources on the SnapServer for two reasons:

- The filesystem of an iSCSI disk differs fundamentally from the SnapServer's native filesystem.
- iSCSI disks are managed from client software rather than the SnapServer's Web Management Interface.

For ease of management and particularly for data integrity and backup purposes, either dedicate the entire SnapServer to iSCSI disks, or if the server is to be used with other shared resources, place the iSCSI disk and the other shared resources on separate volumes.

- **Back up an iSCSI Disk from the Client, not the SnapServer** – An iSCSI disk is not accessible from a share and thus cannot be backed up from the SnapServer. The disk can, however, be backed up from the client machine from which the iSCSI disk is managed.

NOTE: While some third-party, agent-based backup packages could *technically* back up an iSCSI disk on the SnapServer, the result would be inconsistent or corrupted backup data if any clients are connected during the operation. Only the client can maintain the filesystem embedded on the iSCSI disk in the consistent state that is required for data integrity.

- **Do Not Use the GuardianOS Snapshots Feature on a Volume or Storage Pool Containing an iSCSI Disk** – Running a GuardianOS snapshot on a volume or storage pool containing an iSCSI disk will abruptly disconnect any clients attempting to write to the server's iSCSI disk, and the resulting snapshot may contain inconsistent data. Supported Windows servers can create a native snapshot of a SnapServer iSCSI disk using VSS (see [“Configuring VSS/VDS for iSCSI Disks”](#) on [page 6-22](#) for more information).

iSCSI Multi-Initiator Support

The **Support Multi-Initiator** checkbox allows two or more initiators to simultaneously access a single iSCSI target. Multi-Initiator Support is designed for use with applications or environments in which clients coordinate with one another to properly write and store data on the target disk. Data corruption becomes possible when multiple initiators write to the same disk in an uncontrolled fashion.

NOTE: GuardianOS 7.2 supports Windows 2003 and Windows 2008 Server failover clustering.

When the checkbox for Support Multi-Initiator is selected, a warning message **Uncontrolled simultaneous access of multiple initiators to the same iSCSI target can result in data corruption. Only enable Multi-Initiator Support if your environment or application supports it.** appears. It functions as a reminder that data corruption is possible if this option is used when creating an iSCSI disk.

Write Cache Options with iSCSI Disks

NOTE: This section refers only to iSCSI disks. For information about configuring write cache on GuardianOS volumes on a Traditional RAID configuration, see "Volume Properties".

To ensure the fastest possible write performance, SnapServers can buffer up to 1GB of data to efficiently handle data being transmitted to a SnapServer. This widely accepted method of improving performance is not without some risk. For example, if the SnapServer were to suddenly lose power, data still in cache would be lost.

This risk can be minimized by following industry-standard security precautions, such as keeping servers in a secured location and connecting power supplies to the mains using a network- or USB-based UPS. In most environments, taking these simple precautions virtually eliminates the risk of serious data loss from sudden and unexpected power outages.

Of course, the physical conditions and company policies that guide IT decisions vary widely. Power outages are a common occurrence in some areas, and data protection procedures vary from company to company. Administrators who determine that the risk of data loss, even with security cautions in place, outweighs the significant increase in write performance that write cache provides, can disable this feature for individual iSCSI disks.

When working with write cache for iSCSI disks, note the following:

- Write cache can be disabled on an iSCSI-disk-by-iSCSI-disk basis. Disabling write cache for an iSCSI disk does *not* disable write cache for any other iSCSI disk or any other resources on the SnapServer.
- The write cache for an iSCSI disk can be enabled/disabled any time using the Web Management Interface. However, to change it no active sessions can be connected to the iSCSI disk.
- Disabling write cache for an iSCSI disk does not eliminate *all* potential risk of data loss due to an unexpected loss of power as each disk drive contains its own internal cache of 8 MB or more.

Disconnect iSCSI Disk Initiators before Shutting Down the Server

Shutting down the server while a client initiator is connected to an iSCSI disk appears to the client initiator software as a disk failure and may result in data loss or corruption. Make sure any initiators connected to iSCSI disks are disconnected before shutting down the server.

Ignore Volume is Full Message

When an iSCSI disk is created, the volume or storage pool allocates the specified capacity to the disk. If all volume or storage pool capacity is allocated to the iSCSI disk and email notification is enabled, the SnapServer may generate a **Volume is Full** message. This message indicates only that the volume capacity is fully allocated to the iSCSI disk and is not available to other resources. To determine the status of iSCSI disk storage utilization, use the tools provided on the client machine.

iSCSI Disk Naming Conventions

iSCSI disks are assigned formal IQN names. These appear as the iSCSI device names that the user chooses (or types) when connecting from a client initiator to the SnapServer target, and also on the iSCSI Disk details page.

The format of IQN names for GuardianOS iSCSI disks on the SnapServer is:

```
iqn.1997-10.com.SnapServer:[servername]:[diskname]
```

where **[servername]** is the name of the SnapServer, and **[diskname]** is the name of the iSCSI disk on the target SnapServer. For example:

```
iqn.1997-10.com.SnapServer:snap123456:iscsi0
```

NOTE: Users with iSCSI disks created in earlier GuardianOS versions will see a shortened IQN name in the following format:

```
iqn.[servername].[iscsidiskname]
```

The format of IQN names for VSS-based iSCSI disks on the SnapServer is:

```
iqn.1997-10.com.SnapServer:[servername]:[diskname].[nnn]
```

where **[servername]** is the name of the SnapServer, **[diskname]** is the name of the iSCSI disk on the target SnapServer, and **[nnn]** is a sequential number starting from 000. For example:

```
iqn.1997-10.com.SnapServer:snap123456:iscsi0.000
```

The format of IQN names for VDS-based iSCSI disks on the SnapServer is:

```
iqn.1997-10.com.SnapServer:[servername]:[diskname]-snap[n]
```

where **[servername]** is the name of the SnapServer, **[diskname]** is the name of the iSCSI disk on the target SnapServer, and **[n]** is a sequential number starting from 0. For example:

```
iqn.1997-10.com.SnapServer:snap123456:iscsi0-snap0
```

Create iSCSI Disks

Navigate to **Storage > iSCSI** and click **Create iSCSI Disk** to create, edit, or delete iSCSI disks on the SnapServer. Be sure to read [“iSCSI Configuration on the SnapServer”](#) on page 6-17 before you begin creating iSCSI Disks.

ISCSI Disk Name:

Storage Pool: (Traditional RAID uses "Volume")

Size:

Enable Write Cache

Support Multiple Initiators **(Warning:** Uncontrolled simultaneous access of multiple initiators to the same iSCSI target can result in data corruption. Only enable multi-initiator support if your environment or application supports it.)

Enable CHAP Logon

NOTE: You cannot delete or edit an iSCSI disk until all clients have been disconnected from that disk.

The creation process involves first defining iSCSI parameters, then setting up security, and finally confirming your settings.

Step 1: Define the iSCSI parameters.

In the top half of the Create iSCSI Disk page, configure the new disk:

Setting Label	Description of Options
iSCSI Disk Name	Accept the default name or enter a new one. Use up to 20 alphanumeric, lowercase characters.
Storage Pool/Volume	Select the pool or volume on which to create the iSCSI disk. For Traditional RAID, if your configuration includes multiple volumes, select a volume to host the iSCSI Disk. The page refreshes, displaying the capacity of the selected volume and restoring all fields to default values.
Size	Accept the default size of the space remaining on the selected pool or volume, or enter a smaller size. NOTE: If you plan on creating VSS snapshots of the iSCSI disk, be sure to reserve some of the volume space for the iSCSI snapshot. The required Snap volume space for VSS snapshots is 10% of the size of the iSCSI disk per snapshot.

Setting Label	Description of Options
Enable Write Cache	<p>Selected by default, the write cache option significantly enhances performance. However, if a sudden, unexpected power outage occurs, some data may be lost. For more information on how to treat this option, see Write Cache Options with iSCSI Disks.</p> <p>NOTE: Disabling a write cache for an iSCSI Disk does <i>not</i> disable the write cache for any other iSCSI Disk or any other resources on the SnapServer. No active sessions can be connected to the iSCSI disk when enabling or disabling the write cache.</p>
Support Multiple Initiators	<p>Check this box if you want your iSCSI Disk to allow multiple initiator connections.</p> <p>NOTE: Data corruption is possible if this option is checked. See “iSCSI Multi-Initiator Support” on page 6-18 for more information.</p>
Enable CHAP Logon	<p>Select Enable CHAP Logon to enable CHAP authentication for access to the iSCSI Disk. Enter a user name and target secret (password). Currently, GuardianOS supports target CHAP authentication only.</p>

Step 2: If desired, enable CHAP authentication.

Check the **Enable CHAP Logon** box to display the hidden options. Enter a user name and target secret (password) twice. Both are case-sensitive.

- The user name range is 1 to 223 alphanumeric characters.
- The target secret must be a minimum of 12 and a maximum of 16 characters.

Step 3: Confirm your settings.

Click the **Create iSCSI Disk** button. At the confirmation page, verify the settings and click the **Create iSCSI Disk** button again. You are returned to the primary iSCSI page and the new iSCSI disk is displayed in the table there with the following information:

Label	Description
iSCSI Disk Name	The name of the iSCSI disk.
Storage Pool/Volume	The pool or volume on which the iSCSI disk was created.
Status	<p>Current condition of the iSCSI disk:</p> <ul style="list-style-type: none"> • OK – The iSCSI disk is online and accessible. • Not Mounted – The iSCSI disk is offline.
Active Client	The number of current sessions.
Authentication	Either CHAP or none.
Size	The size of the iSCSI disk.

Edit an iSCSI Disk

NOTE: You cannot edit an iSCSI disk if an initiator is connected. The hostname and IQN name of all connected initiators are displayed in the table.

After disconnecting all client initiators, click the iSCSI disk name in the table on the primary **iSCSI** page to display the **iSCSI Disk Properties** page. On this page, you can:

- Increase (but not decrease) the size of the iSCSI disk (if space remains).
- Enable or disable the write cache.

- Enable or disable support for multiple initiators.
- Enable or disable the CHAP logon.

Click **OK** to accept the changes (or **Close** to cancel).



CAUTION: The consistency of the internal filesystem on the iSCSI disk is primarily the responsibility of the file and operating systems on the iSCSI client used to format and manage the disk. Growing an iSCSI disk is handled differently by different operating systems and may lead to unexpected results on some client types.

Delete an iSCSI Disk

NOTE: You cannot delete an iSCSI disk if an initiator is connected. The hostname and IQN name of all connected initiators are displayed in the table.

After disconnecting all client initiators, click the iSCSI disk name in the table on the primary **iSCSI** page to display the **iSCSI Disk Properties** page. Click **Delete iSCSI Disk** (which is followed by a confirmation page) to delete the iSCSI disk.

Configuring VSS/VDS for iSCSI Disks

GuardianOS 7.2 provides VSS and VDS hardware providers to support Microsoft Volume Shadow Copy Services (VSS) and Virtual Disk Service (VDS) for iSCSI disks.

- The VSS hardware provider provides a mechanism for taking application-consistent native snapshots of iSCSI disks without performing full application (or system) shutdown. A snapshot of an iSCSI disk can be automatically created by a backup job run by a VSS-compatible backup application, so that the job backs up the snapshot volume rather than the main production volume.

NOTE: VSS iSCSI snapshots are managed by the Windows client and represent the iSCSI disk, not the Snap volume on which the iSCSI disk resides. They are not related to GuardianOS snapshots as described in "Snapshots". The VSS iSCSI snapshot rollback feature is not currently supported.

- The VDS hardware provider allows administrators to natively manage SnapServer iSCSI disks, using any VDS-compliant management console application.

SnapServers support VSS and VDS on the following platforms:

Platform	VSS	VDS
Windows Server 2003	X	
Windows Server 2003 R2	X	X
Windows Vista		X
Windows Server 2008 R2	X	X

Backing up an iSCSI Disk using VSS Snapshots. Windows VSS-compatible backup applications can create snapshots of SnapServer iSCSI disks to perform consistent backups of application data without stopping the application, using the snapshot instead of the live volume as the backup source.

NOTE: To use Symantec's Backup Exec as your VSS-compatible backup application, you must first modify the registry of the Backup Exec server and agents. For instructions, see [Using Backup Exec for VSS-based Snapshots of SnapServer iSCSI Disks](#).

Each VSS snapshot of an iSCSI target requires additional space on the pool or volume on which the iSCSI disk resides. The required space is 10% of the size of the iSCSI disk per snapshot. If this amount of free space is not available on the pool or volume, the VSS snapshot will not be created and an error will be reported by the SnapServer VSS hardware provider to the Windows event log.

When creating iSCSI disks for later VSS snapshot use, be sure to leave at least 10% of the size of the iSCSI target free on the Snap volume.

NOTE: VSS snapshots can only be taken of Windows volumes that fully consume the iSCSI disk. Snapshots of iSCSI disks that contain multiple Windows volumes are not supported.

1. Add the VSS client to the SnapServer.

- a. From the **Storage > iSCSI** page, click the **VSS/VDS Access** button.
- b. Click **Add**.
- c. Add the **hostname** of the VSS client you wish to grant access and click **Add** (the hostname is not case-sensitive).

The client hostname should appear in the VSS/VDS Clients box.

NOTE: Use the short hostname (*myclientname*) of the client only. Do not use the IP address or fully-qualified name (for example, *myclientname.mydomain.com*).

- d. When you have finished adding VSS clients, click **OK**.

2. Install the VSS hardware provider on the Windows iSCSI client.

- a. Depending on the Windows client, locate *SnapServerToolInstall32.exe* or *SnapServerToolInstall64.exe* on the Overland website:
<http://docs.overlandstorage.com/snapserver>
- b. Double-click the **executable** (.exe) to run the Installation Wizard on the VSS client and select the VSS/VDS hardware providers option. This will add the SnapServer hardware provider to the Windows iSCSI client.

3. Configure VSS-based backups of the iSCSI disk.

- a. Connect the client **iSCSI initiator** to the Snap iSCSI disk and create a volume (if necessary). Add data or configure applications to use the iSCSI volume for the data repository.
- b. Configure a VSS-based **backup** of the iSCSI disk. Where applicable, choose to use the SnapServer VSS hardware provider in the backup job configuration. When the backup job is run, the snapshot of the iSCSI disk is automatically created and hosted by the SnapServer as a virtual iSCSI disk (named after the main iSCSI disk with *snap[n]* appended), and the backup application performs the backup using the snapshot iSCSI disk. The snapshot will be deleted after the backup completes.

NOTE: VSS snapshots are not supported on SnapServer iSCSI disks that have been configured into multiple Windows volumes.

Creating and Managing iSCSI LUNs Using VDS

1. Create the **volume** and **RAID set** for the iSCSI disk on the SnapServer using the Web Management Interface (**Storage > Volumes**).

The volume and RAID set must be created on the SnapServer before the iSCSI disk can be created using a VDS application such as Microsoft's *Storage Manager for SANs*.

2. Add **VDS clients** to the SnapServer.
 - a. From the Storage > iSCSI page, click the **VSS/VDS Access** button.
 - b. Click **Add**.
 - c. Add the hostname of the VDS client you wish to grant access and click **Add** (the hostname is not case-sensitive). The client hostname should appear in the VSS/VDS Clients list.

NOTE: Use the short hostname (*myclientname*) of the client only. Do not use the IP address or fully-qualified name (for example, *myclientname.mydomain.com*).

- d. When you have finished adding VDS clients, click **OK**.
3. Install the **VDS hardware provider** on the Windows client.
 - a. Depending on the Windows client, locate *SnapServerToolInstall32.exe* or *SnapServerToolInstall64.exe* on the Overland website:
<http://docs.overlandstorage.com/snapserver>
 - b. Run the **Installation Wizard** on a VDS client and select the VSS/VDS hardware providers option. This will add the SnapServer hardware provider to the Windows client.
 4. Create and configure the **iSCSI disk** using *Storage Manager for SANs* (or other VDS-compliant application).

NOTE: RAID set terminology differs somewhat between GuardianOS and *Storage Manager for SANs*. The following table shows the equivalents:

GuardianOS RAID Set Level	Storage Manager for SANs Equivalent
0	Stripe
1	Mirror
5/6	Stripe with Parity
10	Stripe Mirror

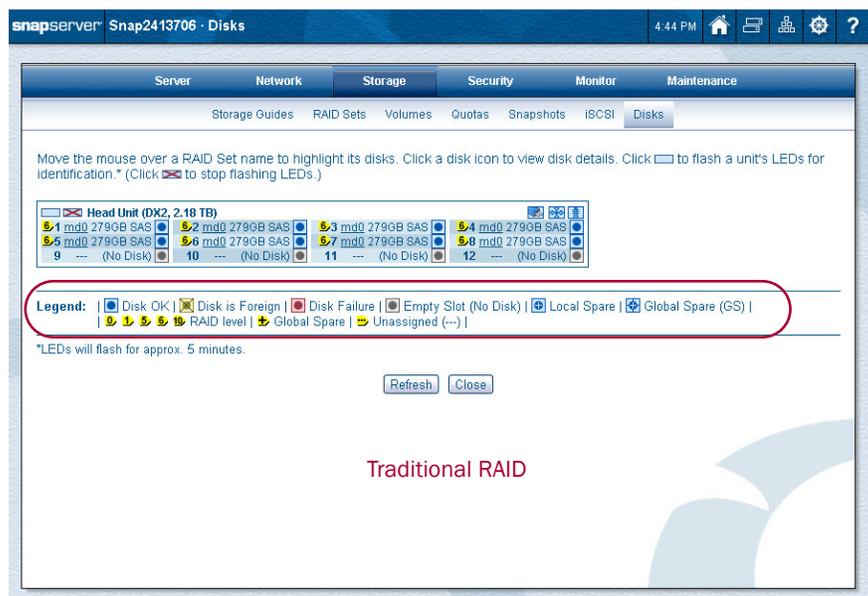
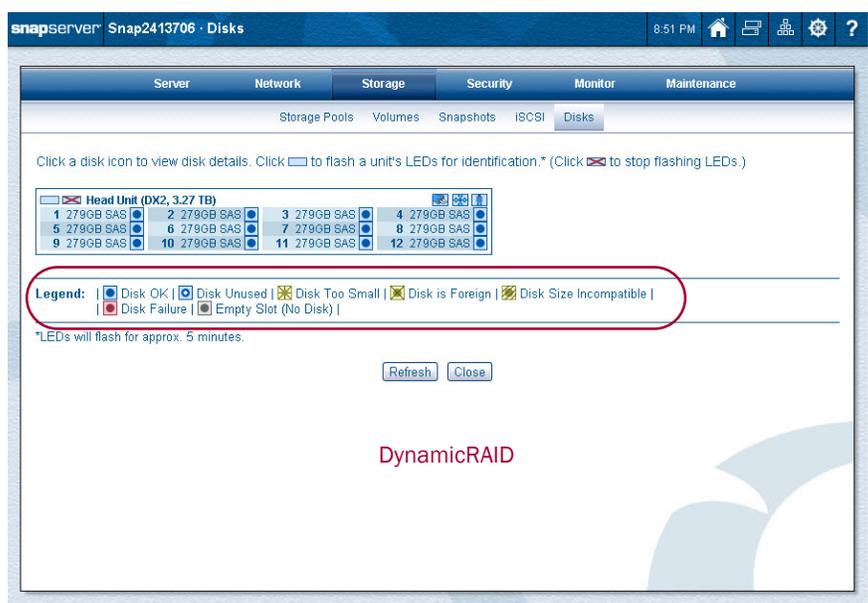
RAID set types listed in *Storage Manager for SANs* when creating an iSCSI disk reflect the types of RAID sets already configured on the SnapServer. Once a RAID set type is selected, the SnapServer automatically chooses a SnapServer RAID set of the selected type and volume to create the iSCSI disk on.

Deleting VSS/VDS Client Access

1. From the **Storage > iSCSI** page, click the **VSS/VDS Access** button.
2. Select the **VSS/VDS client** you want to delete from the VSS/VDS Clients list, and click **Delete**.
3. Click **Yes** to confirm the deletion, then click **OK**.

Disks

The Disks page is a graphic representation of the RAID set or storage pool configuration and disk status on your server. The legend on the **Storage > Disks** page explains the meaning of each icon.



- Click a disk icon (📀) to view disk details.
- Click a unit's LED icon (📶) to flash the unit's status and drive status LEDs for identification. The LEDs flash amber. Click the LED stop icon (🛑) to stop the flashing.

NOTE: The LEDs will continue to flash for five minutes unless stopped. To stop flashing LEDs for all units, click either the master LED stop icon (🛑) or link located below the legend.

- Hover the mouse over a RAID set name of one of the drives to display the RAID level next to all the disks within the RAID set (Traditional RAID only).
- Click a RAID set name to view or edit the RAID set (Traditional RAID only).

Expansion arrays, if attached to your server, will also be displayed on this page.

- If GuardianOS detects an expansion unit that is not integrated with the SnapServer, a message displays across the top of the Disks page with a link to information about the orphaned expansion unit. Also, the orphaned expansion unit will be highlighted on the page.

Replacing Disk Drives

Should a disk drive fail (solid red LED), usually it can be replaced without shutting down the SnapServer appliance (hot-swapped).

A failed disk drive can be removed and replaced anytime if two or more disks are installed in the SnapServer. However, only one disk at a time can be replaced. While dual parity allows two disks to be swapped out simultaneously, they will only be incorporated one at a time.

The following procedures assume that you are installing a new, Overland-approved disk drive as a replacement for a failed drive.

NOTE: Failed drives cannot be added back in to a RAID set.

DynamicRAID Mode

If a disk drive fails, the Administration page displays a **Disk Failure** message and an icon with a link to the Disks page. Both the Storage Pools and Storage Pool Properties pages show the degraded status. If single parity mode is being used, no parity protection message is shown. In dual parity mode, just a degraded status is shown.

NOTE: If a working disk is removed, the same changes occur as when a disk fails.

Once a disk is removed, a new disk can be inserted into any empty slot and DynamicRAID will recognize it as a replacement. The system still shows the storage as degraded but a new message appears on both the Storage Pools and Storage Pool Properties pages saying **New Disks Detected (click to repair)**. At the same time, Storage Pool Disks and Disks pages show **OK - New/Unused Disk** in that slot. To add the disk, click the repair link.

NOTE: Disk drives that have been previously configured can be added; they are indicated in the Storage > Disks list by the 🗑️ icon and a message stating that the disk has previously been used in a different system. If you want to use the drive, add it to the RAID as you would any other drive.

If there are no errors, after the new disk is incorporated any LEDs are turned off and statuses are updated.

Traditional RAID Mode

If a disk drive fails, the Administration page displays a **Disk Failure** message and an icon with a link to the Disks page. This section describes how to remove and replace drives in a RAID set of a SnapServer configured in Traditional RAID mode.

When removing a working disk drive, note the following:

- **RAID 0 (nonredundant) set** – Removing a disk drive from a RAID 0 set causes the RAID set to fail. This action renders any data residing on its drives inaccessible and is not recommended. If a RAID 0 disk drive is inadvertently removed, reinserting it should restore file access.
- **RAID 1, 5, 6, or 10 (redundant) set** – Removing a disk drive from a RAID 1, 5, 6, or 10 set places the RAID set into degraded mode. While operating in degraded mode, users can access or even update data. However, the array loses its redundant characteristics until all drives of the array are available and operating properly (except for RAID 6 set, which can tolerate a two-drive failure before it loses redundancy).

NOTE: If you configure a RAID 1, 5, 6, or 10 set with a spare, the array automatically starts rebuilding with the spare when one of the disk drives fails or is removed.

If a disk drive fails, the Traditional RAID Administration page changes to show the Disk storage as Degraded and provides a link to the RAID Sets page. Both the RAID Sets and RAID Set Properties pages show the degraded status.

NOTE: If a working disk is removed, the same changes occur as when a disk fails.

After a fresh drive is inserted, if auto-incorporation is not enabled, you must use the Web Management Interface to add it to a RAID set:

1. Go to **Storage > RAID Sets** and click the RAID Set **name**.
2. Click the **Repair RAID** button.
3. Select a drive from the list shown, and click **Repair RAID** again to incorporate it into the RAID as a replacement for a failed member drive.

The screenshot shows the SnapServer RAID Set Properties page for RAID Set md1. The RAID set is in a Degraded state. The page displays a table of disks for the RAID set, showing two 20GB SAS disks in OK status. The 'Repair RAID' button is highlighted with a red box.

RAID Set	Status	Group	Size	Unallocated
md1	Degraded Members (active / configured): 2 / 3 Spares (active / configured): 0 / 0	group0	15.49 GB	-

Disks for RAID set md1. (*Note: Usable Space specifies the amount of disk space usable by the RAID set.)

Disk	Location	Status	Usable Space*	Action
<input checked="" type="checkbox"/> 20GB SAS	Head Unit, disk 6	OK	7.74 GB	-
<input checked="" type="checkbox"/> 20GB SAS	Head Unit, disk 7	OK	7.74 GB	-

Number of global spares: 0

Refresh Repair RAID Close

Note: This RAID set is a member of a RAID Group and cannot be deleted individually. To delete this RAID set, you must delete the RAID Group group0.

NOTE: The Repair RAID button only appears when a drive has failed or been removed, and the RAID is in degraded mode.

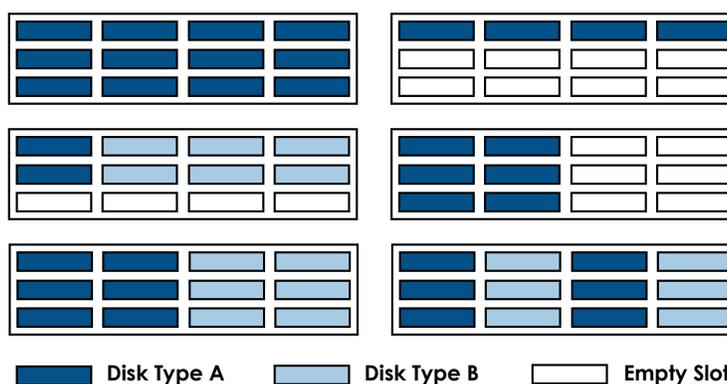
The RAID set status changes to **Resyncing** while the new drive is incorporated into the RAID set. It reads **OK** once the incorporation is complete.

Adding Disk Drives

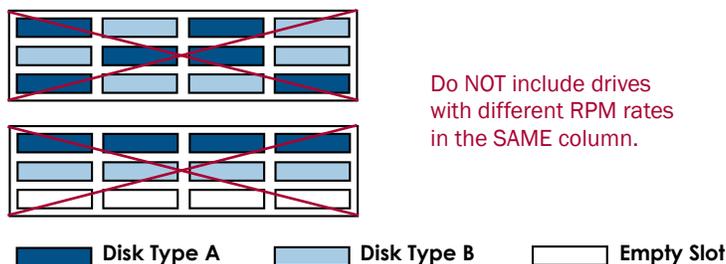
If empty slots are available, you can add an Overland-approved disk drive to expand the storage pool/volume on your SnapServer or SnapExpansion unit.

Drives of different rotational speed (for example, SAS and SATA drives) can be combined in the same server. However, they cannot be combined in the same column, and it is recommended that columns of same-type drives be grouped together. If you are combining drives with different rotational speeds, use the figures below to plan where to place the disk drives.

Recommended Disk Drive Configurations



Unsupported Disk Drive Combinations



DynamicRAID Mode

When adding additional disk drives, keep the following in mind:

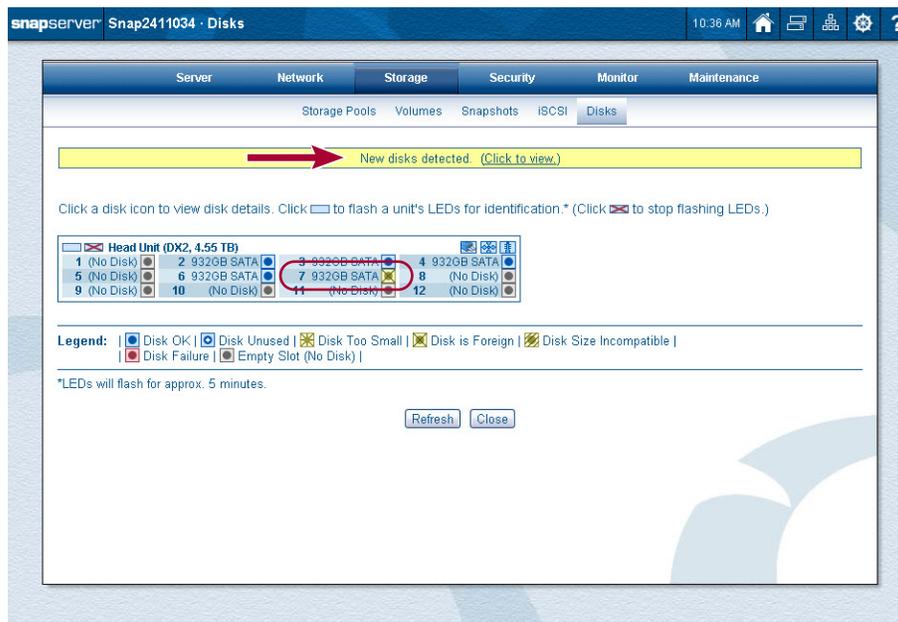
- While disk sizes within a Storage Pool can vary, the type of disks used must be the same (such as, SAS 7200 RPM or SAS 15K RPM).
- If a non-compatible disk of a different partition size is added to a Storage Pool, it is indicated in the **Storage > Disks** list by the  icon.
- If only a single disk is in a Storage Pool, the second disk added must be of equal or greater size.

- A move from dual parity to single parity is allowed at any time, provided the storage pool is healthy. A move from single parity to dual parity is only allowed when a new disk drive is added that is large enough to support the new parity mode.

To add a new disk drive to a DynamicRAID:

1. Insert the **drive** into an empty SnapServer slot.

It appears in the **Storage > Disks** map as “Disk Unused” (for a new disk) or “Disk is Foreign” (reused, clean disk). A “New disk detected” banner is shown.



NOTE: Disk drives that have been previously configured can be added; they are indicated in the Storage > Disks list by the  icon (Disk is Foreign) and a message stating that the disk has previously been used in a different system and all data will be deleted. If you want to use the drive, continue to add it to the RAID as you would any other drive.

2. Click the **link** in the banner (or navigate to the **Storage > Storage Pools** page).
3. In the **Status** column, click the link to open the Storage Pool Properties page.

- At the bottom of the properties page, verify or change the **Parity Mode** and **Snapshot Pool** settings.

Storage Pool Properties for Head Unit Storage. The interface shows a table for Storage Pools with columns for Storage Pool, Status, Parity, Data Pool Usage, and Snapshot Pool Usage. Below this, a table lists detected disks. The Parity Mode is set to Dual-parity protection, and the Snapshot Pool is set to 20%.

Storage Pool	Status	Parity	Data Pool Usage	Snapshot Pool Usage
Head Unit Storage Head Unit Model: DX2	OK	Dual	<1% (133 MB / 1.25 TB)	0% (0 MB / 320.57 GB)
Pending Changes: (Click OK to review & confirm)			Data Pool Usage (New)	Snapshot Pool Usage (New)
· Incorporate 1 new disk.			<1% (133 MB / 1.46 TB)	0% (0 MB / 374 GB)

Disk	Slot	Status
279GB SAS	9	OK - New / unused disk.

Storage Pool Name: Head Unit Storage

Parity Mode

- Single-parity protection - protects your data in the event of a single disk failure.
- Dual-parity protection - uses more disk space than single-parity, yet protects your data in the event of up to 2 disk failures.

Snapshot Pool

Percentage of this Storage Pool reserved for snapshots: 20%

Buttons: OK, Refresh, View Disks, Delete Storage Pool, Cancel

(Click OK to review and confirm your changes.)

- Click **OK** to continue.
- At the confirmation page, click **Save Changes**.
- At the success notice page, click **OK**.

To speed up the storage pool synchronization, it is recommended that you exit the Web Management Interface.

Storage Pools page for Snap2411034. The interface shows a table with one storage pool entry. The status indicates a resync operation is in progress.

Storage Pool	Status	Parity	Data Pool Usage	Snapshot Pool Usage
Head Unit Storage Head Unit Model: DX2	Resync: 0% complete. (Estimated size after operation completes: 2.87 TB)	Single	<1% (133 MB / 2.15 TB)	0% (0 MB / 551.43 GB)

Buttons: Refresh, Close

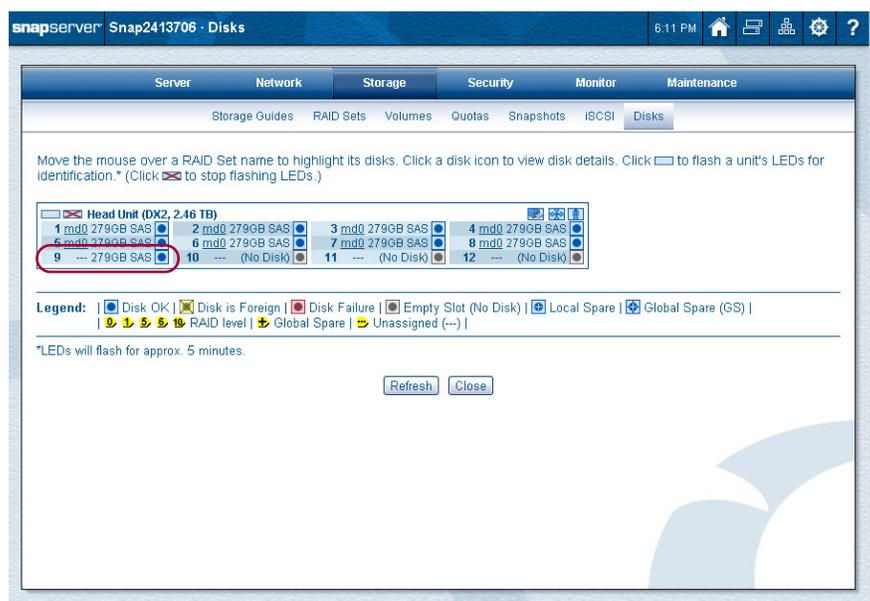
Traditional RAID Mode

This section describes how to safely add drives to an existing RAID 1, 5, 6, or 10 set. On SnapServers, after a fresh drive is inserted into a drive bay, if auto-incorporate is not enabled (see “[Automatic Incorporation of Hot-Swapped Drives](#)”), you must use the Web Management Interface (**Storage > RAID Sets**) to add it to a RAID set.

- **RAID 0 set (nonredundant)** – You cannot add a drive to a RAID 0 set. To reconfigure a RAID 0 set, you must delete the RAID set and then recreate it.
- **RAID 1 set (redundant)** – You can add a new drive to a RAID 1 set as either a spare or as a new member. Adding a disk drive to a RAID 1 set does not add storage capacity. The new member simply creates an additional copy of the original drive.
- **RAID 5, 6, or 10 set (redundant)** – You can add a spare to a RAID 5, 6, or 10 set. However, you cannot add a new drive as a new member.

To add a new disk drive as a Local Spare for a **Traditional RAID** set:

1. Insert the **drive** into an empty SnapServer slot.
It appears in the **Storage > Disks** map as “Disk Unused” (for a new disk) or “Disk is Foreign” (reused, clean disk).



NOTE: Disk drives that have been previously configured as part of a RAID can be added. They are indicated in the Storage > Disks list by the  icon and a message stating that the disk has previously been used in a different system (foreign). If you want to use the drive, add it to the RAID as you would any other drive.

2. Navigate to the **Storage > RAID Sets** page.
3. Click the **name** of the RAID set to which you want to add a drive.
4. On the RAID set page that opens, click **Add Disk**.
If you are adding to a RAID 1 set, select either **Spare** or **Member** at the top of the page.
5. From the Available Disks list, select one or more **drives** to add to the configuration, and then click **Next**.
6. On the confirmation page, click **Add Disk**.
The disk is added as a Local Spare to the selected RAID set.

To add a new disk drive as a Global Spare for a **Traditional RAID** set:

1. Insert the **drive** into an empty SnapServer slot.
It appears in the **Storage > Disks** map as “Disk Unused” (for a new disk) or “Disk is Foreign” (reused, clean disk).

NOTE: Disk drives that have been previously configured as part of a RAID can be added. They are indicated in the Storage > Disks list by the  icon and a message stating that the disk has previously been used in a different system (foreign). If you want to use the drive, add it to the RAID as you would any other drive.

2. Navigate to the **Storage > RAID Sets** page.
3. Click **Global Spares**.
4. From the Available Disks list, select one or more **drives** to add to as Global Spares, and then click **OK**.

The disk is added as a Global Spare.

If you don't want to use a drive as a Global Spare, repeat the process and uncheck the box on the Available Disks list.

Reintegrate Orphaned Disk Drives

An orphaned disk drive occurs in the following circumstances: (1) a working drive from a RAID set is accidentally removed from the server; or (2) the RAID set or system is started with a drive missing. In either case, the drive becomes suspect and is considered an orphan. To remedy the problem, click the RAID set name in the **Storage > RAID Sets** page, then click the **Repair** link next to the drive in question.

Expansion Units

To increase the capacity of a SnapServer DX-series server, use one or more SnapExpansion units.



NOTE: If GuardianOS detects an expansion unit that is not integrated with the SnapServer, a message is displayed across the top of the Disks, Admin Home, RAID, and Storage Pool pages with a link to information about the orphaned expansion unit.

The SnapExpansion is a 2U, 12-bay SAS-connected expansion unit for expanding either SnapServer DX1 or DX2 systems. The SnapExpansion supports SAS and SATA hard drives and allows capacity growth by creating additional DynamicRAID storage pools.

- By adding the optional SAS expansion card to the SnapServer DX1, you can attach up to **three** SnapExpansion chassis to grow storage up to 120TB.
- The SnapServer DX2 includes the SAS expansion card that allows for **seven** SnapExpansion chassis to be attached, to scale up to 288TB in total space.

A SnapExpansion is accessed and managed through the SnapServer head unit to which it is connected. The expansion unit has no physical connection to the network.

Managing Expansion Unit Storage

The **Storage > Disks** page displays the head unit and any expansion units attached to the head unit. For more information about the Disks page, please see “Disks” on page 6-25.

Click a disk icon to view disk details. Click to flash a unit's LEDs for identification.* (Click to stop flashing LEDs.)

Head Unit (DX1, 17.58 TB)											
1	1000GB SATA	2	1000GB SATA	3	2000GB SATA	4	2000GB SATA	5	1000GB SATA	6	1000GB SATA
7	2000GB SATA	8	2000GB SATA	9	1000GB SATA	10	1000GB SATA	11	2000GB SATA	12	2000GB SATA

Expansion Unit 1 (E2000, 10.26 TB)											
1	1501GB SAS	2	1501GB SAS	3	1501GB SAS	4	1501GB SAS	5	1501GB SAS	6	1501GB SAS
7	1501GB SAS	8	(No Disk)	9	(No Disk)	10	(No Disk)	11	(No Disk)	12	(No Disk)

Expansion Unit 2 (E2000, 9.77 TB)											
1	501GB SAS	2	501GB SAS	3	501GB SAS	4	501GB SAS	5	1501GB SAS	6	1501GB SAS
7	1501GB SAS	8	1501GB SAS	9	501GB SAS	10	501GB SAS	11	501GB SAS	12	501GB SAS

Expansion Unit 3 (E2000, 0 GB)											
1	(No Disk)	2	(No Disk)	3	(No Disk)	4	(No Disk)	5	(No Disk)	6	(No Disk)
7	(No Disk)	8	(No Disk)	9	(No Disk)	10	(No Disk)	11	(No Disk)	12	(No Disk)

Expansion Unit 4 (E2000, 3.52 TB)											
1	1800GB SAS	2	1800GB SAS	3	(No Disk)	4	(No Disk)	5	(No Disk)	6	(No Disk)
7	(No Disk)	8	(No Disk)	9	(No Disk)	10	(No Disk)	11	(No Disk)	12	(No Disk)

Expansion Unit 5 (E2000, 5.86 TB)											
1	2001GB SAS	2	2001GB SAS	3	2001GB SAS	4	(No Disk)	5	(No Disk)	6	(No Disk)
7	(No Disk)	8	(No Disk)	9	(No Disk)	10	(No Disk)	11	(No Disk)	12	(No Disk)

Expansion Unit 6 (E2000, 2.44 TB)											
1	2501GB SATA	2	(No Disk)	3	(No Disk)	4	(No Disk)	5	(No Disk)	6	(No Disk)
7	(No Disk)	8	(No Disk)	9	(No Disk)	10	(No Disk)	11	(No Disk)	12	(No Disk)

Legend: Disk OK Disk Unused Disk Too Small Disk is Foreign Disk Failure Format Slot (No Disk)

The disk drives of expansion units are completely integrated into the head unit's logic. Their access is determined by the type of RAID system being used.

DynamicRAID

Each unit in the SnapServer system has its own storage pool which DynamicRAID manages.

Traditional RAID

The head unit and expansion disk drives can be combined as necessary. For example, to create one large RAID set, you could delete the existing RAID sets on both the head unit and the expansion unit, then combine all drives into one high-capacity storage system.

This configuration of one large RAID reduces administrative complexity and overhead, but the failure of any one unit in the system (due to a cable coming loose, for example) will render the entire RAID set inaccessible. This configuration also increases the potential for multiple drive failures in a single RAID set.

Integrating Orphaned Expansion Units

Expansion units that have been discovered by GuardianOS (for example, are physically connected to the SnapServer) but have not been integrated with the SnapServer are listed on the Orphan Expansion Units Table accessed by the message link at the top of the Disk, Admin Home, RAID, and Storage Pools pages.

Property	Description
Expansion Unit	A description of the unit
Status	The status of the unit (for example, orphan)
Serial Number	The expansion unit's serial number
Origin	The serial number of the server with which the expansion unit was last incorporated

If you want to use the expansion unit with the SnapServer, click the checkbox next to the orphaned expansion unit you want to integrate, and click **OK**.



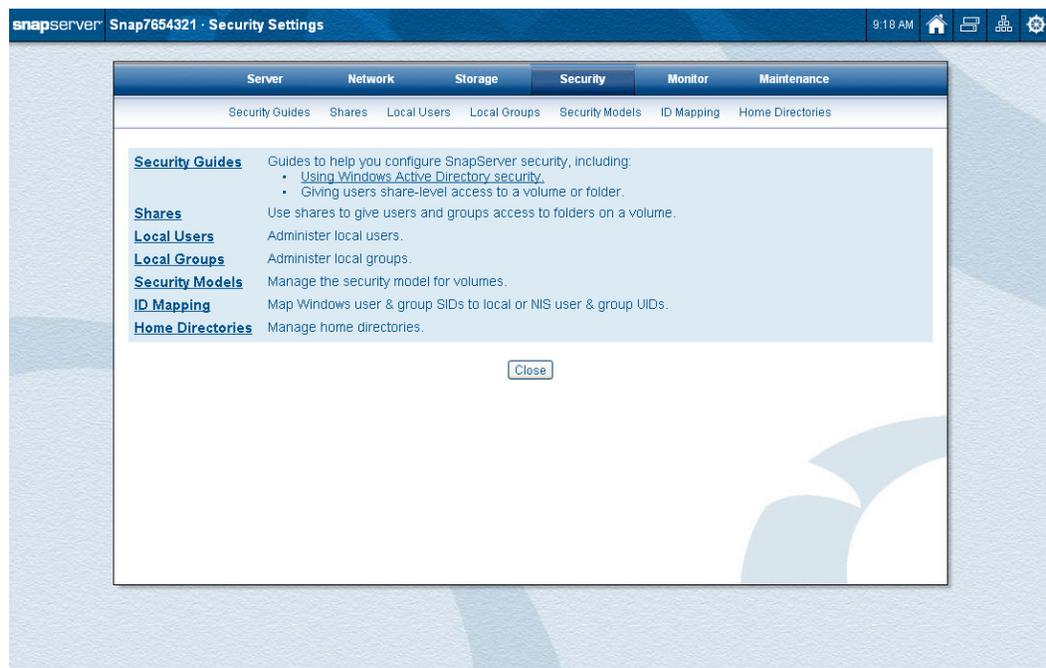
CAUTION: Before integrating an orphaned expansion unit, be sure that it is compatible with the SnapServer (for example, data on the expansion unit is compatible with the SnapServer configuration, Unicode settings are the same, etc.).

This section covers Security options for users, groups, shares, and file access.

Topics in Security Options

- [Overview](#)
- [Security Guides](#)
- [Shares](#)
- [Local Users](#)
- [Local Groups](#)
- [Security Models](#)
- [ID Mapping](#)
- [Home Directories](#)

Overview



Authentication validates a user's identity by requiring the user to provide a registered login name (User ID) and corresponding password. SnapServer appliances ship with predefined local users and groups that allow administrative (admin) and guest user access to the server via all protocols.

Administrators may choose to join the SnapServer to a Windows Active Directory domain, and CIFS/SMB and AFP clients can then authenticate to the server using their domain credentials. To accommodate NFS clients, the SnapServer can also join an NIS domain, and the SnapServer can look up user IDs (UIDs) and group IDs (GIDs) maintained by the domain. For authentication control beyond the guest account, Mac and FTP client login credentials can be created locally on the server. See “[User and Group ID Assignments](#)”.

SnapServer default security configuration provides one share to the entire volume. All network protocols for the share are enabled, and all users are granted read-write permission to the share via the guest account. By default, the `guest` user is disabled in SMB.

Network clients can initially access the server using the guest account, but if you require a higher degree of control over individual access to the filesystem for these clients, you must create local accounts (or use Windows Active Directory security for CIFS/SMB and AFP clients).

Local users or groups are created using the **Security > Local Users** and **Security > Local Groups** pages in the Web Management Interface. Local users are also used for administrative access to the server through the server's Web Management Interface, SnapServer Manager, or SSH.

A local user or group is one that is defined locally on the SnapServer using the Web Management Interface. The default users and groups listed below cannot be modified or deleted.

- **admin** – The local user admin account is used to log into the Web Management Interface. The default password for the admin account is also `admin`.
- **guest** – The local user guest account requires no password.
- **admingrp** – The Admin group account includes the default admin user account. Any local user accounts created with admin rights are also automatically added to this group.

Guidelines for Local Authentication

These password authentication guidelines are for both users and groups.

Duplicating Client Login Credentials for Local Users and Groups. To simplify user access for Windows Workgroup or Mac clients, duplicate their local client logon credentials on the SnapServer by creating local accounts on the SnapServer that match those used to log on to client workstations. This strategy allows users to bypass the login procedure when accessing the SnapServer.



CAUTION: This strategy applies only to local users. Do not use duplicate domain user credentials if joined to an Active Directory domain.

Default Local Users and Groups . Default users and groups `admin`, `guest`, and `admingrp` appear on the list of users or groups on the User or Group Management pages, but they cannot be deleted or modified (although the admin password can be changed). The default local users and groups do appear on the **Share Access** and **Quotas** pages.

Changing Local UIDs or GIDs. The SnapServer automatically assigns and manages UIDs and GIDs. Because you may need to assign a specific ID to a local user or group in order to match your existing UID/GID assignments, the SnapServer makes these fields editable.

Password Policies. To provide additional authentication security, set password character requirements, password expiration dates, and lockout rules for local users.

Local users can also be individually exempted from password expiration and character requirement policies. The built-in *admin* user is exempt from all password policies.

Local Account Management Tools. The SnapServer offers the following tools for creating, modifying, and editing local user and group accounts.

Function	Navigation Path
Local User Management	Navigate to the Security > Local Users page, from which you can create, view, edit, and delete local users. You can also set user password policy, including password character requirements, maximum number of allowed logon failures, and password expiration settings.
Local Group Management	Navigate to the Security > Local Groups page, from which you can create, view, edit, and delete local groups.

User and Group ID Assignments

The SnapServer uses the POSIX standard to assign UIDs or GIDs, in which each user and group must have a unique ID. This requirement applies to all users and groups on the SnapServer, including NIS, Windows domain, and local users and NIS groups.

If you join the SnapServer to a Windows or NIS domain, IDs are assigned using available IDs only. Consider the following when creating users and groups:

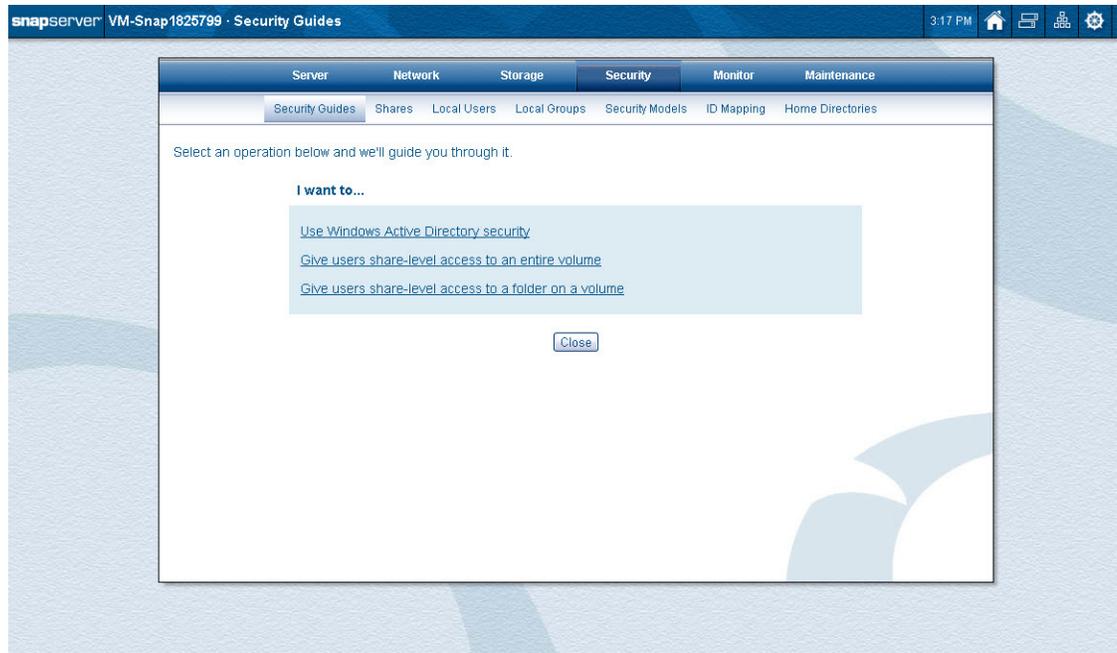
- UIDs and GIDs from 0 to 100 are unavailable for use. If you try to assign a UID or GID that is less than 101 (or in use by NIS or the Windows domain), you will get an error message.
- When the server automatically generates UIDs or GIDs for imported Windows domain users or groups, UIDs or GIDs that are already in use by local and NIS users will be skipped.
- When NIS domain users and groups are imported, the SnapServer will discard any UIDs that are less than 101 or are in conflict with UIDs already in use by local or Windows domain users and groups.

The `nfsnobody` and `nobody` user IDs (UID 65534 and 65535, respectively) and GIDs are reserved. They are not mappable to other IDs, nor is another ID mappable to `nfsnobody` or `nobody`.

Security Guides

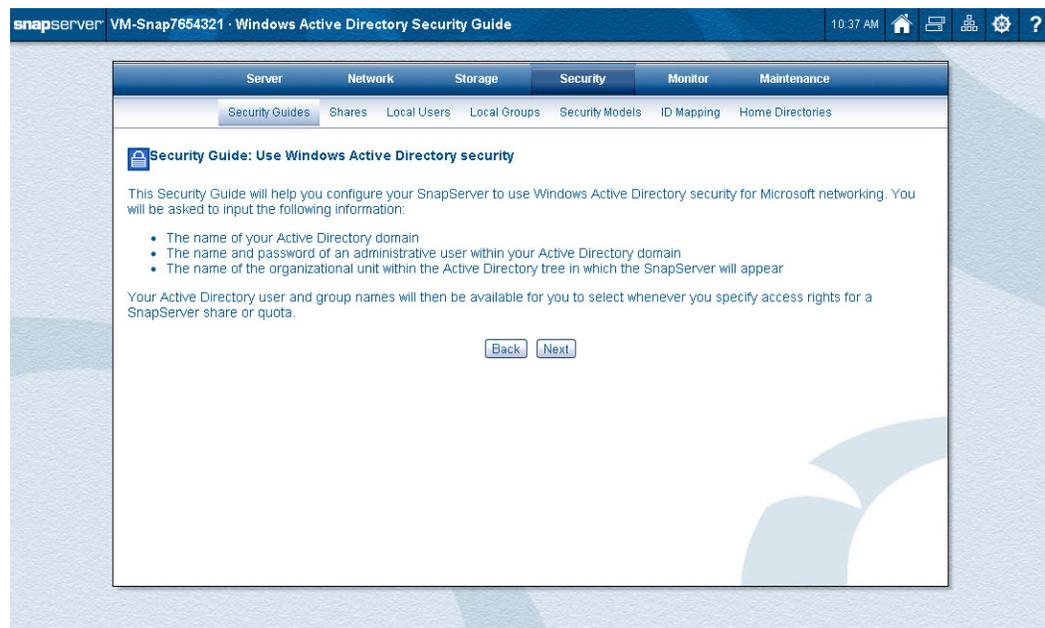
Security Guides are special wizards to guide you through:

- Setting up Windows Active Directory security.
- Giving users or groups share-level access to an entire volume.
- Giving users or groups share-level access to a folder on a volume.



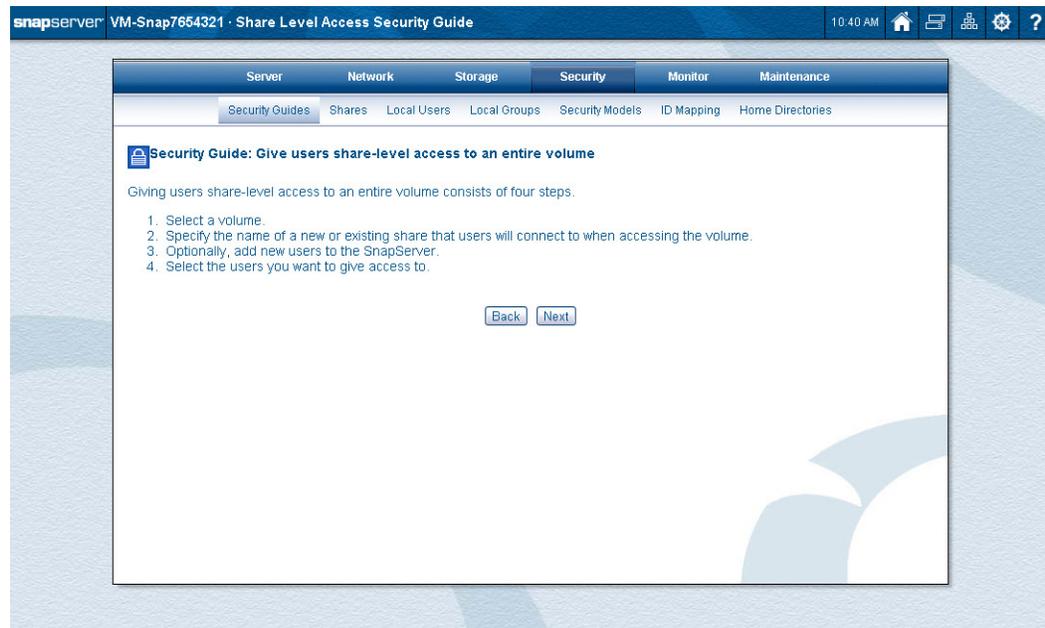
Windows Active Directory Security Guide

This wizard guides you through the setup of Windows Active Directory on your server.



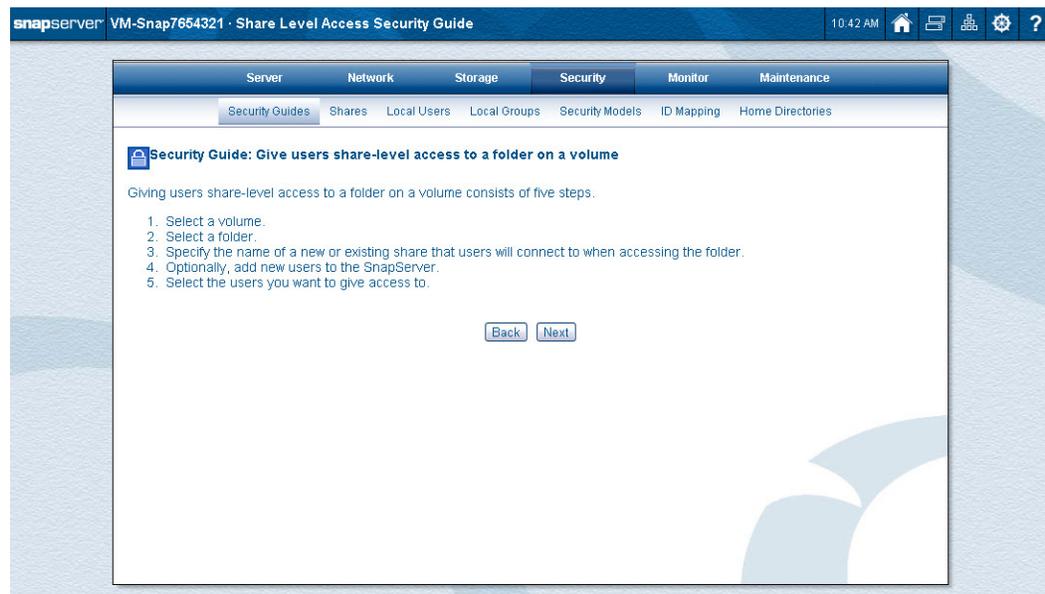
Entire Volume Security Guide

This wizard guides you through the four steps it takes to give share-level access to a folder on a volume.



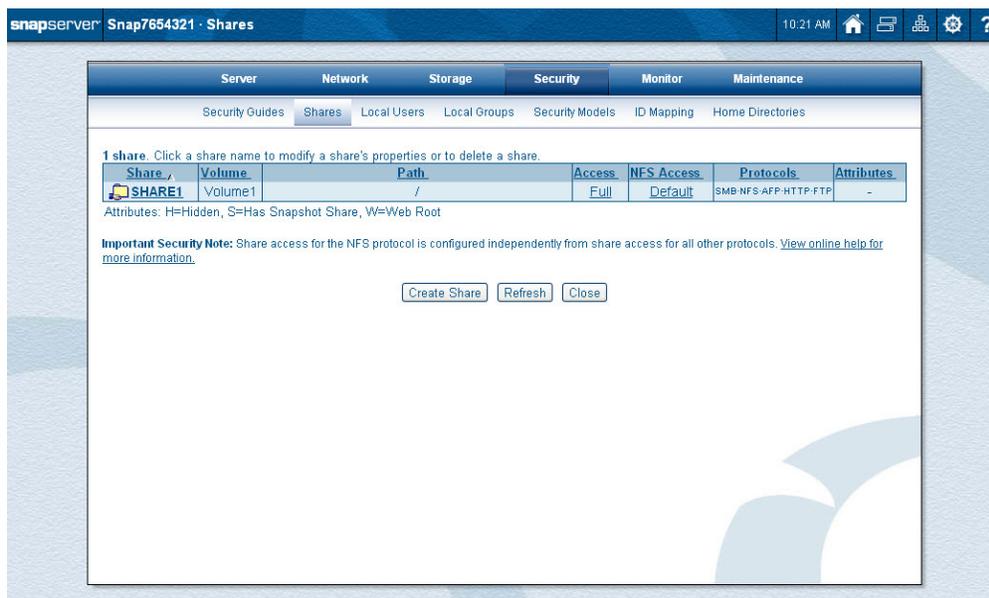
Folder on Volume Security Guide

This wizard guides you through the five steps it takes to give share-level access to an entire volume.



Shares

SnapServer provides full integration with existing Windows Active Directory Domain or UNIX NIS user and group databases. At the share level, administrators can assign read-write or read-only share access to individual Windows (and local) users and groups. Administrators can also edit the NFS exports file to control how shares are exported to NFS client machines.



Share Security Overview

SnapServers support file access in Windows, UNIX, and Apple networks, as well as access via FTP and HTTP. New shares are created by default with full read-write access to all users, subject to the filesystem permissions on the share target directory (see “[File-level Security](#)”). The first step to securing a SnapServer is to specify access at the individual share level. Administrators can assign read-write or read-only share access to individual Windows (and local) users and groups.

For more details and information about using Shares, see [Appendix C, “Security and Access.”](#)

Create Shares

To create a new share, you need, at a minimum, to specify the share name, volume, and folder path.

By clicking the Advanced Share Properties link, additional options are displayed. Use these options to hide the share from network browsing, select the protocols supported, and create a snapshot share associated with this share.

Creating a Share

Creating a share involves selecting the volume, security model, and directory path for the share and then defining share attributes and network access protocols.

1. Accept the default **share name** or enter a new one.
To ensure compatibility with all protocols, share names are limited to 27 alphanumeric characters (including spaces). Choose the volume you need from the drop-down menu.
2. Select from the following **options**:
 - **To create a share to the entire volume** – The current Path field defaults to the root path of the volume. Simply leave it blank if this is the desired configuration.
 - **To create a share to a folder on the volume** – Browse to the folder you want to which you want to point the share, click the folder name, and click OK.

NOTE: If you want to create a new folder inside any other folder, type the folder name into New Folder Name and click Create Folder.

3. If desired, enter a **description** to clarify the purpose of the share.

- Choose a **security model** by selecting either **windows/mix** or **UNIX** from the drop-down list.

The **security model** option is only available under the following circumstances:

- Traditional RAID – When pointing the share at the root of a volume or one directory down from the root of the volume.
- DynamicRAID – When pointing the share to the root of a volume.

If available, the option defaults to the current security model at the specified path. If changed to a different security model, the change will propagate to all files and subdirectories underneath. For more information, see [“Security Models” on page 7-25](#).

- Choose the user-based **Share Access** option desired.

Choose either **Create share with full read and write access for all users**, or **Create share with Admin-only access and proceed to Share Access page** to configure the share access. For more information, see [“Share Access Behaviors” on page 7-11](#).

NOTE: If selecting **Create share with Admin-only access** and if the share has NFS enabled, be sure to configure the **NFS Access settings** afterward.

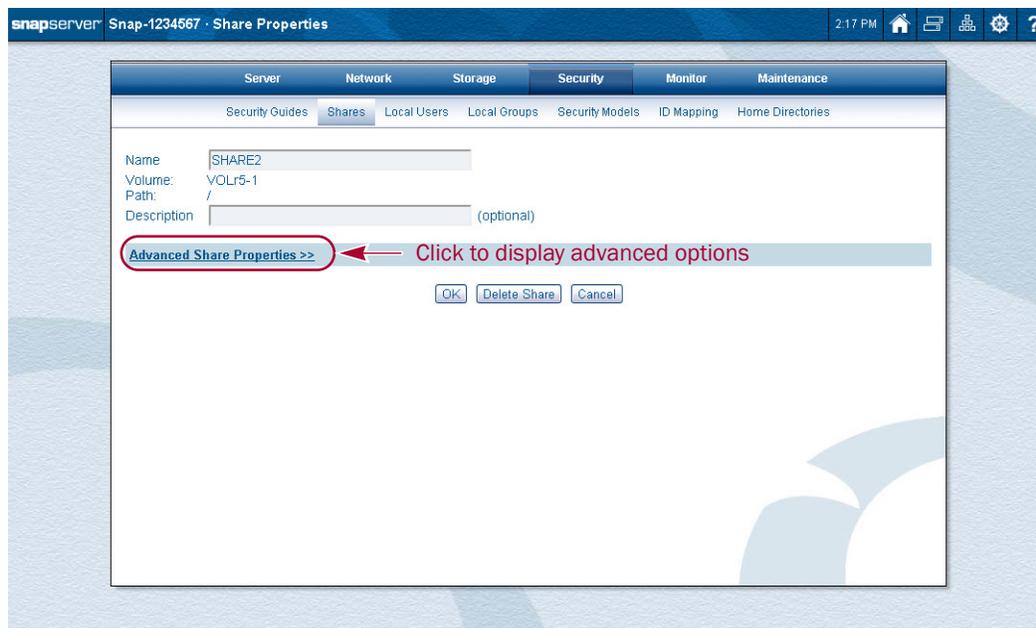
- To further configure the share, click **Advanced Share Properties**, and enter any of the following:

Option	Description
Hide this Share	Select this option if you want the share to be hidden from network browsing using SMB, HTTP/HTTPS, AFP, and FTP (but not NFS) protocols.
Protocols	Select the access protocols for the share: Windows (SMB), Linux/UNIX (NFS), Apple (AFP), Web View (HTTP/HTTPS), and FTP/FTPS.
Snapshot Share	To create a snapshot share, select the Create Snapshot Share checkbox. Optionally, do either of the following: <ul style="list-style-type: none"> If desired, enter a unique name for the Snapshot Share Name field. Use up to 27 alphanumeric characters (including hyphens and spaces). To hide the snapshot share from the SMB, HTTP, AFP, and FTP protocols, select the Hide Snapshot Share checkbox.

- Click **Create Share** to complete the process.

Edit Share Properties

Once a share has been created, you can change its name, description and the advanced properties. To edit the properties, go to **Security > Shares > Share Properties** (displayed by clicking the share name).



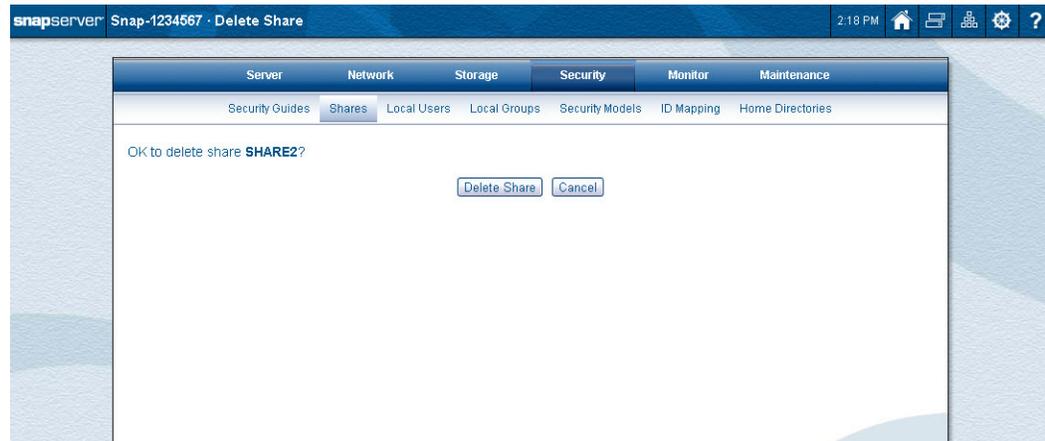
You cannot change the volume (or path). If you need to change the volume, you must delete the share and create a new one on the other volume.

Option	Description
Name	Accept the default share name or enter a new one. If you change the default, observe the following guidelines: <ul style="list-style-type: none"> • Make sure the share name is unique to this server • To ensure compatibility with all protocols, share names are limited to 27 alphanumeric characters (including hyphens and spaces).
Description	If desired, enter a description of the share. This is an opportunity to clarify the purpose of the share.
Hide this share	Select this option if you want the share to be hidden from network browsing.
Protocols	Select the access protocols for the share: Windows (SMB), Linux/UNIX (NFS), Apple (AFP), Web (HTTP, HTTPS), FTP, FTPS.
Snapshot Share	The option that displays depends on whether a snapshot share currently exists. <p>To create a snapshot share, select the Create Snapshot Share checkbox.</p> <ul style="list-style-type: none"> • If desired, enter a unique name for the Snapshot Share Name field. Use up to 27 alphanumeric characters (including hyphens and spaces). • To hide the snapshot share from the SMB, HTTP, AFP, and FTP protocols, select the Hide Snapshot Share checkbox. <p>To remove a snapshot share, do the following:</p> <ul style="list-style-type: none"> • Select the Remove Snapshot Share checkbox.

Delete Shares

To delete a share, go to **Security > Shares > Share Properties** (displayed by clicking the share name).

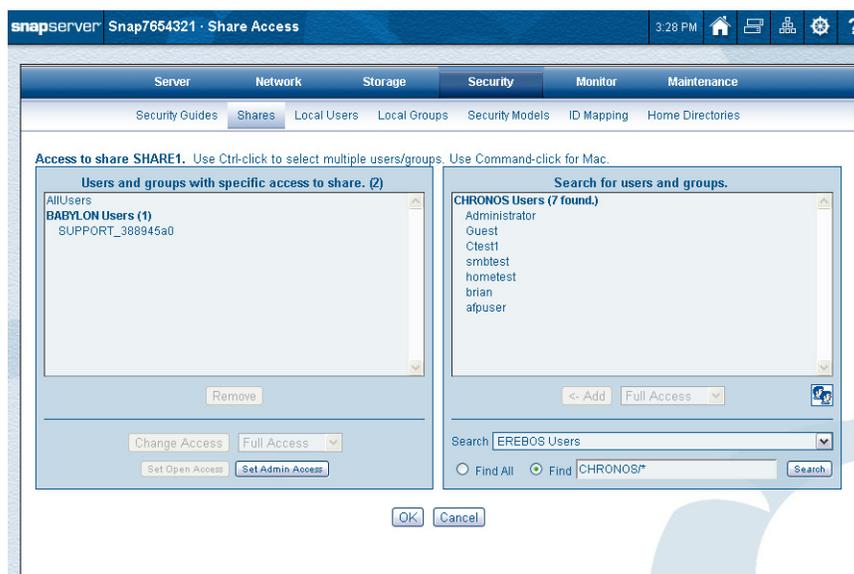
1. Click the **Delete Share** button at the bottom.
2. At the Delete Share confirmation page, click the **Delete Share** button again.



Configuring Share Access

In **Security > Shares**, in the **Access** column, click the link next to the share you want to configure. The Share Access page displays. You can set access levels for the share, as well as grant or deny access to specific users and groups.

NOTE: To add a new user to a share, you must first create the user, then add that user to the share. Please see “Local Users” on page 7-17 for information on creating new users.



Share Access Behaviors

Administrators tasked with devising security policies for the SnapServer will find the following share access behaviors informative:

- **Share access defaults to full control** – The default permission granted to users and groups when they are granted access to the share is full control. You may restrict selected users and groups to read-only access.
- **User-based share access permissions are cumulative** – An SMB, AFP, HTTP, or FTP user's effective permissions for a resource are the sum of the permissions that you assign to the individual user account and to all of the groups to which the user belongs in the Share Access page. For example, if a user has read-only permission to the share, but is also a member of a group that has been given full-access permission to the share, the user gets full access to the share.
- **NFS access permissions are not cumulative** – An NFS user's access level is based on the permission in the NFS access list that most specifically applies. For example, if a user connects to a share over NFS from IP address 192.168.0.1, and the NFS access for the share gives read-write access to * (All NFS clients) and read-only access to 192.168.0.1, the user will get read-only access.

- **Interaction between share-level and file-level access permissions** – When both share-level and file-level permissions apply to a user action, the more restrictive of the two applies (see “[File-level Security](#)”). Consider the following examples:

Example A: More restrictive file-level access is given precedence over more permissive share-level access.

Share Level	File Level	Result
Full control	Read-only to File A	Full control over all directories and files in SHARE1 <i>except</i> where a more restrictive file-level permission applies. The user has read-only access to File A.

Example B: More restrictive share-level access is given precedence over more permissive file-level access.

Share Level	File Level	Result
Read-only	Full control to File B	Read-only access to all directories and files in SHARE1, <i>including</i> where a less restrictive file-level permission applies. The user has read-only access to File B.

Setting User-based Share Access Permissions

Share permissions for Windows, Apple, FTP, and HTTP users are configured from **Security > Shares** by clicking the link in the **Access** column of the share you want to configure. Share permissions for NFS are configured and enforced independently. See “[NFS Access for Shares](#)” on page 7-14 for more information.



User-based share access permissions apply to users connecting over SMB, AFP, HTTP, and FTP. Users and groups with assigned share access permissions appear in the list on the left (*Users and groups with specific access to share*). To search for those without assigned access, use the box on the right (*Search for users and groups*).

The default permission granted to users and groups when they are granted access to the share is Full access. You may restrict selected users and groups to Read-only access.

Share-Level Access Permissions	
Full access	Users can read, write, modify, create, or delete files and folders within the share.
Read-only	Users can navigate the share directory structure and view files.

1. Display the **Share Access** page (**Security > Shares > access_link**).
2. To **add** share access permissions for a user or group:
 - a. At the bottom, using the drop-down list, select the **domain** or **local user/group list** to search.

NOTE: For domains that require authentication (showing an "(A)" after the name), after selecting the domain name, enter the User Name and Password for that domain. The user name and password can be for any user in the domain and are used to retrieve basic information (like the user & group lists) from the domain.

- b. Enter the **search string** (or select Find All).

When entering a search string:

- Returned results will include all users and groups whose name **begins** with the string entered in the Search field.
- The search results returned may be limited. Fine tune your search by using a more specific string to return the names desired.
- On the rare occasion you need to search for a domain that's not listed ("remote domain"), select a domain from the Search drop-down list through which to search, then enter in the Find box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, **remote_domain\user_name**). After you click Search, another authentication prompt may be presented to authenticate with the remote domain.

- c. Click **Search** to display any matches.

After you click Search, another authentication prompt may be presented to authenticate with the remote domain.

- d. Select one or more **names** in the list.

Users that already have access are shown in purple font with a plus sign (+) in front of their name.

- e. Choose either **Full Access** or **Read Only** from the drop-down list.
 - f. Click **Add**.

NOTE: To display recent user or group picks, click the faces  icon. A list with a green background is displayed. Click the now green icon to return to the normal search box.



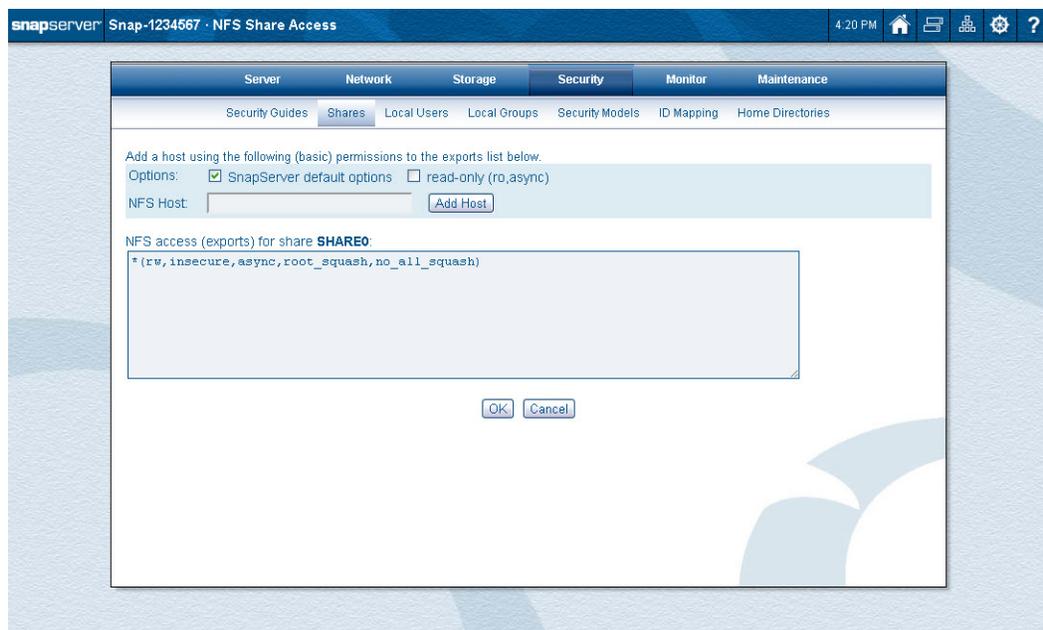
3. To **remove** share access permissions for a user or group:
 - a. Select one or more **users or groups** in the left box.
 - b. Click **Remove**.
4. To change **access permissions** for a user or group, select one or more users or groups in the left box, then select either **Full Access** or **Read Only** from the drop-down list, and click the **Change Access** button.
5. To quickly specify either Open or Admin-only **access** for the entire share, click either the **Set Open Access** or **Set Admin Access** button.
6. Click **OK** to save share permissions.

NFS Access for Shares

NOTE: Multiple shares pointing to the same target directory must have the same NFS access settings. The Web Management Interface applies the same NFS access for all shares pointing to the same directory.

Click the link in the **NFS Access** column next to the share you want to configure. The NFS Share Access page is displayed. You can configure NFS access to the share using standard Linux “exports” file syntax.

On the Shares page, click the name of the access type listed in the NFS Access column to open the NFS Share Access page.



The NFS Access text box is a window into the client access entries in GuardianOS's *exports* file. This file serves as the access control list for filesystems that may be exported to NFS clients. You can use the Add Host controls as described below to assist in making entries to the file, or you can directly edit the text box. After all entries are made, click **OK** to return to the **Security > Shares** page.

NOTE: The syntax used in this file is equivalent to standard Linux exports file syntax. If the SnapServer detects any errors in syntax, a warning message appears. You can choose to correct or ignore the error warning.

The SnapServer Exports File Default Options. The default entry provides read-write access to all NFS clients (including NFSv4, if Kerberos security is not enabled).

```
* (rw,insecure,async,root_squash,no_all_squash)
```

The entry options are explained in the following table:

Entry Code	Meaning
Asterisk	All NFS clients
ro	The directory is shared read only (ro).
rw	The client machine will have read and write (rw) access to the directory.
insecure	Turns off the options that require requests to originate on an Internet port less than IPPORT_RESERVED (1024).
root_squash	Forces users connected as root to interact as the "nobody" user (UID 65534). This is the GuardianOS default.

Entry Code	Meaning
no_root_squash	no_root_squash means that if root is logged in on your second machine, it will have root privileges over the exported filesystem. By default, any file request made by user root on the client machine is treated as if it is made by user nobody on the server. (Exactly which UID the request is mapped to depends on the UID of user nobody on the server, not the client.) If no_root_squash is selected, then root on the client machine will have the same level of access to the files on the system as root on the server. This can have serious security implications, although it may be necessary if you want to perform any administrative work on the client machine that involves the exported directories. You should not specify this option without a good reason.
async	Tells a client machine that a file write is complete – that is, has been written to stable storage – when NFS has finished handing the write over to the filesystem.
no_all_squash	Allows non-root users to access the nfs export with their own privileges.

Configuring Export Strings for NFSv4 with Kerberos Security. Share access for NFSv4 clients can be enforced either by the traditional NFS host method (described in “[The SnapServer Exports File Default Options](#)”) or via Kerberos.

If Kerberos is enabled, access is applied uniformly to all Kerberos-authenticated NFSv4 clients connected using the matching Kerberos option. Host-based access as described in [The SnapServer Exports File Default Options](#) still applies to NFSv2 and v3 clients when Kerberos is enabled, but it does not apply to NFSv4 clients.

When UNIX Kerberos security is enabled for NFSv4, the following entries are automatically added to the NFS Access settings for each NFS-enabled share:

```
gss/krb5(rw,insecure,async,root_squash,no_all_squash)
gss/krb5i(rw,insecure,async,root_squash,no_all_squash)
gss/krb5p(rw,insecure,async,root_squash,no_all_squash)
```

These give read-write access to Kerberos-authenticated NFSv4 users connecting via:

- Standard Kerberos (**gss/krb5**)
- Kerberos with data integrity checksumming (**gss/krb5i**)
- Kerberos with protection/encryption (**gss/krb5p**).

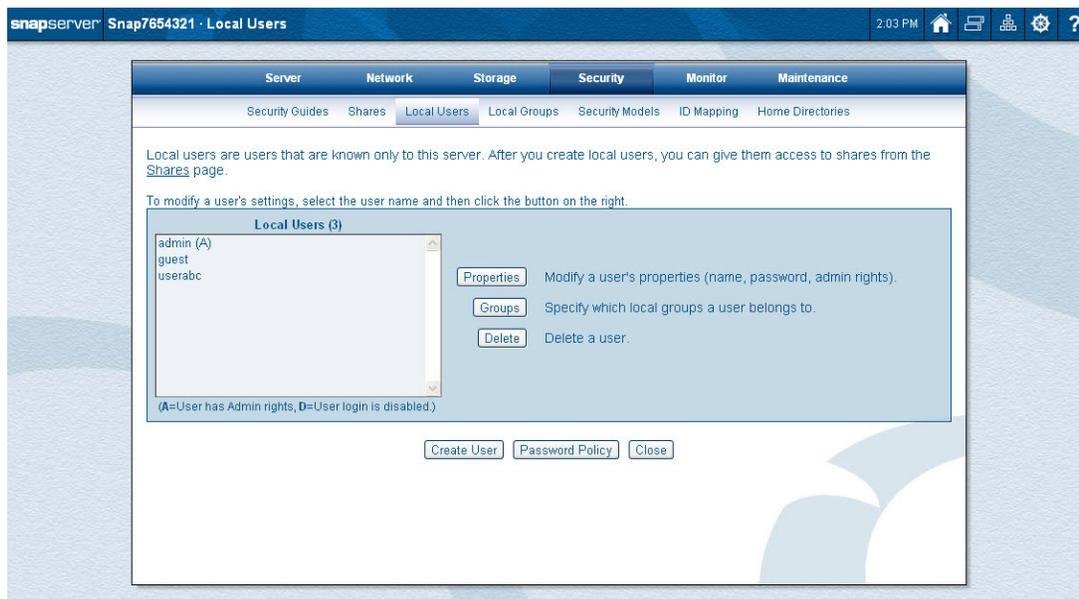
These entries can be independently removed, added, and modified on each NFS-enabled share.

Using the Add Host Controls. Follow these steps:

1. Select **one** of the following options:
 - SnapServer **Default Options** – Inserts the default options as described above
 - **Read Only** – Inserts the read only option only
 - **Both** – Inserts default options, but substitutes read only for read/write
2. Do **one** of the following in the NFS host text box:
 - **To apply the options to all NFS hosts** – Leave this field blank
 - **To apply the options to specific hosts** – Enter one or more IP addresses.
3. Click **Add Host**.

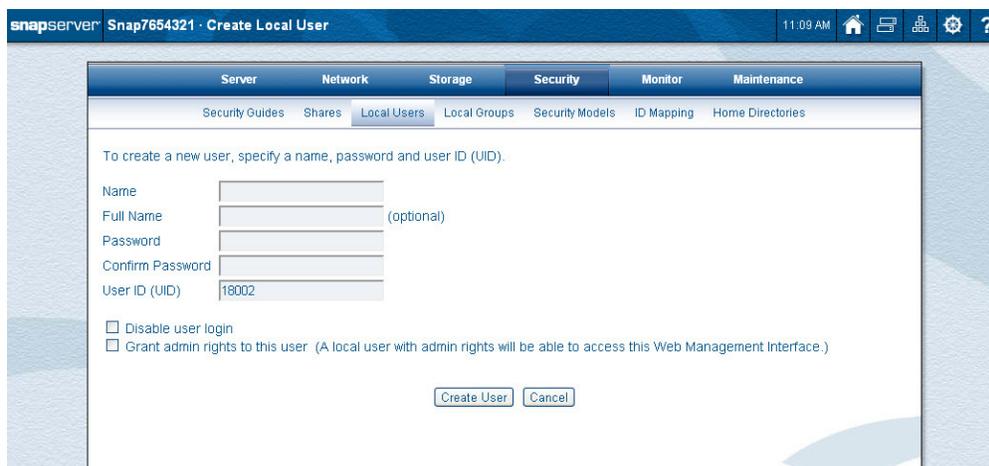
Local Users

The Local Users page (**Security > Local Users**) provides all the options to manage local users. Local users are users that are known only to the server being accessed. Each server running GuardianOS comes with two predefined users: admin and guest. The admin user has full Administrator rights. Go to **Security > Local Users** to view settings or make changes.



Create a User

Click the **Create** button to create a new user on this server. Enter the user data, select any special options, and click the **Create User** button again.



To Create a Local User

1. On the **Local Users** page, click **Create**.
2. On the **Create Local User** page that opens, enter the requested **information**:

Option	Description
Name	Use up to 31 alphanumeric characters and the underscore.
Full Name	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
Password	Passwords are case-sensitive. Use up to 15 alphanumeric characters without spaces.
Password Verify	Type the chosen password again for verification.
User ID (UID)	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see “User and Group ID Assignments” on page 7-3.
Disable User Login	Select this checkbox to disable the user login. The user's information will remain in the system, but login rights will be denied. The user login can be enabled by deselecting the checkbox. This checkbox can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.
Exempt from Password Expiration and Character Requirements	This checkbox is only visible if Password Policy is enabled. Select this checkbox to exempt this user from password expiration and character requirement policies.
Grant Admin Rights To This User	Select this checkbox to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click **Create User** again to create the user account.

Edit User Properties

Use the **Properties** button to open the Local User Properties page to make changes.

The screenshot displays the 'Local User Properties' configuration page. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'Security' section is active, with sub-tabs for 'Security Guides', 'Shares', 'Local Users', 'Local Groups', 'Security Models', 'ID Mapping', and 'Home Directories'. The 'Local Users' tab is selected, showing the following fields and options:

- Name: Fred
- Full Name: Fredrick Flintrock (optional)
- Password: (blank)
- Password Verify: (blank)
- User ID (UID): 18001
- Disable user login
- Grant admin rights to this user (A local user with admin rights will be able to access this Web Management Interface.)

Buttons for 'OK' and 'Cancel' are located at the bottom of the form.

To Edit Local User Properties

1. On the **Security > Local Users** page, select the user you want to edit and click **Properties**.
2. On the Local User Properties page that opens, enter or change the following **information**:

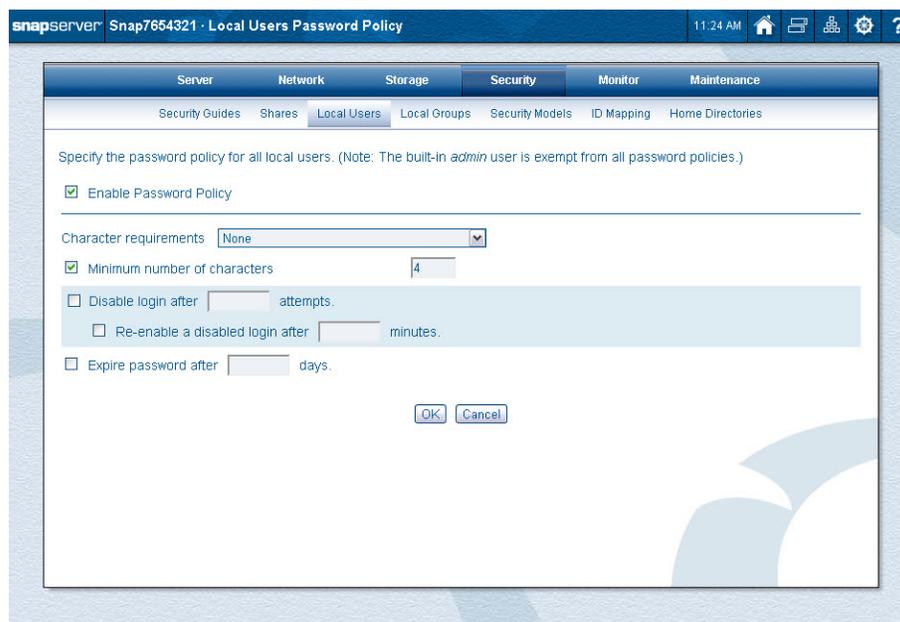
Option	Description
Name	Cannot be modified.
Full Name	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
Password	Passwords are case-sensitive. Use up to 15 alphanumeric characters. Leave this field blank to keep the existing password.
Password Verify	Type the chosen password again for verification. Leave this field blank to keep the existing password.
User ID (UID)	<p>Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see “User and Group ID Assignments” on page 7-3.</p> <p>NOTE: Changing a user's UID may alter filesystem access permissions that apply to that UID. In addition, any existing permissions for a UID previously assigned to a user that are changed to a different UID may become active if another user is created with the same UID. Carefully consider security configuration on existing files and directories before changing the UID of a user.</p>
Disable User Login	<p>Select this checkbox to disable the user login. The user's information will remain in the system, but login rights will be denied. The user login can be re-enabled by deselecting the checkbox.</p> <p>This checkbox can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.</p>
Exempt from Password Expiration and Character Requirements	<p>NOTE: This checkbox is only visible if Password Policy is enabled.</p> <p>Select this checkbox to exempt this user from password expiration and character requirement policies.</p>
Grant Admin Rights To This User	Select this checkbox to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click **OK**.

User Password Policies

NOTE: Local users can be individually exempted from password expiration and character requirements. This may be necessary for some special users, such as users configured to perform backups. See [“To Create a Local User”](#) for procedures to set password policy for local users. Also, the built-in *admin* user is automatically exempt from all password policies.

Click the **Password Policy** button to make changes to all the local user password settings.



To Set Password Policy for Local Users

1. On the **Security > Local Users** page, click the **Password Policy** button.
2. On the Local Users Password Policy page, check the **Enable Password Policy** box.
3. Enter the following **information**:

Option	Description
Character Requirements	Select the alpha/numeric/special character requirements for the password from the drop-down list.
Minimum Number of Characters	Check the checkbox to enable the policy, then enter the minimum number of characters required for the password.
Disable Login After <i>n</i> Attempts	Check the checkbox to enable the policy, then enter the number of times a user can fail to login before the system locks the user out. NOTE: To unlock a user, clear the Disable User Login checkbox for the user in the Local Users page.
Re-enable a Disabled Login After <i>n</i> Minutes	If you have defined a limit to the number of times a user can fail to log in, you can also check this checkbox and enter a time period after which the system will allow the user to log in again. NOTE: This will save the administrator from having to manually re-enable the user.
Expire Password After <i>n</i> Days	Check the checkbox to enable the policy, then enter the number of days before the password must be changed. NOTE: Local users with expired passwords can change their passwords at: <a href="http://<servername>/changepassword">http://<servername>/changepassword .

4. Click **OK** to save the settings and return to the Local Users page.

Assign User to Group

Use the Groups for Local Users page (**Security > Local Users > Groups**) to make changes to a local group membership.



To Add or Remove Users from Groups

1. On the **Groups for Local User** page, select a **user**.
2. Click **Groups**.
The group settings for the selected user are shown.
3. To add the user to a group, select the group from the right-side list and click **Add**.
4. To delete the user from a group, select the group from the left-side group and click **Remove**.
5. Click **OK** to save your changes and return to the Local Users page.

Delete Local User

On the Local Users page, click the **Delete** button to remove a user.

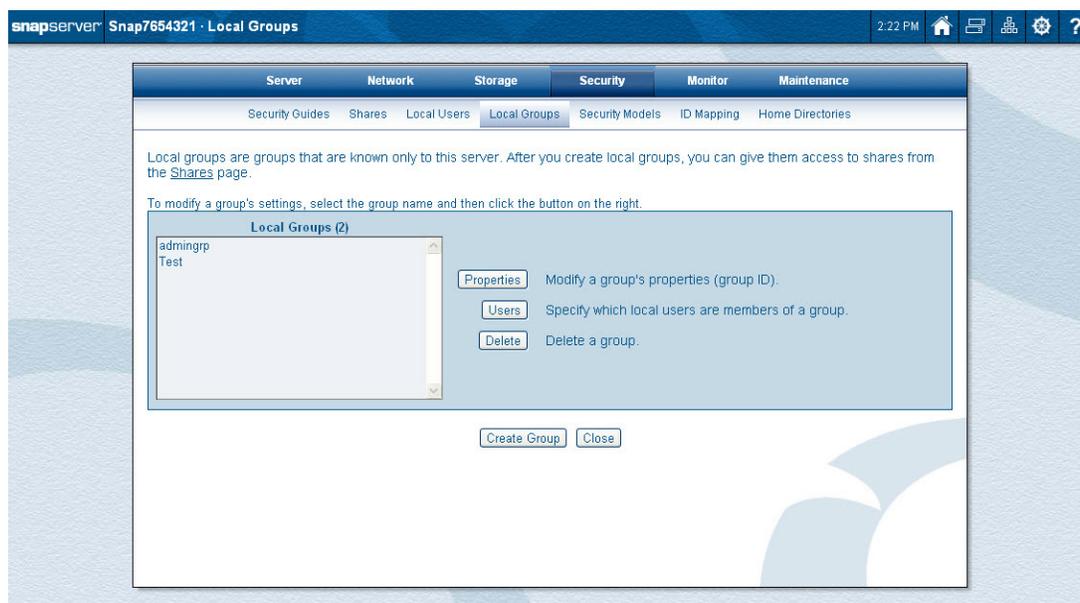


To Delete a Local User

1. On the **Security > Local Users** page, select the user to be deleted.
2. Click **Delete**.
3. The confirmation page will display. Click **Delete User** to delete the selected user (or click **Cancel**).

Local Groups

The Local Groups page (**Security > Local Groups**) provides all the options to manage local groups. Local groups are groups of local users that are known only to the server being accessed. Each server running GuardianOS comes with one predefined group: `admingrp`.



Create New Group

Click the **Create** button to create a new group on this server. Enter the group name, accept or change the Group ID (GID), and click the **Create Group** button.



To Create a New Local Group

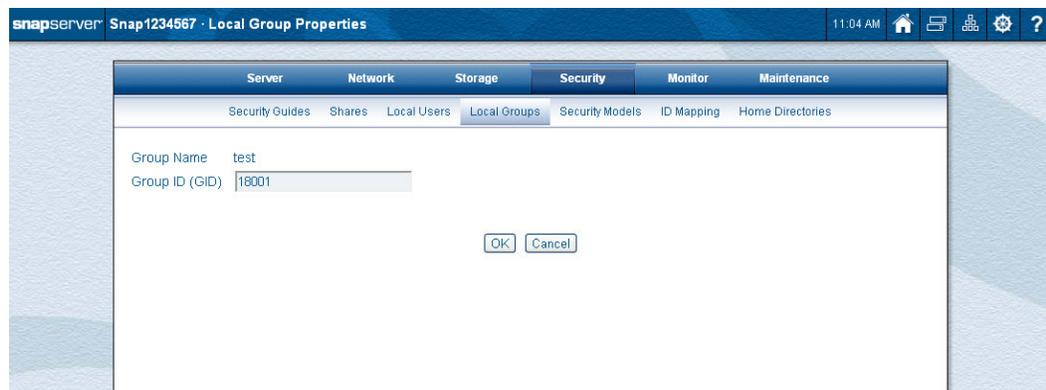
1. On the **Local Groups** page, click **Create**.
2. On the **Create Local Group** page that opens, enter the following information:

Option	Description
Group Name	Use up to 31 alphanumeric characters and the underscore.
Group ID (GID)	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see "User and Group ID Assignments" on page 7-3.

3. Click **Create Group** when finished. The Users for Local Group page is displayed, allowing you to add users to your new group.
4. Click **Close** when you are finished with local groups.

Edit Group Properties

Use the **Properties** button to open the Local Group Properties page to make changes to the options there.



To Edit Local Group Properties

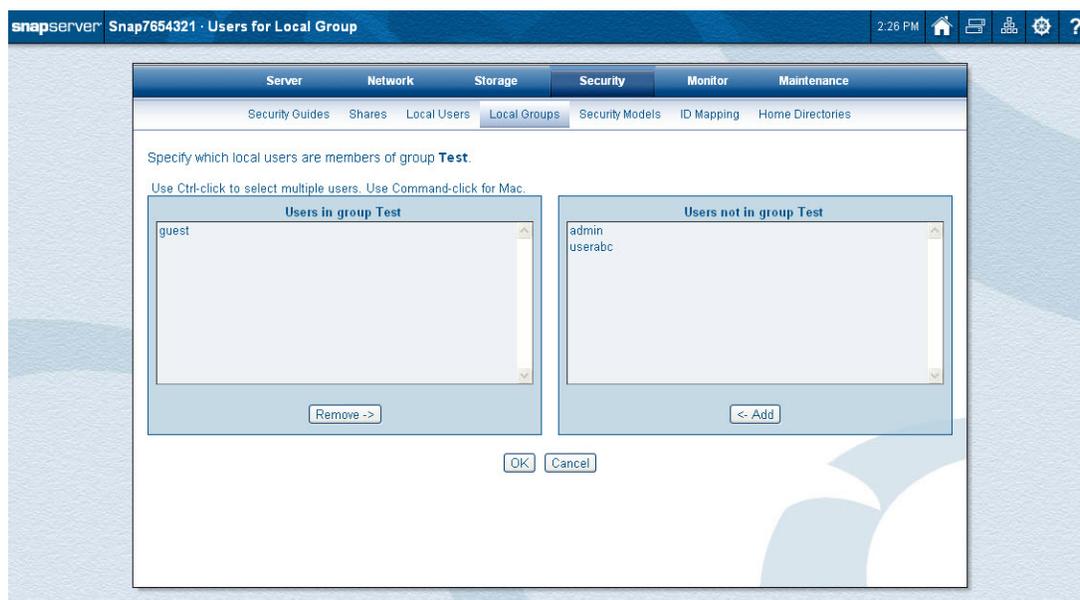
1. On the **Security > Local Groups** page, select the group you want to edit and click **Properties**.
2. On the page that opens, you can change the GID. For information on available UID ranges, see “[User and Group ID Assignments](#)” on page 7-3.

NOTE: Changing a group's GID may alter filesystem access permissions that apply to that GID. In addition, any existing permissions for a GID previously assigned to a group that are changed to a different GID may become active if another group is created with the same GID. Carefully consider security configuration on existing files and directories before changing the GID of a group.

3. Click **OK**.

Specify Users in Group

Use the Users for Local Group page (**Security > Local Groups > Users**) to make changes to a local group membership.



To Add or Remove Group Users

1. After creating a new group, or when editing an existing group, add and remove users by selecting the desired group and clicking **Users**.
2. Add users by selecting the user and clicking **Add**.
3. Delete users by selecting the user and clicking **Remove**.
4. Click **OK** when finished.

Delete Group

To Delete a Group

1. On the **Local Groups** page, select the group to be deleted and click **Delete**.

2. The confirmation page will display. Click **Yes** to delete the selected group, or click **No** to cancel the deletion.

Security Models

There are two file-level security models that can be used by a SnapServer: Windows/Mixed and UNIX. In Traditional RAID, the security model can be configured on volumes and the first set of folders created in the root of the volumes. In DynamicRAID, the security model only can be configured on the volumes.

The security model determines the rules regarding which security personality will be present on files and folders created by the various protocols and clients, and whether the personality of files and folders can be changed by changing permissions.

Folders created in a volume default to the security model of that volume. The folder's security model may differ from the personality of the folders (for example, folders with a Windows/Mixed security may have a UNIX personality).

For more information about Security Models, see [Appendix C, "Security and Access."](#)

Managing Volume Security Models

1. Select **Security > Security Models**.
2. Click the **security model name** (Windows/Mixed or UNIX).
3. From the drop-down list, select the **security model type** desired, and click **OK**.
4. At the confirmation message, click **Apply Security Model**.

If there are files and directories under the volume, you are prompted whether you want to recursively apply the change. When done, the main page displays a conversion status.

Managing Folder Security Models

NOTE: This is only available with Traditional RAID.

1. Select **Security > Security Models**.
2. Click the **volume name**.
3. To **create** a folder:
 - a. Click **Create Folder**.
 - b. Enter the **folder name**.
 - c. From the drop-down list, select the **security model type** desired.
 - d. Click **OK**.

4. To **change** the security model of a folder:



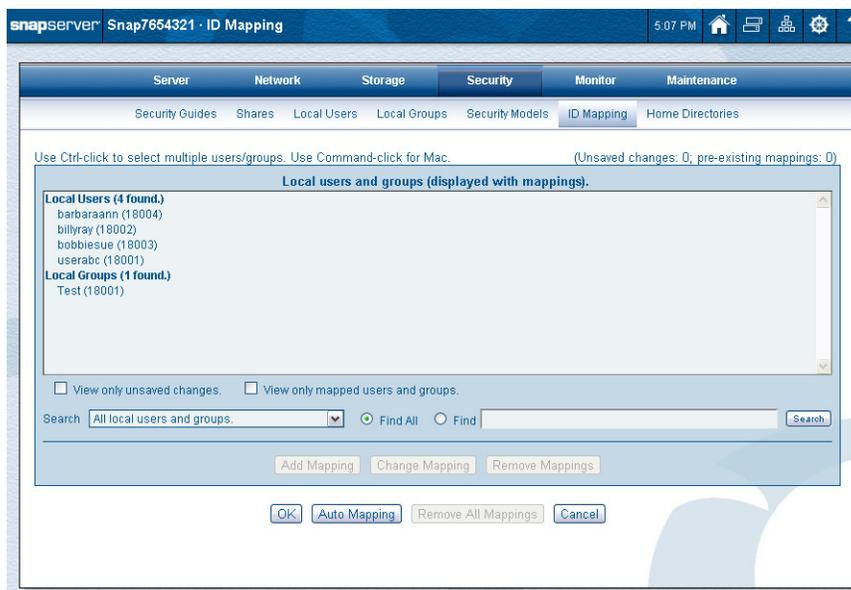
- a. Click the security model **icon** (“W” or “U”) of the folder.
- b. From the drop-down list, select the **security model type** desired, and click **OK**.
- c. At the confirmation message, click **Apply Security Model**.

If there are files and directories under the volume, you are prompted whether you want to recursively apply the change.

5. When done, click **Close**.

ID Mapping

ID mapping allows users and groups that exist on Windows domains to share user IDs with local or NIS users and groups. This results in the same permissions and quota consumption applying to both the Windows domain user and the local or NIS user.



Example: John Smith is a local user on a SnapServer, as well as having a user ID on a Windows domain. John’s quota for the SnapServer has been set to 200 MB. The administrator of the SnapServer maps the Windows domain user’s UID for John Smith to the local UID for John Smith, giving both users access to John’s 200 MB.

Select a local or NIS user or group from the displayed list on the default page. You can then click **Add Mapping** to map the user's UID or group's GID to that of a Windows domain user or group. **Change Mapping** is used to change existing mappings. **Remove Mappings** removes one or more mappings while **Remove All Mappings** removes all mappings that had been previously established.

To simplify the discovery of a desired user or group to manage their ID mapping search options are presented at the bottom of the selection pages. On the search results page, you can narrow the list by using the following options:

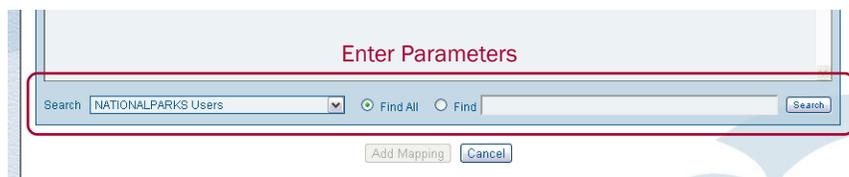
- Check **View only unsaved changes** to display only mapping changes that have not yet been applied.
- Check **View only mapped users and groups** to display only local or NIS users and groups that have been mapped to a Windows domain user or group.

Add Mapping

1. If the desired user or group to be mapped to does not appear in the default page list, use the **search option** to locate them.



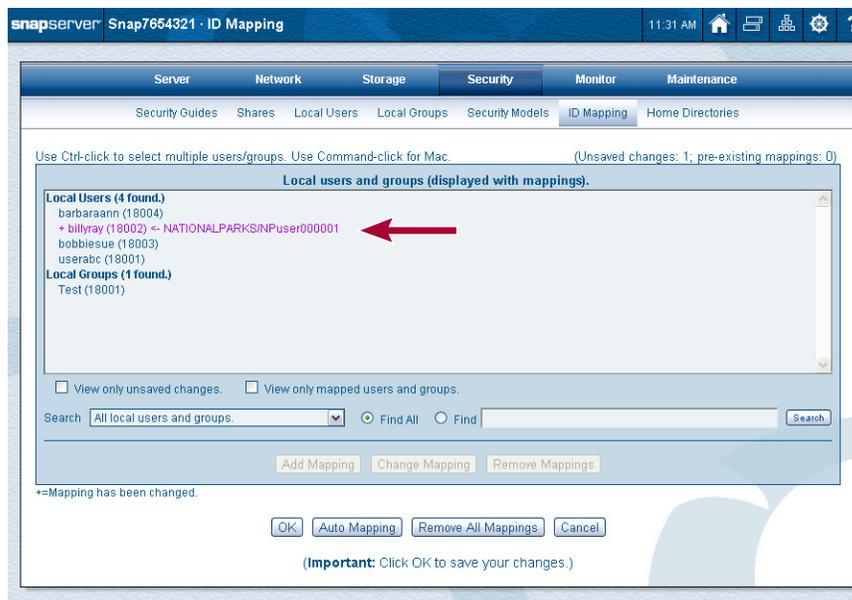
- a. At the bottom of the list, using the Search drop-down list, select the **local or NIS user or group list** to search.
 - b. Enter the **search string** (or select Find All).
Enter the exact **name** (or a string with a wildcard "*" before or after).
 - c. Click **Search** to display any matches.
2. Select a user or group from the results list, and click **Add Mapping**.
 3. At the Add Mapping page, select the Windows domain **user or group** list, and click **Search**.
 - To search for a specific user or group, use either Find All or a Find search string (wildcard "*" before or after allowed).



- For domains that **REQUIRE** authentication (showing an "(A)" after the name), select the domain name, enter the User Name and Password for that domain, and use either Find All or a search string (use the beginning of the user/group name).

- On the rare occasion you need to search for a Windows domain that's not listed (“remote domain”), select a Windows domain from the Search drop-down list through which to search, then enter in the Find box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, `remote_domain\user_name`). After you click Search, another authentication prompt may be presented to authenticate with the remote domain.
4. From the search results, select the Windows domain user you want to map the local or NIS user to, and click **Add Mapping**.

The mapping result is shown on the default page.



Check **View only unsaved changes** to display only changes that have not yet been applied. Check **View only mapped users and groups** to display only local or NIS users or groups that have been mapped to a Windows domain user or group.

5. Repeat [Steps 1–4](#) to add any other mappings.
6. Save your **changes**:
 - a. Click **OK** to save changes (or **Cancel** to reset).
 - b. At the confirmation page, click **Save Changes**.
 - c. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

See [“Filesystem Updates”](#) for more details.



IMPORTANT: Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

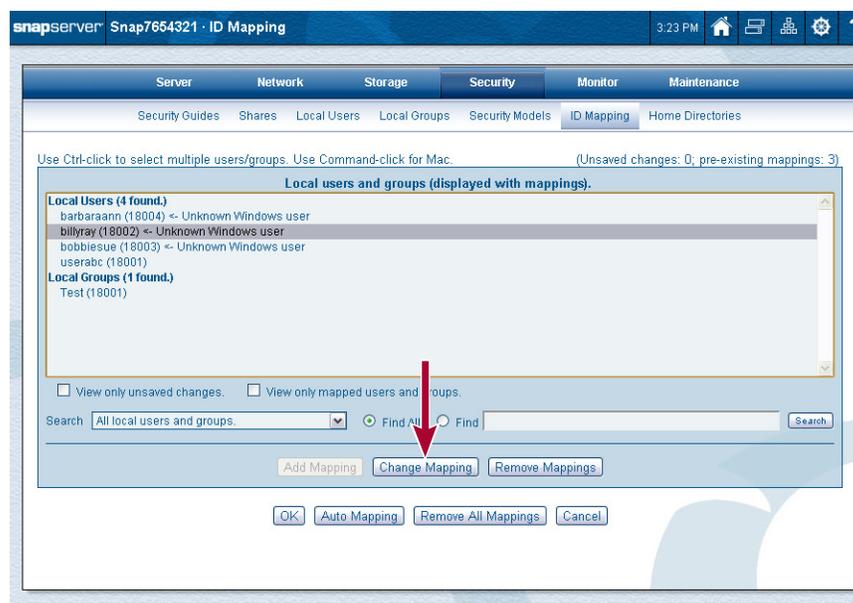
Change Mapping

To re-map a mapped local or NIS user or group to a different Windows domain user or group:

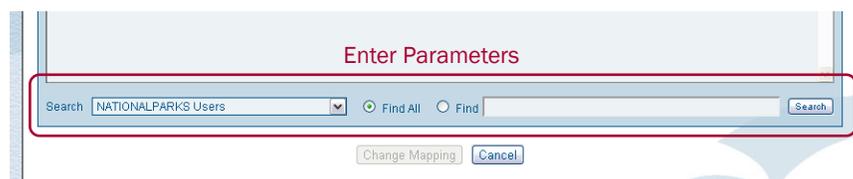
1. If the desired user or group to be changed does not appear in the default page list, use the **search option** to locate them.



- a. At the bottom of the list, using the Search drop-down list, select the **local or NIS user or group list** to search.
 - b. Enter the **search string** (or select Find All).
 Enter the exact **name** (or a string that uses the beginning characters of the user/group name).
 - c. Click **Search** to display any matches.
2. Select a mapped user or group to be changed, and click **Change Mapping**.



3. At the Change Mapping page, select the Windows domain **user or group** list, and click **Search**.
 - To search for a specific user or group, use either Find All or a Find search string (wildcard "*" before or after allowed).



- For domains that REQUIRE authentication (showing an "(A)" after the name), select the domain name, enter the User Name and Password for that domain, and use either Find All or a search string (use the beginning of the user/group name).

- On the rare occasion you need to search for a Windows domain that's not listed ("remote domain"), select a Windows domain from the Search drop-down list through which to search, then enter in the Find box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, `remote_domain\user_name`). After you click Search, another authentication prompt may be presented to authenticate with the remote domain.
4. From the search results, select a Windows/SMB domain user to map to and click **Change Mapping**.

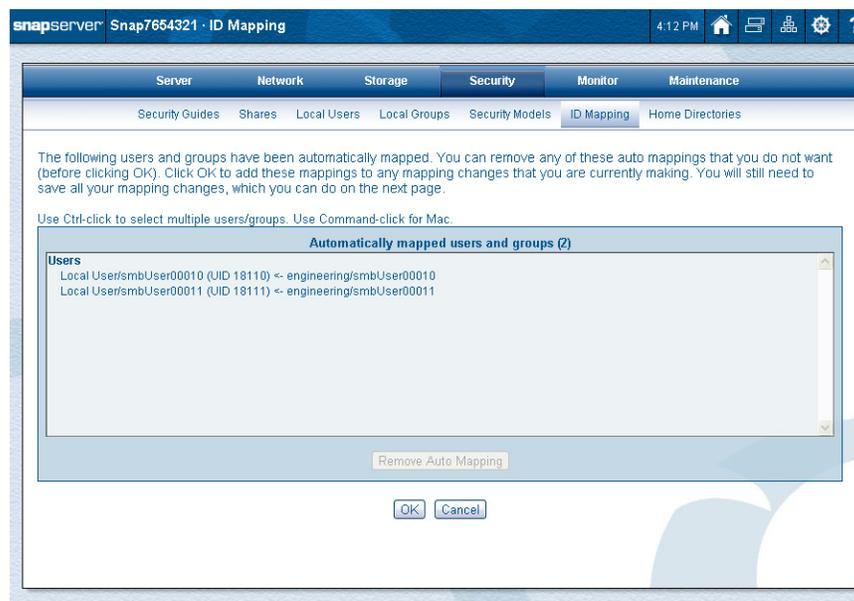
5. Repeat [Steps 1–4](#) until all **changes** are made.
6. Save your **changes**:
 - a. Click **OK** to save changes (or Cancel to reset).
 - b. At the confirmation page, click **Save Changes**.
 - c. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.
See "[Filesystem Updates](#)" for more details.



IMPORTANT: Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

Auto Mapping

1. Click **Auto Mapping** to generate a list of Windows domain users/groups that have the same name as your Local or NIS users and groups:
Domain, local, and NIS user/group lists are compared. The matches are automatically queued. Users and groups already mapped will not be affected
2. At the Auto Mapping confirmation page, click **View Auto Mappings** to continue.
A page is displayed summarizing your changes.



3. Save your **changes**:
 - a. Click **OK** to save changes (or **Cancel** to reset).
 - b. At the confirmation page, click **Save Changes**.
 - c. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.
See [“Filesystem Updates”](#) for more details.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

Remove Mappings

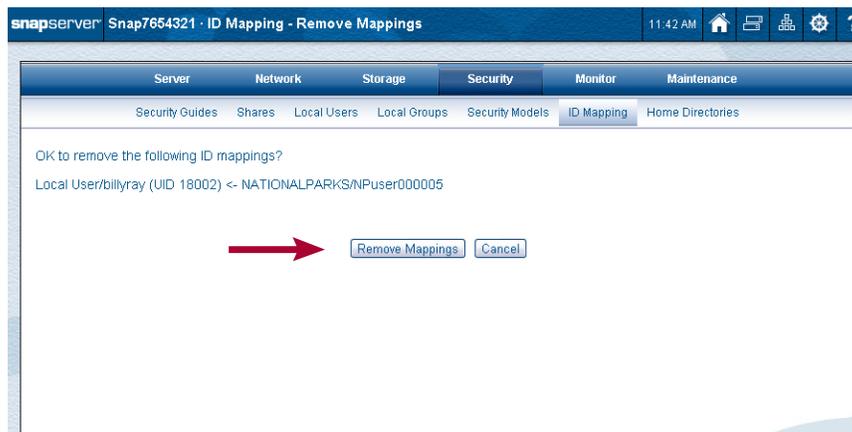
User mappings can be removed individually or all at once. Once removed, they can not be restored but must be added back using [“Add Mapping.”](#) You also have the option to update the filesystem after removing the ID mappings.

Remove a Mapping

1. At the default page, select one or more users/groups you wish to unmap and click **Remove Mappings**.

Check **View only mapped users and groups** to display only local or NIS users or groups that have been mapped to make it easier to find ones to remove.

2. At the Remove Mappings page, verify the users/groups on the list, and click **Remove Mappings**.



The mappings are removed and the default page is displayed.

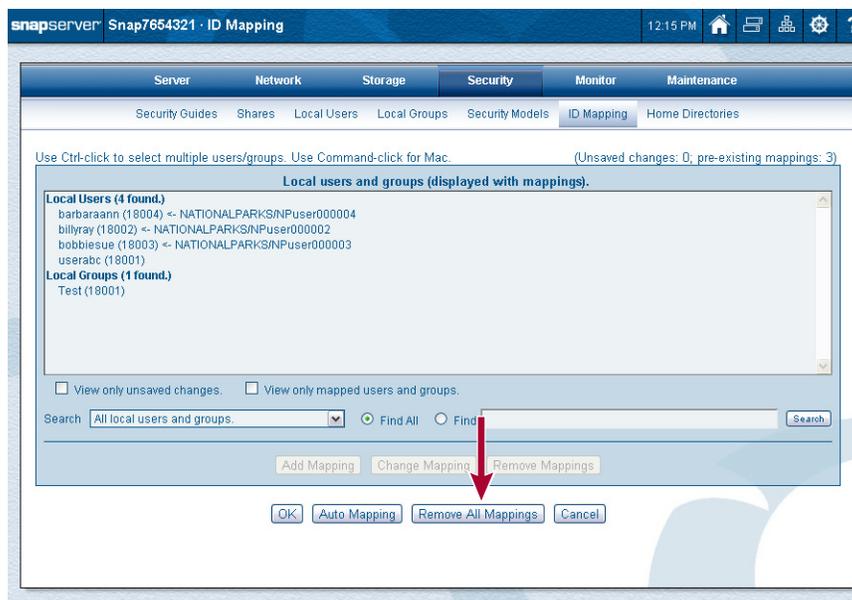
3. Repeat [Steps 1–2](#) until all **changes** are made.
4. Save your **changes**:
 - a. Click **OK** to save changes (or **Cancel** to reset).
 - b. At the confirmation page, click **Save Changes**.
 - c. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.
See [“Filesystem Updates”](#) for more details.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

Remove All Mappings

The **Remove All Mappings** button allows you to remove all ID mappings on the SnapServer. Click this only if you want to remove all ID mappings. If there are no mappings, the button is grayed out.

1. At the default page, click the **Remove All Mappings** button.



Check **View only unsaved changes** to display only mapping changes that have not yet been applied. Check **View only mapped users and groups** to display only local or NIS users/groups that have been mapped to a Windows domain user or group.

2. A confirmation page appears. Click **Remove All Mappings**.
3. Save your **changes**:
 - a. Click **OK** to save changes (or **Cancel** to reset).
 - b. At the confirmation page, click **Save Changes**.
 - c. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

See “[Filesystem Updates](#)” for more details.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

Remove Missing ID Mappings

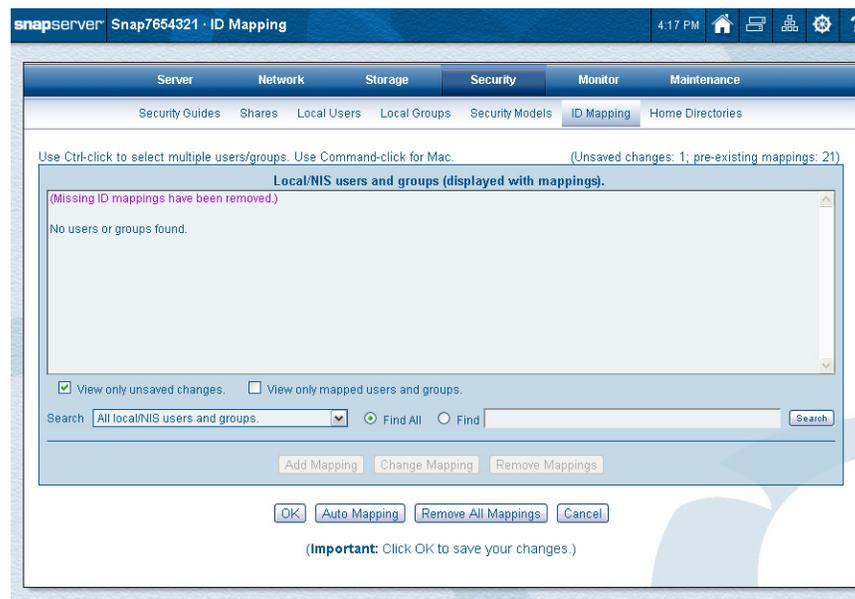
If the SnapServer has mappings for users or groups that no longer exist, the following warning message may be displayed at the top of the main ID Mappings page:



1. Click the **Click here** link in the warning message to display the following page:



2. Click **Remove Missing Mappings** to clear them from the system. A confirmation is shown on the ID Mapping main page.



3. Click **OK** to save changes.

Filesystem Updates

After making any changes to ID mappings, you are presented with a filesystem update option page, where you can choose either **Update Filesystem** or **Do Not Update Filesystem** options.

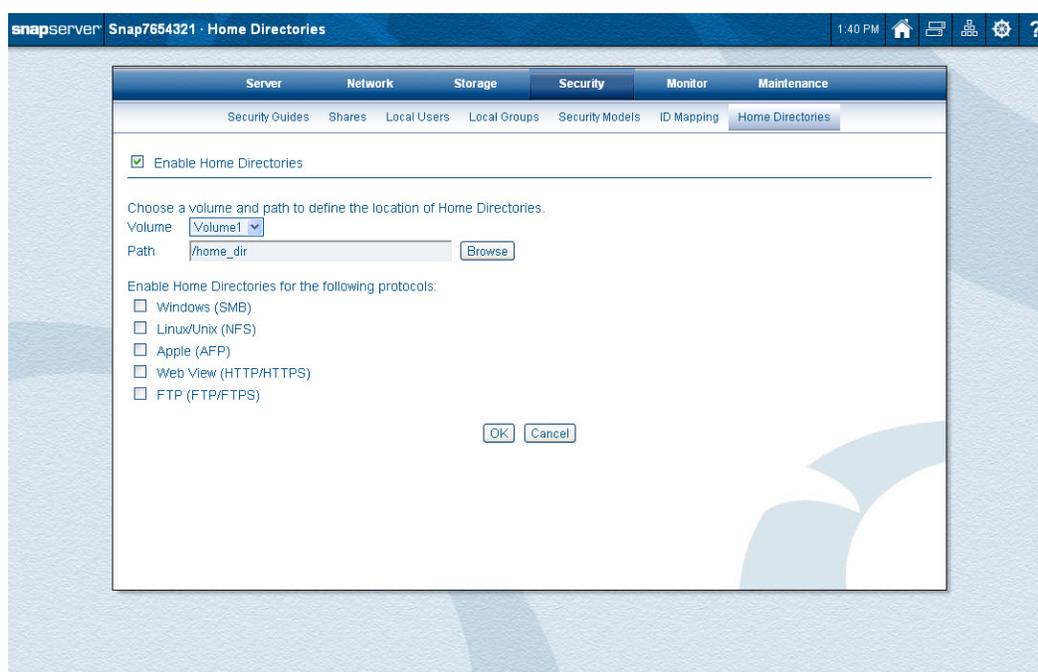


If you choose Update Filesystem, UID and GID ownership on files and SIDs in ACLs will be updated to reflect the ID mapping operation.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

Home Directories

To enable Home Directories, go to **Security > Home Directories** and check **Enable Home Directories**. Choose the volume, path, and protocols you want.



The Home Directories feature creates a private directory for every local or Windows domain user that accesses the system. When enabling Home Directories (from the **Security > Home Directories** page), the administrator creates or selects a directory to serve as the home directory root. When a user logs in to the server for the first time after the administrator has enabled Home Directories, a new directory named after the user is automatically created inside the home directory root, and is configured to be accessible only to the specific user and the administrator.

Depending on the protocol, home directories are accessed by users either via a user-specific share, or via a common share pointing to the home directory root.

Home directories are supported for SMB, NFS, AFP, HTTP/HTTPS, and FTP/FTPS. They are accessed by clients in the following manner:

- For SMB, AFP, and HTTP/HTTPS, users are presented with a virtual share named after the user name. The virtual share is visible and accessible only to the user. Users are not limited only to their virtual shares; all other shares on the server continue to be accessible in the usual fashion.

- For NFS, the home directory is exported. When a user mounts the home directory root, all home directories will be visible inside the root, but the user's home directory will be accessible only by the user and the administrator.

NOTE: If desired, UNIX clients can be configured to use a Snap Home Directory as the local user's system home directory. Configure the client to mount the home directory root for all users, and then configure each user account on the client to use the user-specific directory on the SnapServer as the user's home directory.

- For FTP/FTPS, local users will automatically be placed in their private home directory when they log in. Access to the home directory is facilitated through a share pointing to a parent directory of the home directory, so users can still change to the top-level directory to access other shares.

If **ID Mapping** is enabled, domain users and local users mapped to the same user will be directed to the domain user's home directory. In some cases, data in the local user's home directory will be copied to the domain user's home directory:

- If a local user home directory accumulates files before the local and domain users are mapped, and if the domain user's home directory is empty, the local user's files will be copied to the domain user's home directory the first time the local user connects after the users are mapped.
- If both the local and domain user home directories accumulate files before the local and domain users are mapped, the files in the local user's home directory will not be copied to the domain user's home directory.

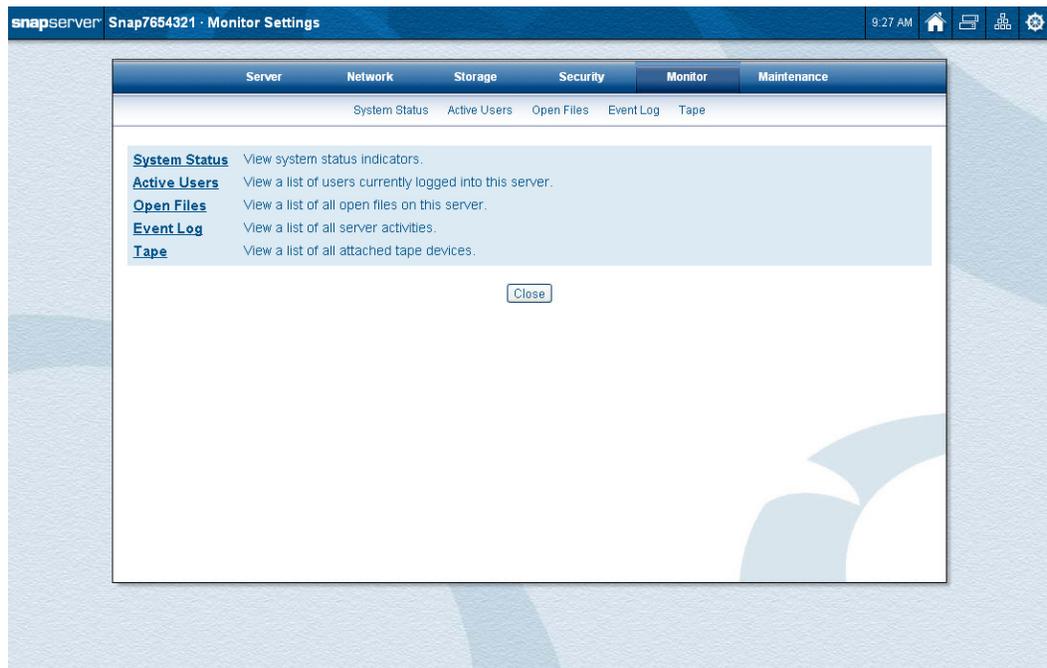
Configure Home Directories

Complete the following fields and click **OK**.

Field	Description
Enable Home Directories	Check to enable Home Directories for local users. Remove the check to disable.
Volume	Select the volume where the Home Directories will be located. NOTE: Be sure the volume you select has enough disk space. Once Home Directories are placed, they cannot be moved.
Path	Provide the path to the Home Directories or click Browse to create a new folder. The default path is <code>/home_dir/</code> .
Protocols	Check each of the protocols where Home Directories will be enabled.

NOTE: Do not put Home Directories on a volume that might be deleted. If you delete the volume, you will also delete the Home Directories.

This chapter addresses the options for monitoring the SnapServer. Here you can view the system status and other activities.



Topics in System Monitoring:

- [System Status](#)
- [Active Users](#)
- [Open Files](#)
- [Event Log](#)
- [Tape Monitor](#)

System Status

Use the System Status page (**Monitor > System Status**) to assess the status of the SnapServer and any attached expansion units.



SnapServer Status

The following status fields will be displayed for the SnapServer (head unit). Any critical messages will be displayed in **red**.

Field	Description
Server Name	Current name of the server. The default server name is SNAPnnnnnn, where nnnnnn is your server number (for example, SNAP112358).
Server Model	Server model (such as SnapServer DX1).
OS Version	The version of GuardianOS currently loaded on the SnapServer.
Server Number	Number derived from the MAC address of the primary Ethernet port, used as part of the default server name.
Serial Number	Unique number assigned to the SnapServer.
Uptime	The amount of time the server has been up (since the last reboot) in “days:hours:minutes” format.
CPU 1	The type of central processing unit for the server’s first CPU.
CPU 2	The type of central processing unit for the server’s second CPU.
Memory	Amount of system RAM.
Ethernet 1	Details on the server’s primary Ethernet connection.
Ethernet 2	Details on the server’s secondary Ethernet connection.
Ethernet <i>n</i>	Details on the server’s other Ethernet connections, if any optional Ethernet card is installed.

Field	Description
Ambient Temp.	The temperature of the space around the SnapServer.
CPU Temp.	Current CPU temperature.
Power Supply	The status of power supply modules
Fan Status	The status of fan modules.

Expansion Unit Status

The following status fields will be displayed for each expansion unit.

Field	Description
Expansion Unit	EXTN1, EXTN2, etc.
Expansion Model	SnapServer Expansion S50, etc.
Serial Number	The serial number of the expansion unit
Ambient Temperature	The temperature of the space around the expansion unit.
Power Supply	The status of the power supply
Fan Status	The status of fan modules.

Click **Refresh** to update the information. Click **Close** to return to the main Monitor page.

Active Users

Use this page to view read-only details on the active users logged on to the server. Information available on this page includes user names of all active users, their workstation names, authorization, the number of open files they have on the share, the protocol, and when they logged on. Columns can be sorted in ascending or descending order by clicking the column head.

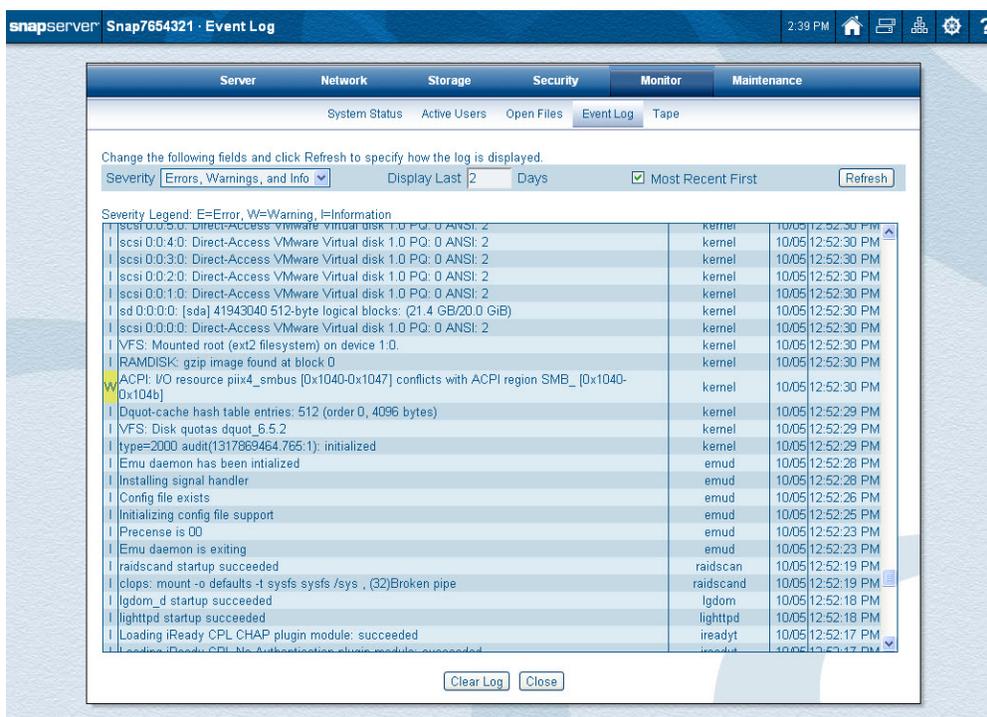
NOTE: Active users are not displayed for HTTP or NFS.

Open Files

Use this page to view read-only details on the open files in use on this server.

Event Log

Use the **Event Log** page to view a log of operations performed on the server.



Entries are color coded according to severity as described in the following table:

Background Color	Entry Type
Red 	Errors (E)
Yellow 	Warning (W)
(no color)	Informational or Unclassified (I)

Filter the Log

Edit the following fields as appropriate, then click **Refresh**.

Option	Description
Severity	Select the type of entries you want to view.
Display Last	Enter the number of days' entries (24-hour periods) you want to view
Most Recent First	Select to start the list with the most recent entry, deselect to start the list with the oldest entry.

Erase All Log Entries

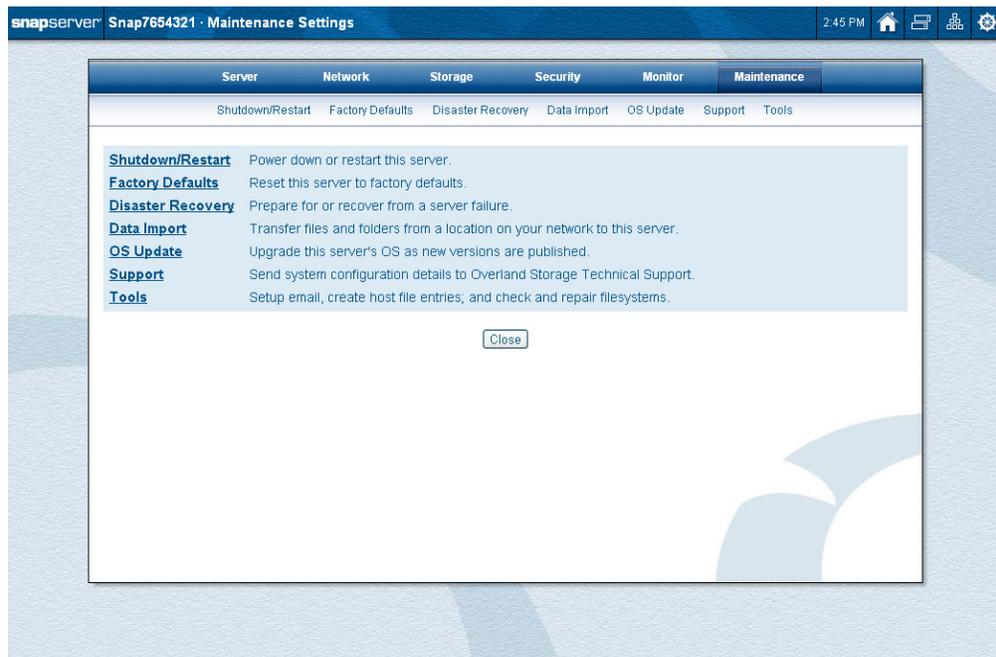
Click **Clear Log** to erase all log entries.

Tape Monitor

Use the **Tape Monitor** page to view read-only details on the SCSI and USB tape devices attached to the server.

Field	Description
Device Model	The manufacturer's model for the device.
Device Type	Type of tape device: either Sequential-Access (tape drive) or Medium-Changer (for example, robotic arm for a tape library).
Device Name	Name of the device node to which the device is bound.
Connection	Identifies the connection type: SCSI or USB.
Bus	Bus number indicating which physical interface (for example, SCSI card) the device is connected to.
ID	ID number (SCSI only)
LUN	LUN identifier (SCSI only)

Clicking the **Maintenance** tab on the Web Management Interface displays seven options used to maintain your SnapServer appliance and the GuardianOS 7.2 software.

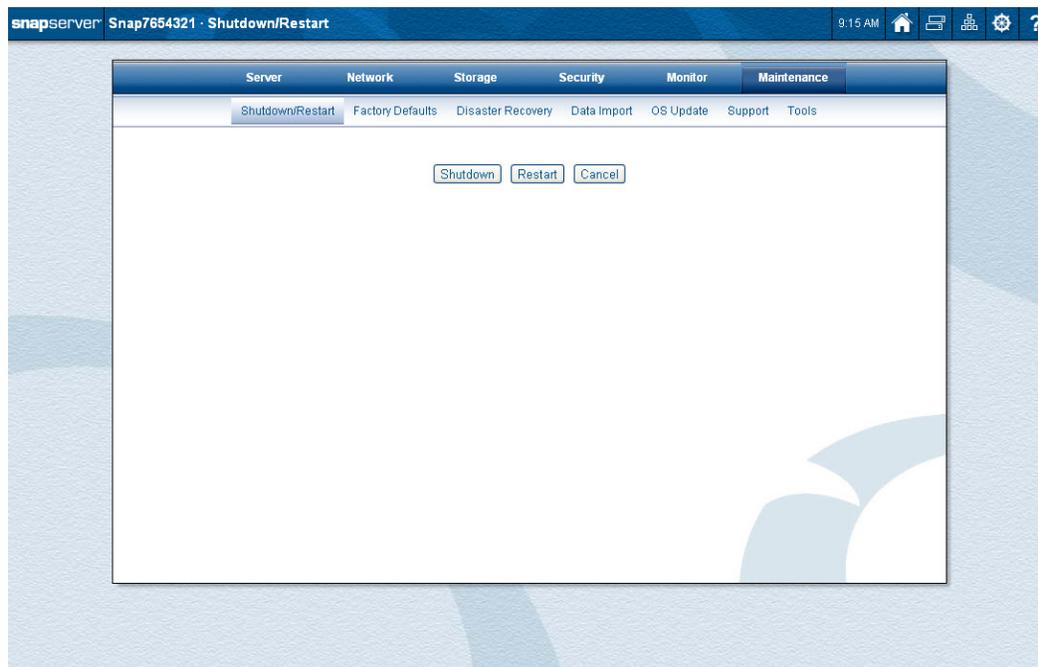


Topics in Web Management Interface

- [Shutdown and Restart](#)
- [Factory Defaults](#)
- [Disaster Recovery](#)
- [Data Import](#)
- [OS Updates](#)
- [Support](#)
- [Maintenance Tools](#)
 - [Email Notification](#)
 - [Host File Editor](#)
 - [To Check the Filesystem on a Volume](#)
 - [To Check the Root Filesystem](#)

Shutdown and Restart

Use the **Shutdown/Restart** page to reboot or shut down the server.



Click one of the following buttons:

- **Shutdown** – Shuts down and powers off the server.
- **Restart** – Reboots the server via a controlled shutdown and restart.

Manually Powering SnapServers On and Off



CAUTION: To prevent possible data corruption or loss, make sure all users are disconnected from the SnapServer before powering down the server.

Use the power button on the front of the server to power on and power off the server:

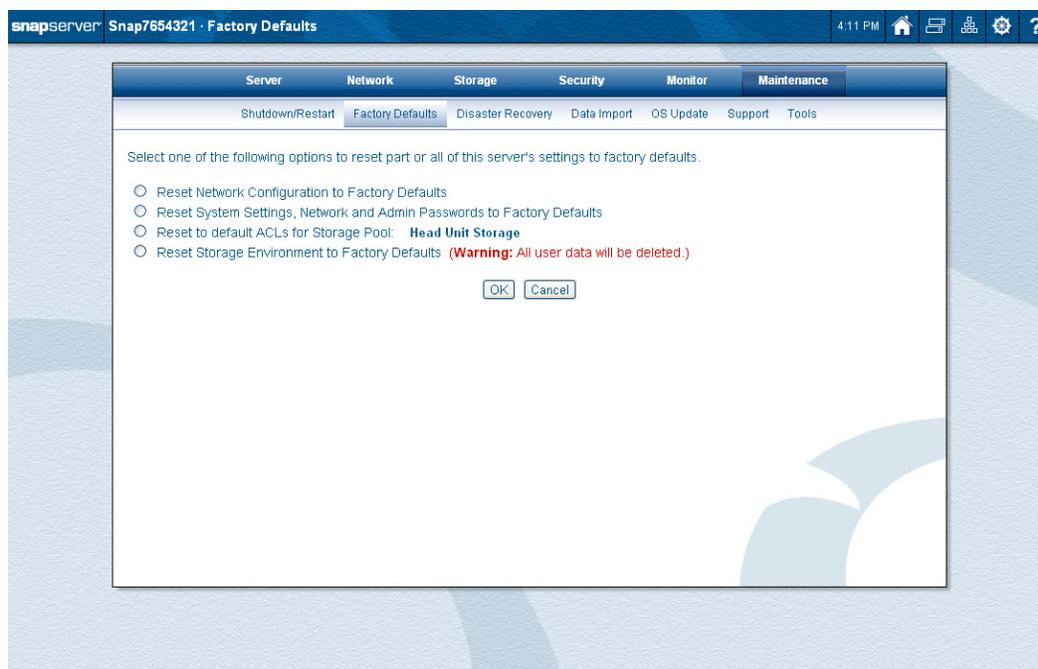
- To turn the server **ON**, press the power button on the front of the server. The server takes a few minutes to initialize. A green system/status LED indicates that the system is up and running.
- To turn the server **OFF**, press and release the power button to begin the shutdown process. Do not depress this button for more than four seconds.

NOTE: SnapServers have a persistent power state. When a physical loss of power occurs, the SnapServer returns to the same operation it had when the power went out. Therefore, if the system is powered down prior to a power loss, it will remain powered down when the power is restored.

Factory Defaults

GuardianOS allows you to reset different components of the system back to the original factory defaults. You can reset some or all of the factory settings using the different options available on the **Factory Defaults** page.

 **CAUTION:** Each reset option requires a restart of the server. To prevent possible data corruption or loss, make sure all users are disconnected from the SnapServer before proceeding.



Navigate to the **Maintenance > Factory Defaults** page in the Web Management Interface, select one of the following options, and then click **OK**:

- **Reset Network Configuration To Factory Defaults** – Returns TCP/IP and other network protocol settings to factory defaults.
- **Reset System Settings, Network, and Admin Passwords To Factory Defaults** – Returns the admin and root passwords to the default value, returns TCP/IP and other network protocol settings to factory defaults, eliminates all shares to all volumes, and returns settings for server name, date and time, users, groups, Windows and NIS domain memberships, quotas, and the activation and configuration of CA Antivirus to factory default values. Storage configuration and data is retained.

When the server finishes rebooting, the Login dialog box opens. Enter the default admin password of **admin**, and click **OK**. The Initial Setup Wizard runs, allowing you to reset the server name, admin password, and IP address.

NOTE: Resetting system settings will disable Snap EDR. After reset, you will need to uninstall, reinstall, and reconfigure Snap EDR.

- **Reset To Default ACLs For Volume: <volume name>** – Resets the file and directory security on selected volumes. Volumes are all set to the Windows/Mixed security model. All files and directories are set to the Windows personality with a Windows ACL that gives full access to Administrators, read access to Everyone, file/directory create access to Everyone (for directories), and full access to the owner (owners are retained in the reset operation).

NOTE: Rebooting or shutting down the server in the middle of an ACL reset will halt the operation, and it will not recommence on reboot.

- **Reset Storage Environment to Factory Defaults** – Storage configuration is reset and the Initial Setup Wizard is displayed when the SnapServer is restarted.

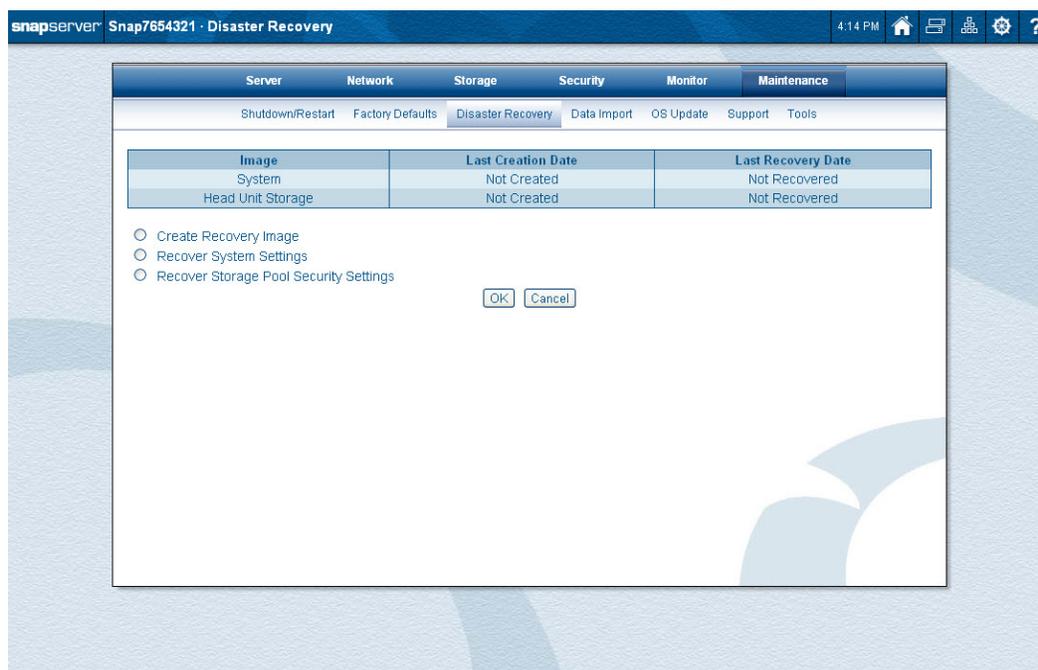
 **CAUTION:** ALL USER DATA WILL BE DELETED on the SnapServer when this option is selected. A confirmation page will be displayed, and the admin password must be entered before the process begins.

NOTE: Use this option to change your RAID configuration standard from DynamicRAID to Traditional RAID or vice versa.

Disaster Recovery

Disaster recovery entails creating the files you need to recover a SnapServer's configuration information, such as network and RAID configurations, as well as volume-specific information, such as ACLs and quota settings.

NOTE: Disaster recovery does not include user data. Backups of user data must be configured and managed separately; see "Backup Solutions" for detailed information on backup options.



It also includes what to do if all access to the data on a SnapServer is cut off due to a hardware or software failure. Focus is placed on the procedures for:

- Reinstalling the SnapServer operating system.
- Restoring the server to its original configuration with data intact.

These files are then used to restore any SnapServer to its original state. The disaster recovery feature can also be used to clone one server to another by restoring the disaster recovery image from one server to another server.

Backing Up Server and Volume Settings

In addition to backing up the data stored on the SnapServer, you may also back up the server's system and volume settings. The **Disaster Recovery** page allows you to create the files you need to restore these settings:

- Server-specific settings such as network, RAID, volume and share configurations, local user and group lists, snapshot schedules, and Snap EDR Management Console settings (if applicable).
- Volume-specific settings such as ACLs, extended attributes, and quota settings.

For information about scheduling these tasks, see [“Scheduling Data Protection Tasks”](#).

The SnapDRImage File and the Volume Files

Details on the SnapServer disaster recovery files and the information they contain are as follows:

- **SnapDRImage** – The SnapServer disaster recovery image saves server-specific settings such as network, RAID, volume and share configuration, local user and group lists, and snapshot schedules, and Snap EDR Management Console settings (if applicable). There is one SnapDRImage file per server, residing in the “.os_private” directory on the root of the first volume in Traditional RAID, or on the root of the first volume on the first storage pool in Dynamic RAID.

NOTE: The SnapDRImage file is in binary form and can be safely used only with the SnapServer Disaster Recovery tool. Other tools will not work and may compromise the integrity of the file.

- **Volume-specific files** – These files, named *backup.acl*, *backup.qta.groups*, and *backup.qta.users*, preserve volume-specific settings such as ACLs, extended attributes, and quota settings. One set of these files exists per volume, and are located as follows:
 - In Traditional RAID, the volume settings specific to each volume are stored in the “.os_private” directory on the root of each volume.
 - In DynamicRAID, the volume settings for an entire storage pool are stored in the “.os_private” directory on the first volume of the storage pool.



CAUTION: The Create Recovery Files option in the snapshot feature automatically updates the volume-specific files when the snapshot is taken. If you do not use snapshots to back up a volume to tape, you must manually regenerate these files whenever you change ACL or quota information to ensure that you are backing up the most current volume settings.

Creating the SnapDRImage and Volume Files

Creating a SnapDRImage that covers the scope of your server's configuration is essential to a successful disaster recovery operation. Create a disaster recovery image on the **Disaster Recovery** page. This DRImage should be created after server configuration is complete, and can be used to recover the server or a replacement server to the configured state.

Before you create the disaster recovery files, make sure you have completed the following activities:

- You have completely configured the SnapServer. If you subsequently make any major changes to the configuration of your server, you must repeat the procedures described in this section to have an up-to-date SnapDRImage.

NOTE: You may want to record, in an off-server location, the following information about the configuration of your server: (1) the server name; (2) the number of RAID's; (3) the number of volumes; and (4) the size of each volume. If the disaster recovery fails, having this information may be useful in recreating the original configuration of the server.

- You have devised and implemented a data backup strategy. It is recommended that you make a backup of your system regularly, from the root of the share for each volume, and store it in an off-server location. This ensures that the most current data is backed up and available for use with a disaster recovery.

Use the following procedure to create and secure the disaster recovery files:

Step 1: Create the disaster recovery files.

Navigate to the **Maintenance > Disaster Recovery** page. Select the **Create Recovery Image** button and click **OK** to create the SnapDRImage file and the volume files in a single operation.

Step 2: Copy the files to a safe place off the server.

Once the recovery image has been made, click the **Download Recovery Image** button to download the SnapDRImage file to a safe location on another server or backup medium. (See The SnapDRImage File and the Volume Files for file names and paths.) This strategy ensures that if the filesystem on the SnapServer is corrupted, the image file will be available to restore server settings.

The DRImage is also automatically placed in the root of the first user volume. These files will be copied to tape as part of your regular backup procedures.

Step 3: Take no action regarding the volume-specific files.

These files will be copied to tape as part of your regular volume backup procedures.

Restoring Original Server and Volume Configurations

To restore the original configurations to the server, two separate operations are required and they must be run sequentially. After you start any of the recovery processes, you will see the **Disaster Recovery Status** page.



CAUTION: Do not try to navigate back from this page during the recovery process. Activity is restricted to this page so that the recovery operation will not be interrupted.

NOTE: You can go to <http://<servername|IP address>/recover> to bypass the initial setup and go straight to Disaster Recovery.

Step 1: Restore server settings.

Select the **Recover System Settings** button and click **OK** to open the **Server Recovery** page. Use the **Browse** button to navigate to the SnapDRImage file. Click **Recover** to start the operation. If the recovery file contains Snap EDR application settings, you are asked if you

want to include those settings. Check the settings you want to recover, and click **Recover**. Once the server configuration recovery operation is complete, you can start the volume configuration recovery operation.

Step 2: Restore volume ACL and quota configurations.

Select the **Recover Volume Security Settings** button and click **OK** to open the **Server Recovery** page. Select the volumes you want to restore. (Volumes that do not have a recovery file attached will appear as unavailable.) The creation date of the recovery file on a volume indicates when the recovery image was generated. Click **Recover** to start the operation, and then follow the on page instructions.

NOTE: Restoring the server and volume configurations will only occur if the system can accomplish it without compromising data. If the current configuration is different than previous SnapDRImage files, the restore will fail. A successful restoration requires either that the current RAID and volume configuration exactly match the saved one or that there is no current configuration (such as, you are using raw drives) so that the saved configuration can be recreated with the available raw drives. View the log file if the recovery operation fails.

Rejoining the Server to a Windows Domain

If you are restoring server settings to either the same physical server or to a replacement server, the server will automatically rejoin the Windows domain it was a member of before the SnapDRImage was applied as long as the servername is the same as the current servername. If you have changed the servername, you will have to manually join the server to the desired Windows domain. Navigate to **Network > Windows/SMB** to rejoin the server to a domain.

SnapDRImage Usage Scenario

The SnapDRImage contains the server configuration settings for a specific server. These settings may be useful in situations other than disaster recovery.

Reset Server Configuration After Swapping Out Components

In the extreme situation that one of the major components of your SnapServer fails (example: the mother board), and new one has been provided to you, it is possible to reset the sever configuration using the SnapDRImage. Follow the recovery procedures to utilize this option.

Cloning Servers

You can use SnapDRImages to copy server configurations from one server to another. For example: if you have a SnapServer DX1 configured for peak performance in your network environment, you can create a SnapDRImage of this server, then apply the SnapDRImage to a new SnapServer DX1. The server settings and configuration would be identical to the first. See “Cloning a Server” for details.

NOTE: Cloning a disk configuration will only succeed if the destination server has sufficient storage resources to duplicate the original configuration.

Replacing a Server

The procedure described in this section for responding to a catastrophic event is general in nature and may result in the loss of data. Should such an event actually occur, the exact procedure to follow will vary according to environmental conditions. Overland Storage strongly recommends that you contact a technical service representative before proceeding.

This section describes a worst-case scenario:

- The operating system has failed, (for example, due to a malicious attack to the root filesystem), and you cannot access the server.
- The data has been corrupted and must be restored from tape.
- Technical support has deemed your server unsalvagable and provided you with a new, unconfigured server.

Restoring Previous Server Settings to a New Server

After Technical Support has supplied you with a new server, you can restore the settings from the previous server to the new server. Any third-party license keys you have not purchased through Overland Storage will be lost. If you have installed data replication or management utilities such as Snap EDR, you will need to re-install and/or relicense them for use with the new server.

NOTE: If you are restoring Snap EDR Management Console settings, you must recreate the RAID and volume configuration that matches the DRI settings, then install and enable the Snap EDR Management Console. As an alternative, you can first restore just the system settings, install Snap EDR, and then restore just the Snap EDR settings.

You will also need to reschedule snapshots as well reconfigure CA Antivirus.

1. Point your browser to `http://<servername | IP address>/recover`.
2. Log in when prompted.
3. Navigate to **Maintenance > Disaster Recovery**, select **Recover System Settings** and click **OK**.
4. Click the **Browse** button and navigate to the SnapDRImage you made of the previous server, then click **Recover**.
5. The server reboots and the settings are restored. To view the log, click the date link on the **Disaster Recovery** page after the server has rebooted.
6. After restoring your server settings, rejoin the server to the Windows domain if necessary.
7. Now you can replace your data from tape backup. If the backup doesn't retain permission and ownership settings, you can restore these by selecting **Recover Volume Security Settings** on the **Maintenance > Disaster Recovery** page.

NOTE: If you are restoring from any backup, you will need to recover the volume settings.

Cloning a Server

The Disaster Recovery process can be used to clone a server in order to apply the same configuration to one or more servers. To clone a server:

1. Create a disaster recovery **image** on the source server (refer to **Creating the SnapDRImage and Volume Files**).
2. Copy the disaster recovery files from the source server to a **client**.

3. Perform a disaster recovery **restore** procedure to each of the clone target servers using the disaster recovery files from the source server (refer to Restoring Previous Server Settings to a New Server).

Data Import

Use the **Data Import** page to import (migrate) data from a legacy SnapServer or other computer that supports CIFS or NFS (v2 or v3) to a new SnapServer. The Data Import feature can be used to copy or move files and folders from a server on the network (Source) to your SnapServer (Target). To access the Data Import utility, navigate to **Maintenance > Data Import**.

The screenshot shows the SnapServer Data Import utility interface. The title bar reads "snapserver: Snap7654321 · Data Import". The interface includes a navigation menu with "Data Import" selected. Below the menu, there is a description: "Use data import to copy or move files and folders from a location on your network (Source) to this server (Target)."

The form is divided into three main sections:

- Source:**
 - Network Protocol: Windows (SMB) (specifies how to communicate with host)
 - Auth. Name: [text input]
 - Auth. Password: [text input]
 - Host: [text input]
 - Share: [text input] [Browse]
 - Path: [/] [Browse]
- Target:**
 - Volume: Volume1
 - Path: [/] [Browse]
- Options:**
 - Import Type: Copy (source data is maintained)
 - Include all sub-folders (if source path specifies a folder)
 - Overwrite existing target files and folders (that have identical names as the source files and folders)
 - Preserve file/folder permissions
 - Verify imported data (takes twice as long)

A note at the bottom states: "Note: You can setup Email Notification (administrative operation event) to be notified when a data import operation is complete." Buttons for "Start Import", "View Log", and "Close" are located at the bottom of the form.

If an error is encountered during the import (for example, a file or folder is locked and cannot be imported), the DM utility records the error in a log, and continues the operation. When the import is completed, the administrator can view the log of import errors. Once the errors have been corrected, the administrator returns to the DM main page, and recreates the import. With the exception of the password, all fields will still be populated with the specifications of the last job.

The following import options can be specified:

- Copy or move data
- Include subfolders
- Overwrite existing files
- Preserve the original permissions settings

NOTE: If you elect to preserve original permissions settings, be sure to review "Preserving Permissions".

- Verify imported data

NOTE: If you elect to verify imported data, all data will be read twice, once for import and once for comparison to the copied data. This could be a lengthy process.

Setting Up a Data Import Job

Before setting up a data import job, be sure to specify a user identity for the operation that will have full access to all files on the source, regardless of permissions set:

- For Windows import, specify an administrator or member of the Windows server/domain administrators group.
- For NFSv2/3 import, consider using the user root, and configuring the NFS export on the source to `no_root_squash` for the IP Address of the SnapServer for the duration of the import.

To create a data import job, perform the following procedure:

NOTE: Only one import job can run at a time.

1. On the **Data Import** page, complete the required **information** for both the source (legacy server) and target (SnapServer).

Option	Description
<i>Source:</i>	
Network Protocol	<p>Protocol that the SnapServer uses to connect to the source server. Select:</p> <ul style="list-style-type: none"> • Windows (SMB for Windows servers or GuardianOS servers with source data on a Windows root directory; default) <p>NOTE: If you are importing via SMB, SMB must also be enabled on the target SnapServer (go to Network > Windows/SMB to enable SMB in GuardianOS).</p> <ul style="list-style-type: none"> • NFS (NFSv2/3 for UNIX/Linux-based servers or GuardianOS servers with source data on a UNIX root directory)
User Identification	<ul style="list-style-type: none"> • If Windows was selected as the protocol, provide the Auth. (Authentication) Name and Password (Windows user name and password to log in to the server over SMB). • If NFS was selected as the protocol, provide the User Name (SnapServer local user name or NIS user, representing the UID used to perform the operation over NFS).
Host	Enter the name or IP address of the source computer you are importing data from.
Share/Export	<p>Specify the share (Windows) or export (NFS) on the source server containing the data you want to import.</p> <p>NOTE: Wildcards are not supported when specifying the source share to import.</p>
Path	<p>Enter the path to the file or folder you want to import. If you are importing the entire share, you can leave the Path field blank.</p> <p>NOTE: Wildcards are not supported when specifying the path to import.</p>
<i>Target:</i>	
Volume	Specify the volume on the target SnapServer where you want the data imported.

Option	Description
Path	Specify the path to the target SnapServer where you want the data imported.
<i>Options:</i>	
Import Type	Options for the import data are to Copy (source data is maintained) or Move (source data is removed during copy). If Verify imported data is enabled, the Move option removes the original data after the verification is complete. The default is Copy. NOTE: If you select to Move rather than Copy data, it is strongly recommended that you also select to Verify imported data.
Include All Sub-folders	If the folder you select for import contains sub-folders, selecting this option will import all files and folders underneath this folder (checked by default). If disabled, <i>only</i> the files directly in this folder will be imported.
Overwrite Existing Target Files & Folders	If files/folders on the target share identical names with files/folders on the source, checking this option overwrites those files/folders during import (checked by default.)
Preserve File/Folder Permissions	Selecting this option will retain the source permissions when the files/folders are imported to the SnapServer target (unchecked by default). NOTE: Before selecting this option, be sure to review Preserving Permissions.
Verify Imported Data	Selecting this option will cause all source data to be read twice, once to write to the target SnapServer and once to perform a binary comparison with the data written to the SnapServer (unchecked by default). If enabled, and if the Import Type is <i>Move</i> , files on the source will only be removed after verification. Otherwise, files will be removed immediately during the copying of them to the SnapServer. If you select to move files rather than copy them, it is strongly recommended that you enable the Verify imported data option. If a file mismatch occurs during verification, the target file is moved to a <i>data_import_verify_failures</i> directory on the root of the same volume. Check the failed file to determine the problem, then run the import again with Overwrite Existing Target Files & Folders deselected (so you don't re-copy files that have already been copied and verified). NOTE: Depending upon how much data is being imported, verifying imported data can be a lengthy process.
Email Notification	Clicking the email notification link will take you to the Email Notification page (for more information, see " Email Notification "). Fill in notification information and check the box next to Administrative Operation Event in order to receive an email when the import operation is complete.

- Once you have completed the import information, click the **Start Import** button to begin the import. You can see the progress of the import, an estimated time until completion, and the Import log on the **Data Import** page as it is compiling.
- When the import is complete, click the **View Log** button to see details of all errors. Click the **Data Import Error Log** link to download the entire log.

Stopping an Import Job

To stop the import at any time, click the **Stop Import** button on the **Data Import** page. If a file was in the process of being copied, the partially-copied file on the target will be removed.

Recreating an Import Job

The Data Import log records all errors that occurred during import. You can import files and folders that were not imported during the original job because of an error condition (for example, the file was locked).

1. Review the Data Import errors log and correct all error conditions.
2. Reopen the **Data Import** page. All fields (except the password) for the last import will still be visible on the page.
3. Click **Start Import** to run the import again. By default, all files will be re-imported. If you want only to import those files that failed to import the first time, you can disable the **Overwrite existing target files** option. However, make sure that any problematic files during the first import are deleted from the target SnapServer so they will be re-imported.

NOTE: If an import failed, it is strongly recommended that you enable the **Verify imported data** option for the re-importation.

Preserving Permissions

The types of permissions retained will differ, depending on which import scenario is applied.

Importing from a Windows Security Model to a Windows Root Directory

If you are importing from a Windows server (or other type of server that follows the Windows security model) to a Windows root directory on a SnapServer, permissions will be retained exactly as they exist on the source. However, as is the case when moving files with permissions between Windows servers, permissions for users that are unknown on the target server will be retained but not enforced. This includes permissions for:

- Local users on the source machine.
- Domain users for domains unknown to the SnapServer (for example, trusted domains, if the SnapServer is not configured to support trusted domains).
- Certain built-in Windows users and groups.

Importing from a UNIX Security Model to a UNIX Root directory

If you are importing from a UNIX server to a UNIX root directory, UNIX permissions for UIDs/GIDs are copied exactly from source to target; thus, identities of the users and groups will be best retained if the SnapServer belongs to the same NIS domain as the UNIX server.

Importing Between Conflicting Security Models

When importing from a UNIX source to a Windows root directory, UNIX permissions will be retained and the security personality on the resulting files and directories will be UNIX.

However, when importing from a Windows source to a UNIX root directory, permissions cannot be retained (since UNIX root directories are required to be UNIX personality throughout). Files and directories will inherit the UNIX personality and will have a set of default UNIX permissions.

Importing from a GuardianOS Server

When importing from one GuardianOS server to another, it is recommended that you maintain the same security model on the target server that you have on the source.

- If your source server uses a Windows root directory and has permissions assigned to Windows domain users, use a Windows connection for import. Windows permissions will be retained exactly as they are on the source, with the same enforcement limitations for unknown users as for import from Windows servers (see [Importing from a Windows Security Model to a Windows Root Directory](#)).

NOTE: If importing from a pre-5.0 GuardianOS server, Windows permissions will be retained verbatim, but may have different meaning due to the differences between the pre-5.0 POSIX ACL security model and the Windows security model introduced in 5.0.

- If your source server uses a UNIX root directory and has permissions assigned to local or NIS users, use an NFS connection for import.

NOTE: Local users that have UNIX permissions on the source will not be created on the target with the same UIDs.

OS Updates

Use this page to install updates to GuardianOS and other installed software, and to configure GuardianOS to automatically check for updates to GuardianOS and Snap EDR.

Information about the last GuardianOS update is listed at the bottom of the page, and may include the status of the update, product and version, and the completion time.



CAUTION: Do not interrupt the update process. You may severely damage the server if you interrupt a software update operation.

Update the GuardianOS Software

1. Click the **Check for Updates** button. If an update is available, follow the instructions on the page to download it.

NOTE: If the server does not have access to the Internet, download the latest GuardianOS image or other software package from the [Overland Storage website](#).

2. Click **Browse** on the **OS Update** page, locate the downloaded file, and select it.
3. Click **OK**. The SnapServer uploads the software package and then prompts you to reboot the server to perform the upgrade. Or click **Cancel** to stop the update.

Software Update Notification

You can configure GuardianOS to display an alert when GuardianOS or Snap EDR updates are available for the server. When enabled, Update Notification checks weekly for GuardianOS or EDR updates that are applicable to the server. If updates are available, a banner alert will display just below the menu bar on all Web Management Interface pages.

NOTE: You can choose to hide the banner by clicking the [Remind me later](#) or [Hide this message](#) link on the banner. If [Remind me later](#), the server will display the banner after the next check for updates; if [Hide this message](#), the server will hide the banner for the update in question until a later version is released.

Configuring Update Notification

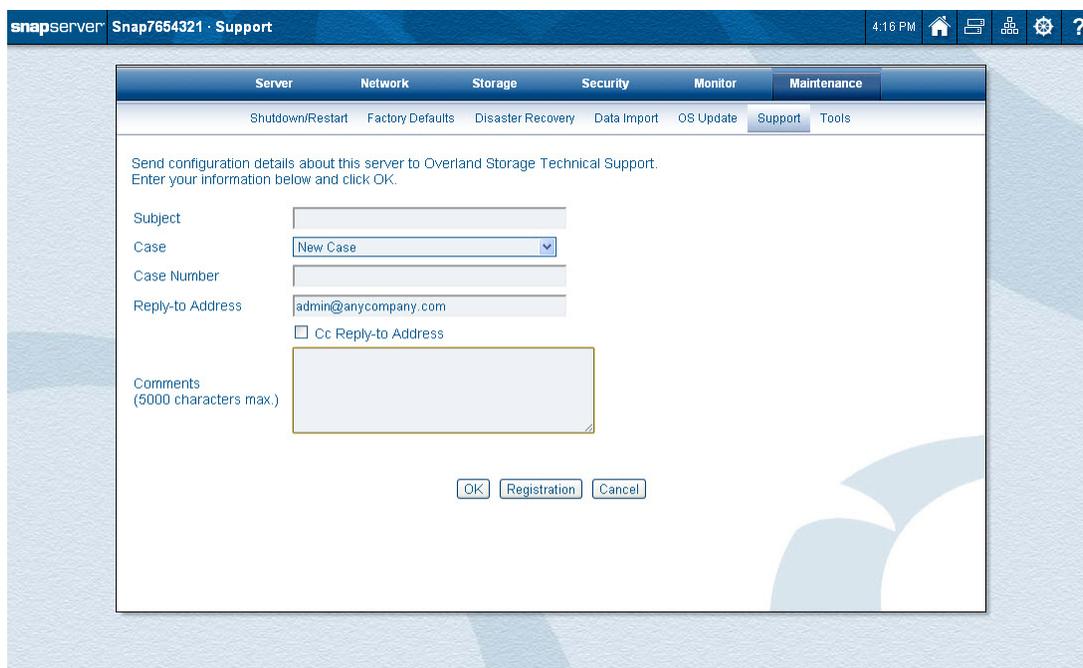
1. Click the **Update Notification** button.
2. Click to put a check in the **Enable Automatic Update Notification** check box.
3. If your environment requires using a proxy server for external web-based communication, check the **Use a proxy server for HTTP communication** check box and complete the Proxy Host and Proxy Port fields.
4. Click **OK**.

Checking for Updates

Click the **Check for Updates** button to force the server to immediately search for applicable updates. If an update is available, it will be displayed with information about it and a link to download the software.

Support

The **Support** page provides an easy way to contact Overland Technical Support.



The screenshot shows the SnapServer web interface. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'Maintenance' tab is active, showing sub-tabs for 'Shutdown/Restart', 'Factory Defaults', 'Disaster Recovery', 'Data Import', 'OS Update', 'Support', and 'Tools'. The 'Support' sub-tab is selected, displaying a form titled 'Send configuration details about this server to Overland Storage Technical Support. Enter your information below and click OK.' The form fields include: 'Subject' (text input), 'Case' (dropdown menu with 'New Case' selected), 'Case Number' (text input), 'Reply-to Address' (text input with 'admin@anycompany.com' entered), a checkbox for 'Cc Reply-to Address', and 'Comments (5000 characters max.)' (text area). At the bottom of the form are 'OK', 'Registration', and 'Cancel' buttons.

Once your SnapServer has been registered, Phone Home Support becomes available for use. Phone Home Support emails system logs and files that contain information useful for troubleshooting purposes to Overland Storage technical support. See [“Phone Home Support”](#).

Registering Your Server

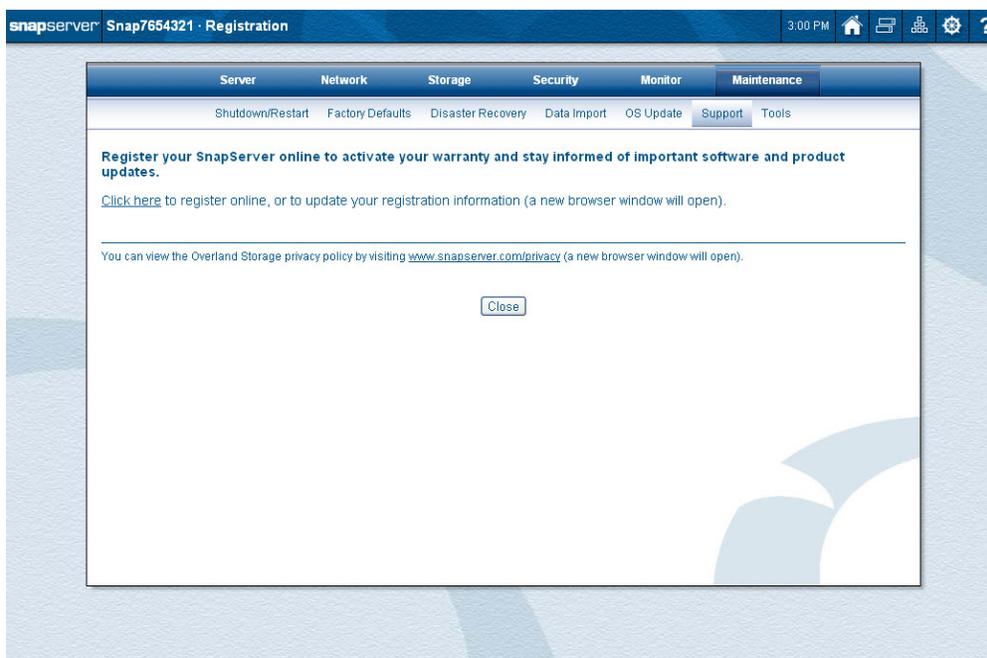
Registering your server activates your warranty and allows you to create and track service requests. Registration also provides access to GuardianOS upgrades, third-party software, and exclusive promotional offers.

NOTE: Warranty information is available at <http://docs.overlandstorage.com/support>

To Register Your Server

NOTE: To use this feature, access to the external Internet is required.

Go to **Maintenance > Support > Registration** and click the **Click here** link to launch the online registration page.



The same page is also used to update your registration information. Once you have registered, you will receive a confirmation email.

Maintenance Tools

The tools under this Maintenance subheading provide general-purpose server maintenance for both volume and root filesystems.

Email Notification

To configure the server to send email alerts in response to system events, navigate to the **Maintenance > Tools > Email Notification** page. To set up email alerts, you will need: (1) the SMTP server's IP address; and (2) the email address of each recipient to receive an alert.

Configuring Email Notification

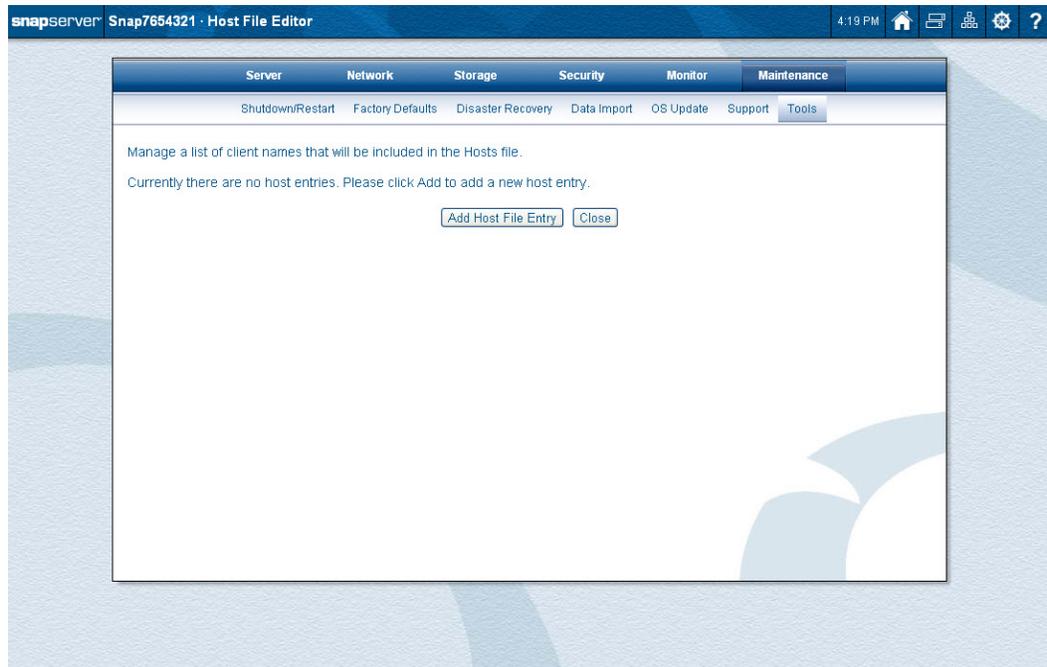
Edit settings as described in the following table, and then click **OK**.

Option	Description
Enable Email Notification	To enable email notification, check the Enable Email Notification check box.
SMTP Server	Enter a valid SMTP server IP address or host name.
SMTP Port	Enter a port number for the SMTP server or accept the default.

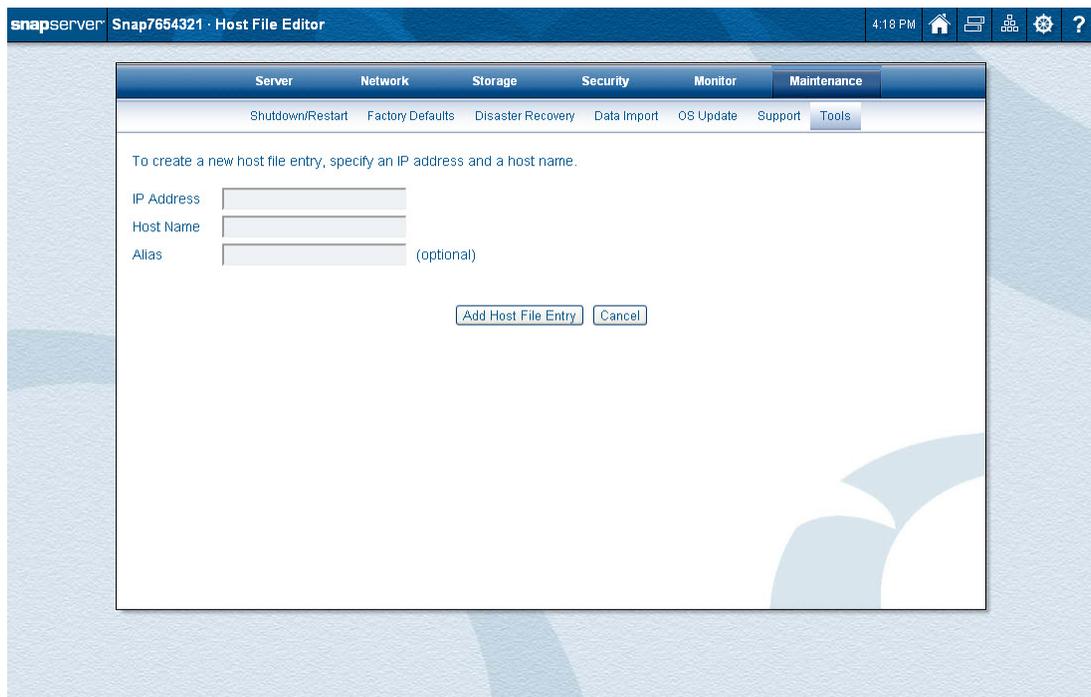
Option	Description
Use Authenticated SMTP	Check this box to require authentication when an email is sent to the SMTP server by the SnapServer. Provide an authentication User Name and Password in the fields that appear when the feature is enabled.
Use Secure Connection	Check this box to encrypt emails from the SnapServer. STARTTLS and TLS/SSL encryption protocols are supported.
Email Address of Sender:	Choose: <ul style="list-style-type: none"> The default address (<i>servername@domain</i>) where the <i>domain</i> is the DNS domain name. If there is no DNS domain name, then the server's IP address for Eth0 will be used (<i>servername@ipaddress</i>) Specify a specific sender.
Recipients	Enter one or more email addresses to receive the notifications. One address is required. Three additional email addresses can be added.
Send Email Notification	Check the boxes next to the events you wish to be notified about: <ul style="list-style-type: none"> Server shutdown/restart – The server shuts down or reboots due to an automatic or manual process. RAID Set event – (1) A RAID 1 or 5 experiences a disk drive failure or a disk drive is removed; or (2) A RAID 1 or 5 configures a spare or a new disk drive as a member. Volume is Full – Storage space on a volume reaches 95% utilization. Hardware event – The internal temperature for the server exceeds its maximum operating temperature or other hardware problems. Printing event – A printer error occurs (for example, the printer is out of paper). Administrative operation event – A Data Import operation has finished or experienced an error. License event – One of the trial licenses included on the SnapServer is about to expire. A notification email will be sent 14 days before the license expires. One day before the license expires another email will be sent. It is recommended that, if you are not acquiring a license key for the SnapExtension that is expiring, you turn off the SnapExtension.
Send a Test Alert	To verify your settings, check Send a test email , then click OK.

Host File Editor

Use this page to identify backup or media servers in the SnapServer's hosts file. This page allows you to supply a hostname-to-ip address mapping that persists across system reboots.



Click **Add Host File Entry**, complete the fields as described on the following table, and then click **Add Host File Entry** again.



Option	Description
IP Address	The IP address of the backup server.
Host Name	Enter the fully qualified address for the backup server, using the format: <i>myserver.mydomain.com</i> . NOTE: Your backup software may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.
Alias (optional)	Enter an optional abbreviated address for the backup server, using the format: <i>myserver</i> . NOTE: Your backup software may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.

Checking Filesystems

Filesystems on individual volumes can be checked for errors and repaired, if necessary. The root volume filesystem can also be checked, and any errors found will automatically be repaired. Because GuardianOS automatically checks the root volume for errors if any of a number of triggers occurs (for example, a power outage or failure of the volume to mount), it is recommended that the root filesystem check feature only be used when directed by a Technical Support representative.

To Check the Filesystem on a Volume

Checking Filesystems (**Maintenance > Tools > Check Filesystem**) provides a thorough filesystem check on the volume.



IMPORTANT: To begin the check operation, the volume you select is taken offline and access to the volume's data is unavailable until the operation is complete.

1. In **Maintenance > Tools**, click **Check Filesystem**.
2. From the drop-down list, select the **volume** (Traditional RAID) or **storage pool** (DynamicRAID) to be checked.
3. Choose the **type** of repair operation:
 - **Do not repair errors** – Checks for errors, but does not repair them. It is recommended that you do this periodically, especially following a power outage or any other unconventional incident.
 - **Repair errors** – Repairs standard filesystem errors. It is recommended that you run this level if you suspect filesystem damage may have occurred (for example, if a previous **Do not repair errors** operation reported filesystem errors).
 - **Repair errors (aggressive)** – Attempts to repair severe filesystem corruption.



CAUTION: It is only recommended that you run this level if you have been advised to do so by SnapServer Technical Support, or if **Repair errors** has failed to solve the problem and you are willing to risk loss of data.

4. Click **Check Filesystem**.

Checking a filesystem may require a reboot of the server in some circumstances. If prompted that a reboot is required, click **Yes**.

5. To view a log of the results, click the **View Log** button after the filesystem check completes.

To Check the Root Filesystem

Checking the Root Filesystem (**Maintenance > Tools > Check Root Filesystem**) provides a thorough filesystem check on the root.



CAUTION: Checking the root filesystem requires a reboot of the server.

1. In **Maintenance > Tools**, click **Check Root Filesystem**.
2. On the page that opens, click the **Check Root Filesystem** button.
3. Click **Yes** when informed that a reboot is required.
4. After the server reboots, to view a log of the results, click the **View Log** button.

The GuardianOS site map provides links to all the web pages that make up the Web Management Interface. It also provides, in the last column, special links to higher level options and processes which is the focus of this chapter.

These options are also directly navigable from the various menus in the Web Management Interface, and Home, Snap Finder, SnapExtensions, Site Map, and Help are accessible from any page by clicking their respective icon in the top right corner of the screen (see the table on page 2-11).

snapserver						
Server	Network	Storage	Security	Monitor	Maintenance	Misc.
Server Name	Information	Storage Pools	Security Guides	System Status	Shutdown/Restart	Admin Home
Date/Time	TCP/IP	Volumes	Shares	Active Users	Factory Defaults	Web Home
SSH	Windows	> Create Volume	> Create Share	Open Files	Disaster Recovery	SnapExtensions
UPS	Apple	Snapshots	Local Users	Event Log	Data Import	Snap Finder
Printing	NFS	> Create Snapshot	> Create Local User	Tape	OS Update	> Snap Finder Properties
	NIS	> Snapshot Schedules	> Password Policy		> Update Notification	Change Password
	FTP	iSCSI	Local Groups		> Check for Updates	Mgmt. Interface Settings
	SNMP	> Create iSCSI Disk	> Create Local Group		Support	
	Web	> VSS/VDS Access Control	Security Models		> Registration	
	iSNS	> Add VSS/VDS Hostname	ID Mapping		Tools	
		Disks/Units	Home Directories		> Email Notification	
					> Host File Editor	
					> Add Host	
					> Check Filesystem	
					> Check Root Filesystem	

Topics in Misc. Options

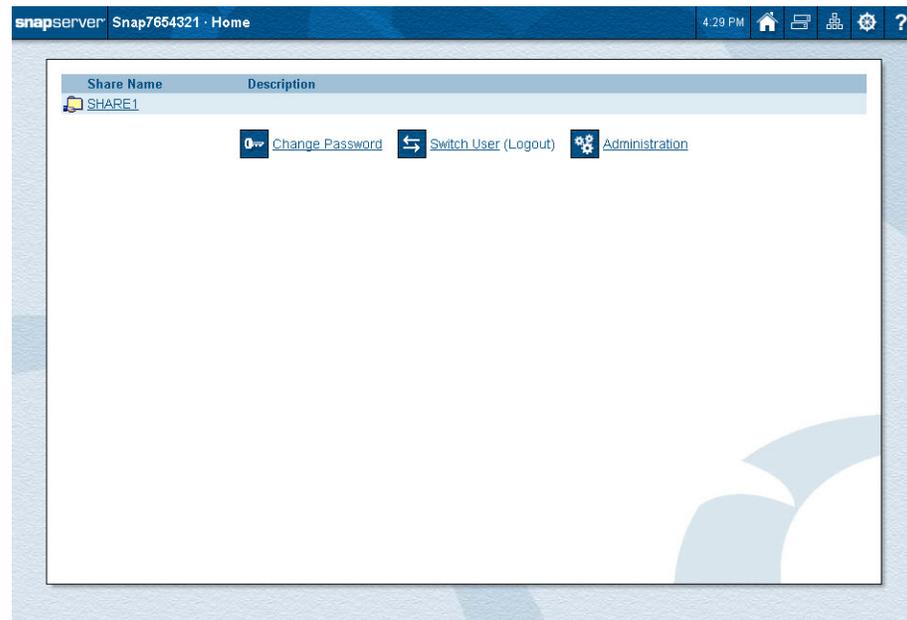
- [Home Pages – Web/Admin](#)
- [SnapExtensions](#)
- [Snap Finder](#)
- [Change Password](#)
- [Mgmt. Interface Settings](#)

Home Pages – Web/Admin

When you first log in to the GuardianOS, the Web Home page is displayed. Once logged in, you can switch between the Web Home page and the Admin Home page using the Home (🏠) icon.

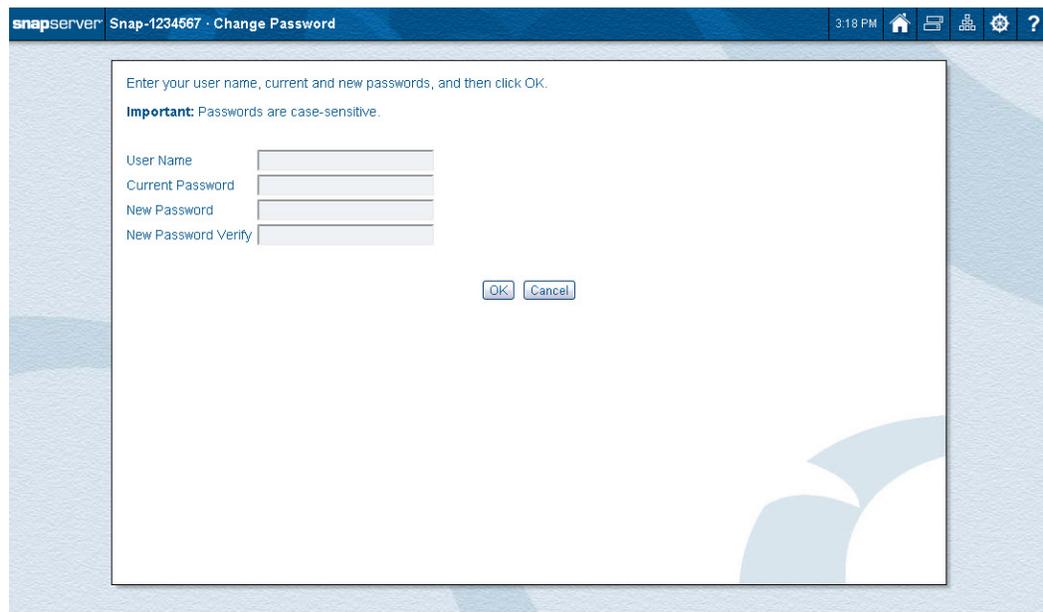
Web Home

The Web Home page shows a list of all the shares on the SnapServer.

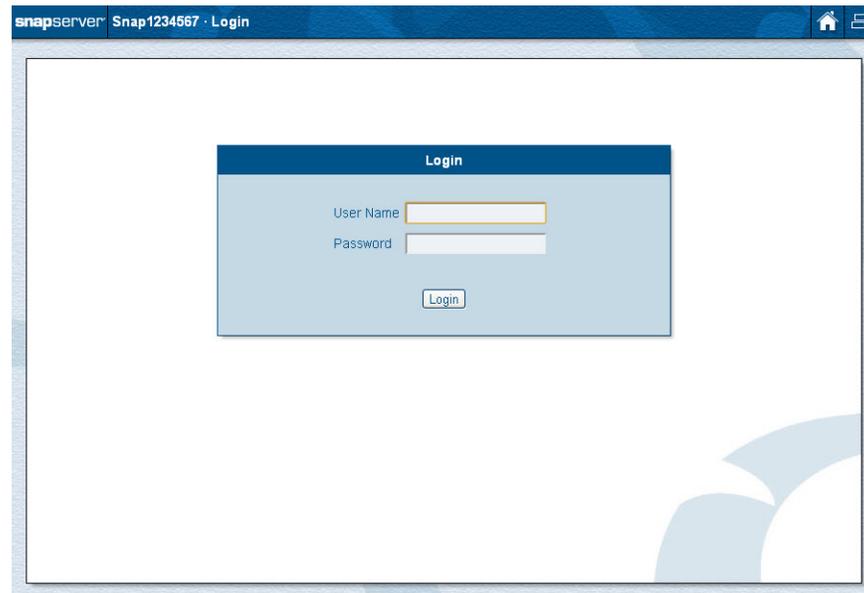


This page also provides three key administrative function links:

- **Change Password** (key icon) – Takes you to the **Change Password** page where you can change your administration password. Enter your **User Name** and **Current Password** for access. See “Change Password”.



- **Switch User (Logout)** (↵) – Automatically logs out the current user and displays the **Login** page for the new user to gain access to the SnapServer.



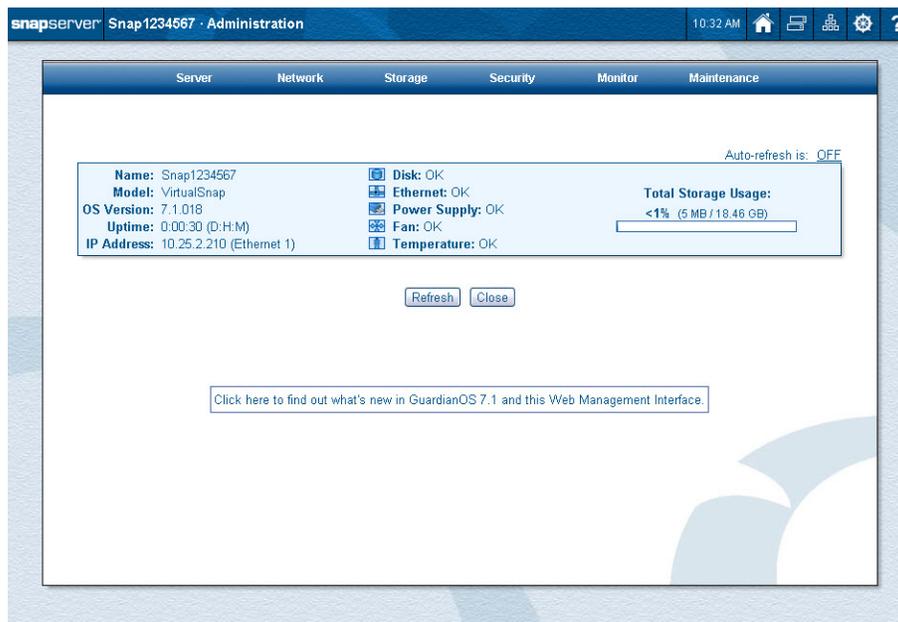
- **Administration** (⚙️) – Displays the Admin Home page (see “Admin Home”). You will be prompted to log in if you have not already done so.

If any of the following conditions are present, you may not be able to access the Home page:

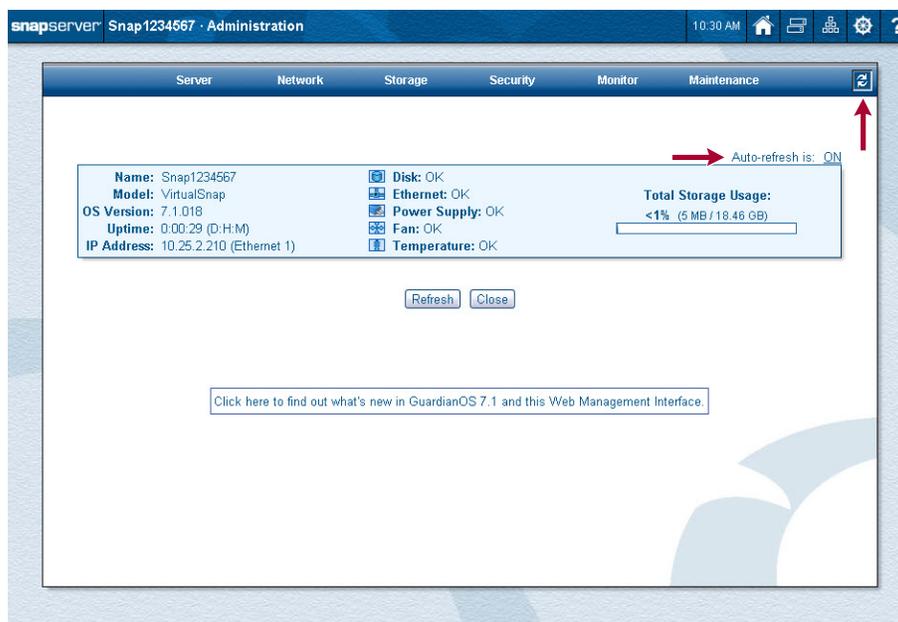
- **Require Web Authentication** is enabled (via **Network > Web > Require Web Authentication**) and you do not have a valid user name and password on the server.
- The server has not completed the **Initial Setup Wizard** (if this is the case, you will not be able to access the Administration page of the Web Management Interface either).
- **Web Root** is enabled (via **Network > Web > Enable Web Root**).

Admin Home

The Admin Home page is accessible by clicking either the Admin Home link in the Site Map or the Administration (⚙️) or Home (🏠) icons on the Web Home page. It provides a high-level view of the SnapServer status, the amount of total storage being used, and a link to find out what's new in GuardianOS. The tabs at the top provide access to the various functions and features of the GuardianOS.



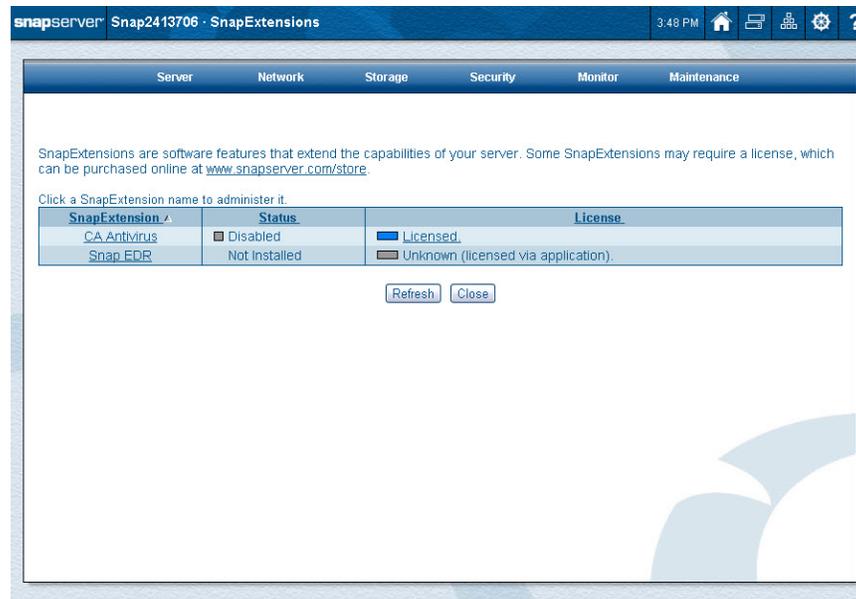
The **Auto-refresh** link on the right just above the Server Status box lets you select **ON** or **OFF**. When Auto-refresh is **ON**, the site information is automatically refreshed every 5 minutes, and an Auto-refresh icon (🔄) is displayed on the right corner just above the Server Status Box. Click the icon to manually refresh the information.



From the Administration page, clicking 🏠 takes you to the Web Home page.

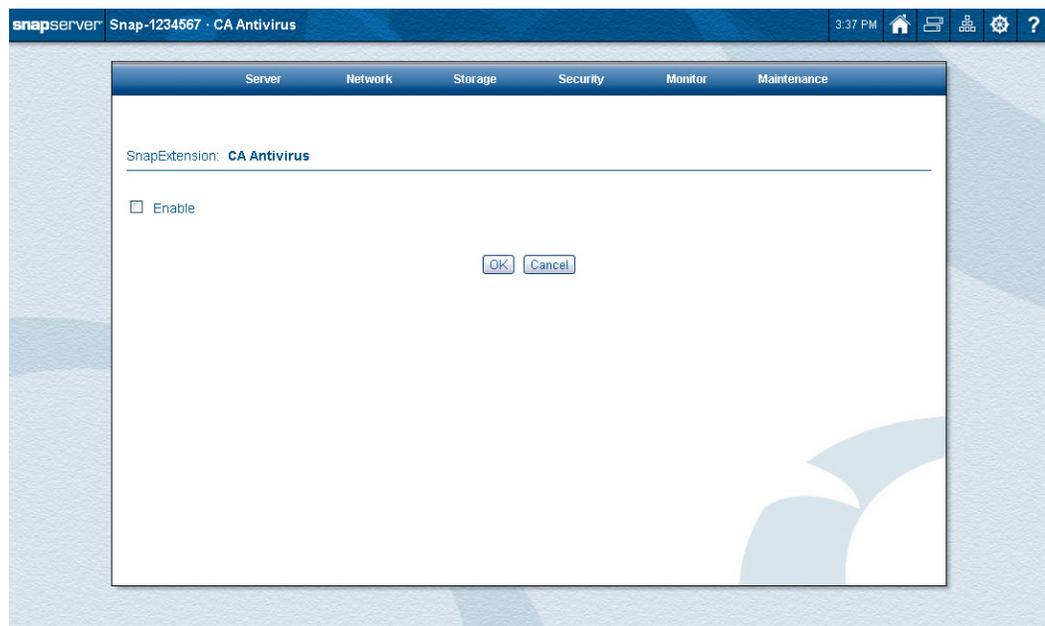
SnapExtensions

The SnapExtensions icon () opens the SnapExtensions page. This page is used to manage the SnapExtensions installed on your SnapServer.



If any SnapExtensions are installed, you can click the SnapExtension name in the table to display the management page for that extension.

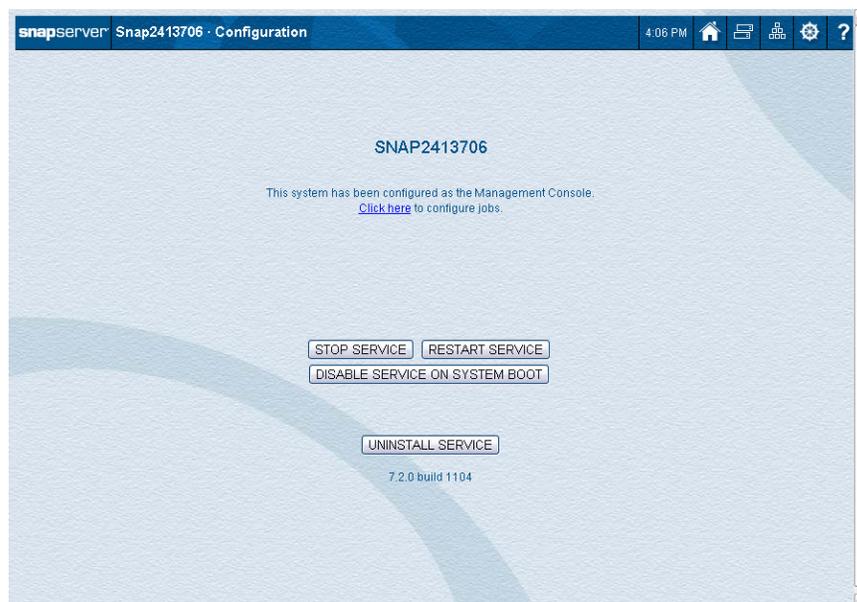
For **CA Antivirus**, click the box to enable it.



For **SnapEDR**, at the configuration page, select either to configure it as the Management Console or as an Agent of another Management Console. If configuring it as an Agent, enter the Name or IP of the Management Console.



After SnapEDR finishes its configuration, the Management Console screen is shown:



Snap Finder

Snap Finder () is a powerful tool that lists all the SnapServer and REO 4600 appliances on your network and on a remote network segment if so configured, and shows the current status. Click the server name (if you have server name resolution) or IP address of a server to access it through the Web Management Interface.

NOTE: You can sort the columns (ascending or descending order) by clicking the column head.

84 servers. Click a server name (if your network has name resolution) or IP address to connect to a server.

Server	Status	IP Address	OS Version	Model	Number	Avail Cap.	Total Cap.
athos	OK	192.168.48.38	6.0.043	4400	762797	182.88 GB	199.75 GB
beryl	OK	192.168.48.166	7.0.125	DX1	2300028	0.00 MB	0.00 MB
bm gala1	OK	10.25.2.78	7.0.125-kdb	VirtualSnap	14391761	3.68 GB	6.09 GB
bm gala2	OK	10.25.2.97	7.2.0-briansled11	VirtualSnap	11696003	2.00 GB	2.94 GB
bobbert	OK	10.25.2.134	5.0.133	4400	1723986	177.02 GB	280.71 GB
brees-110	OK	10.25.2.114	6.1.015	110	2250276	0.00 MB	0.00 MB
CCDragonflyDR	OK	192.168.45.107	6.5.023	N2000	730056	1.12 TB	1.71 TB
CCGalapagos-Q	OK	10.25.2.233	7.0.089	VirtualSnap	1381739	12.78 GB	47.26 GB
CCGalapagos6	OK	10.25.2.235	7.0.085	VirtualSnap	5944123	8.85 GB	8.86 GB
CCStorm920-1	OK	192.168.45.106	6.5.028	520	726420	1.09 TB	1.09 TB
CCStorm920-2	OK	10.25.2.238	4.4.049 SP3	520	719634	1.61 TB	1.61 TB
VM-Snap3752160	OK	10.25.2.219	3.0.0-psteesled11	VirtualSnap	3752160	0.00 MB	0.00 MB
VM-Snap4228931	OK	10.25.2.229	7.1.018	VirtualSnap	4228931	28.41 GB	28.41 GB
VM-Snap4381383	OK	10.25.2.103	7.1.019	VirtualSnap	4381383	28.41 GB	28.41 GB
VM-Snap4468555	OK	10.25.2.30	7.0.087	VirtualSnap	4468555	0.00 MB	0.00 MB
VM-Snap5248249	OK	10.25.2.47	3.0.017	VirtualSnap	5248249	42.07 GB	42.47 GB
VM-Snap6123337	OK	10.25.2.216	3.0.0-psteesled11	VirtualSnap	6123337	47.85 GB	47.87 GB
VM-Snap752896	OK	10.25.2.217	3.0.0-psteesled11	VirtualSnap	752896	47.85 GB	47.87 GB
VM-Snap7960161	OK	10.25.2.215	6.5.027	VirtualSnap	4643692	15.65 GB	15.66 GB
VM-Snap8706295	OK	10.25.2.43	7.1.0 autobuild_Dec19@20:21	VirtualSnap	8706295	28.41 GB	28.41 GB
VM-Snap9114734	OK	10.25.2.39	3.0.017	VirtualSnap	9114734	42.07 GB	42.47 GB
vmstore1	OK	10.25.10.116	5.2.056	520	710412	290.31 GB	1.78 TB
vmstore3	?	10.25.10.125	6.0.043	18000	900148	1.26 TB	2.72 TB

(You must have administrative privileges to view or edit Snap Finder properties.)

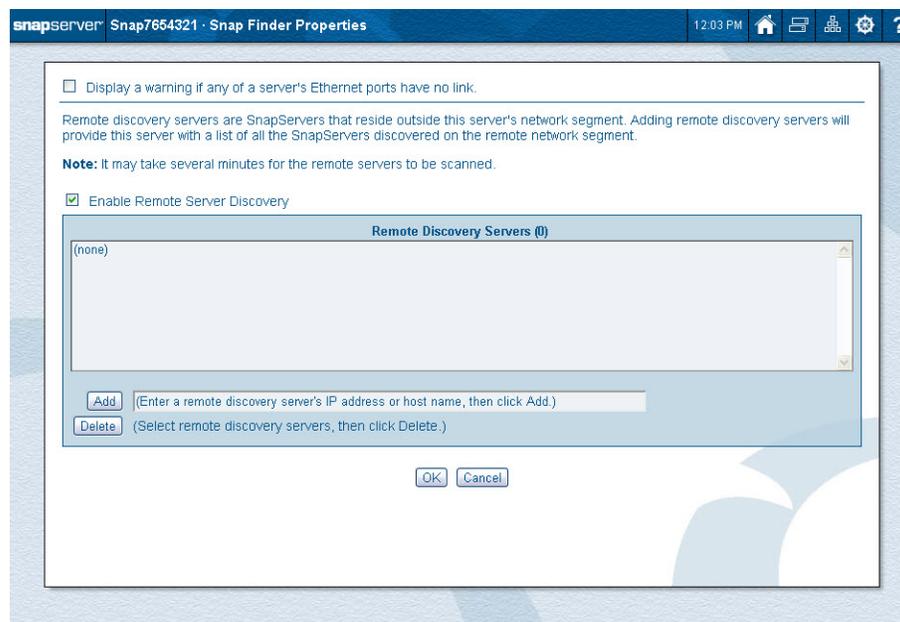
To enable remote discovery of SnapServers on a different subnet or to display a warning icon for servers with an enabled Ethernet port that has no link, click the **Properties** button to open the **Snap Finder Properties** page.

The following table details the columns in the table:

Identification	Description
Server Name	Current name of the server. The default server name is SNAPnnnnnn, where nnnnnn is your server number (for example, SNAP1234567).
Status	The status of the server (for example, OK, fan failure, or power failure).
IP Address	The IP address of the server.
OS Version	The version of GuardianOS currently loaded on the SnapServer.
Model	The SnapServer model.
Number	The Server Number derived from the MAC address of the primary Ethernet port, used as part of the default server name.
Avail Cap	The available capacity on the server.
Total Cap	The total capacity on the server.

Snap Finder Properties

Anyone with administrative privileges can view or edit the Snap Finder properties. Click the **Properties** button to access the **Snap Finder Properties** page.

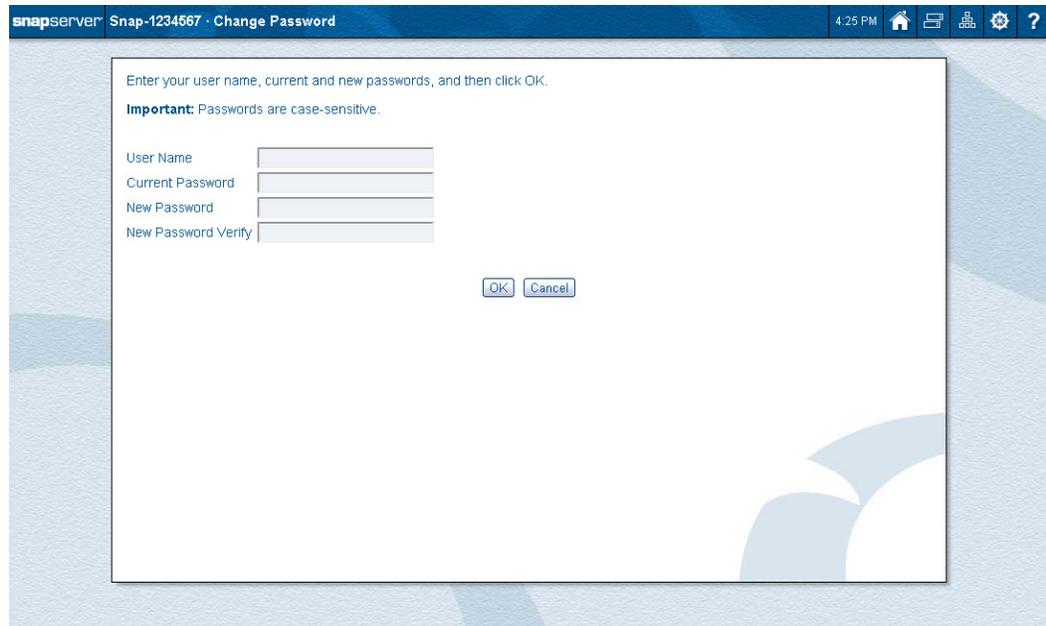


From this screen you can select to display a warning icon for servers with an enabled Ethernet port that has no link and enable remote discovery of SnapServers on a different subnet. Complete the following fields and then click **OK** to return to the Snap Finder screen:

Option	Description
Display warning if any of a server's Ethernet ports have no link	Check to display a warning icon in the Status column for any servers that have an enabled Ethernet port with no link. By default, this box is unchecked.
Enable Remote Server Discovery	Check to enable remote discovery of SnapServers on a different subnet.
Add Server	Enter the server's host name or IP Address in the field to the right of the Add button, and click Add to incorporate it into the list of Remote Discovery Servers.
Delete Server	Select a server in the Remote Discovery Servers field and click Delete . When asked to confirm the deletion, click Delete again.

Change Password

To enhance the security of your SnapServer, it is recommended that users change their passwords regularly. This is done using the **Change Password** page.



Enter your user name, current and new passwords, and then click OK.

Important: Passwords are case-sensitive.

User Name

Current Password

New Password

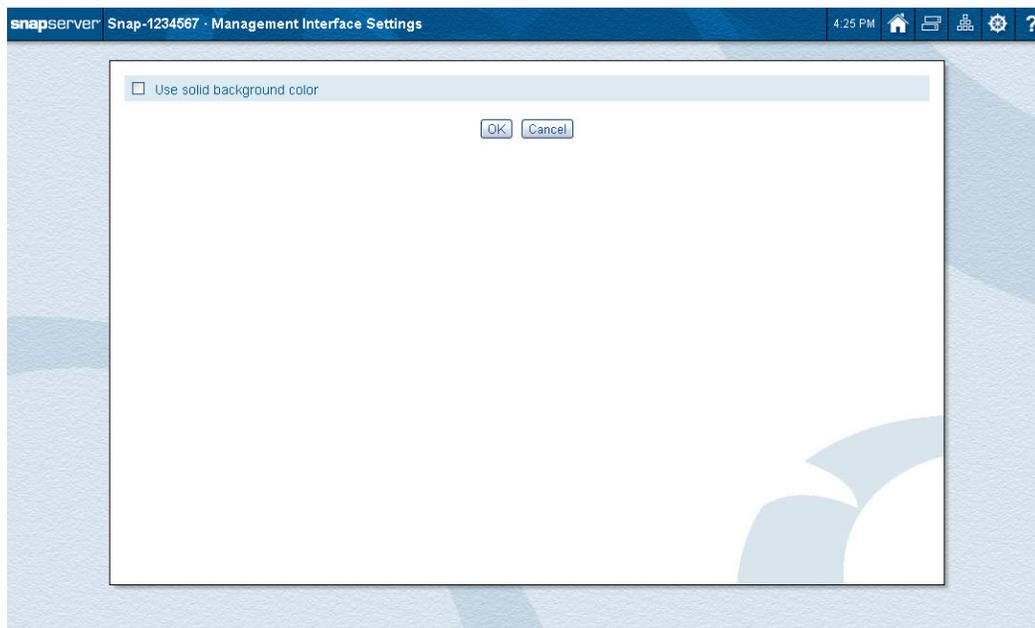
New Password Verify

Changing Your Password

1. On the **Home** page, click the **Change Password** link ().
2. At the **Change Password** page, enter your **User Name** and **Current Password**.
3. Enter and confirm your **new password**.
Passwords are case-sensitive. Use up to 15 alphanumeric characters without spaces.
4. Click **OK**.

Mgmt. Interface Settings

The Web Management Interface default background is light blue with the stylized “O” symbol. This can be changed to a solid blue background on the Web Management Interface Settings page by clicking the Site Map icon (⚙️), then **Mgmt. Interface Settings**, then **Use solid background color**.



The CA Antivirus software is preinstalled on all GuardianOS SnapServers. By default, the software is enabled on most SnapServers, but no scan jobs or signature updates have been scheduled. (The server will, however, check for signature updates whenever the server boots.) These and other antivirus configuration and management tasks are performed using the CA Antivirus GUI, accessed from the **SnapExtensions > CA Antivirus** page of the Web Management Interface. This section outlines the major steps in configuring the antivirus software. See the GUI online help for detailed descriptions of all options.

Topics in CA Antivirus Software:

- [Antivirus Dependencies](#)
- [Launching the CA Antivirus GUI](#)
- [The Local Scanner View](#)
- [Scan Jobs](#)
- [Signature Updates](#)
- [Alert Options](#)
- [The Move Directory](#)
- [Log View](#)

NOTE: Antivirus functions or options not relevant to the SnapServer have been disabled in the configuration GUI.

Antivirus Dependencies

The SnapServer implementation of CA Antivirus software includes the following features:

HTTP Access and Antivirus Configuration . To access the CA Antivirus configuration interface, HTTP must be enabled on the **Network > Web** page.

Resetting the Server Date and Time. If the current server date and time are changed to an earlier date and time (**Server > Date/Time**), the change does not automatically propagate to any scheduled antivirus operations. To synchronize scheduled antivirus operations with the new date and time settings, you must reschedule each operation.

NOTE: New jobs may be affected by the time change. Be sure to check that new jobs have been executed if a date or time change has been made to the server.

Storage Configuration and the Antivirus Software . The antivirus software resides on the largest volume (that existed at the time the software was installed). If you delete this volume, the CA Antivirus software will also be deleted. The SnapServer automatically reinstalls the antivirus software on the largest remaining volume on the system.

NOTE: The antivirus reinstallation process does not preserve custom antivirus configuration settings. Make a note of any such settings before deleting a volume.

Launching the CA Antivirus GUI

The CA Antivirus software on SnapServers is enabled by default. Some situations, such as deleting a volume or performing an upgrade procedure, may require you to re-enable the software. To learn how the antivirus software interacts with other GuardianOS software components, see “[Antivirus Dependencies](#)” on page 11-1.

NOTE: Antivirus functions or options not relevant to the SnapServer have been disabled in the configuration GUI;

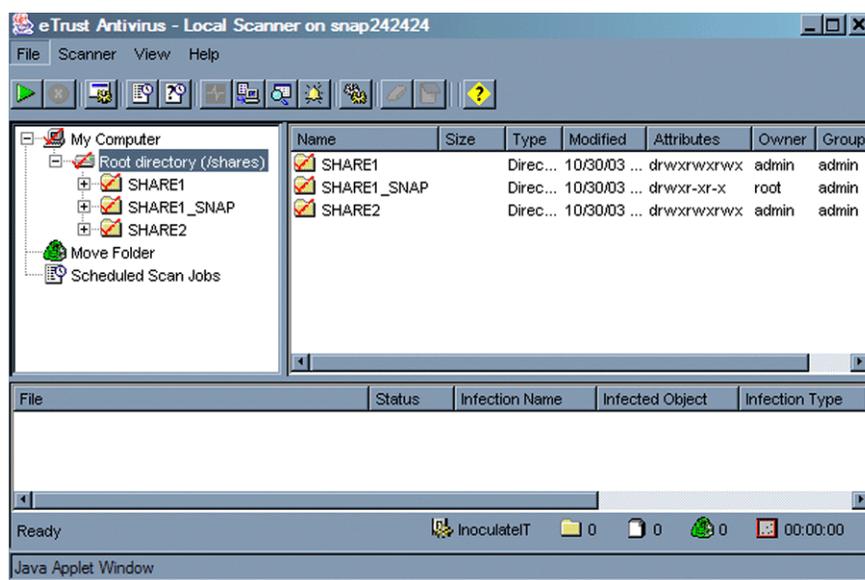
Launching the CA Antivirus Browser Interface

The first time you connect to the CA Antivirus GUI, it may take from 30 seconds to several minutes for the application to load, depending on the speed of your connection.

1. If you need to enable the antivirus software, go to **SnapExtensions > CA Antivirus**, click the checkbox next to **Enable**, and click **OK**.
2. Click the **Configure Antivirus** link. The splash page opens first, followed momentarily by the GUI login dialog box.
3. Enter the same administrative user name and password (case-sensitive) you have established for the GuardianOS Web Management Interface, and then click **Login**. The antivirus GUI box opens.

The Local Scanner View

Use the Local Scanner view to scan a SnapServer for infected drives, folders, files, or disks on demand.



Component	Description
Root Directory	Displays the directory structure of the SnapServer. As in Windows Explorer, click folder icons to navigate the structure and display subfolders and files in the right-hand pane.
Move Folder	May contain infected files. The administrator can instruct the software to automatically move infected files to this directory. For more information, see “Scan Jobs” on page 11-3 .
Scheduled Scan Jobs	Scan Jobs you schedule appear in this folder. For more information, see “Scheduling a Scan Job” on page 11-4 .

Scan Jobs

You can run scan jobs on demand or you can configure scan jobs to run periodically. This section outlines the process of configuring and running manual and scheduled scans. For detailed descriptions of all scanning options, see the CA Antivirus online help.

NOTE: You may not want to include Snapshot shares as part of your virus scan. Because access to an archived version of the filesystem provided by a snapshot share is read-only, you cannot treat or move any infected file; you would have to delete the entire snapshot to effect a cure. A more useful approach is to always scan your filesystem for viruses before running a snapshot. Adjust your antivirus scan schedule to synchronize with your snapshot schedule so that any infected files are cured (repaired) or removed before the snapshot is scheduled to run.

Defining Scan Jobs

This section provides an overview of the major choices available in configuring scan jobs. Access these options by selecting **Local Scanner Options** from the Scanner Menu.

Choosing an Infection Treatment (Scan Tab)

You can instruct the software to perform one of the following file actions when an infected file is found:

File Actions	Description
Report Only	(Default) Reports when an infection is found.
Delete File	Deletes an infected file.
Rename File	Renames an infected file with an AVB extension. Infected files with the same name are given incremental extensions (for example, FILE.0.AVB, FILE.1.AVB, and so on). After a file is renamed with an AVB-type of extension, it is not scanned subsequently.
Move File	Moves an infected file from its current directory to the Move directory for quarantine.
Cure File	Attempts to cure an infected file automatically. Choosing this setting enables the File Options button. Click this button to display the Cure Action Options and specify how the Cure File option performs.

NOTE: The *System Cure* option is not available on SnapServers.

Setting the Type of Files to Scan (Selections tab)

Use the Selections tab options to choose the types of objects to scan, the types of file extensions to include or exclude from a scan, and the types of compressed files to scan.

- **File Extensions** – You can choose to scan files regardless of extension, or select specific types of extensions to include or exclude.
- **Compressed Files** – To scan compressed files, select the **Scan Compressed Files** checkbox, and then click **Choose Type** to specify the compressed file extension types.

Filtering File Information for Logs (Manual Scans Only)

You can specify the types of events that are written to a log. Check the *Infected files* option to put information in the log about files that are found to be infected. Check the *Clean files* option to put information in the log about files that are scanned and are not infected. Check the *Skipped files* option to put information in the log about files that have been excluded from the scan.

Running a Manual Scan Job

Before running a local scan job, confirm that the scanner options are correctly configured as described in the previous section [“Defining Scan Jobs” on page 11-3](#).

Step 1: In Local Scanner View, select the folders you want to scan.

The left-hand pane displays the directory structure of the SnapServer. A red check mark on a folder or file indicates that it is selected for scanning. (By default, all directories and files are selected for scanning.) Click folders or files to toggle file/folder selection on or off.

Step 2: Run the scan.

Select **Scanner > Start Scanning**. The interface is unavailable for further configuration while the scan is in progress. The scan results display in the lower pane of the Local Scanner View, and the action taken with each file is listed in the Status column.

Scheduling a Scan Job

A scan job is configured and scheduled in the **Schedule New Scan Job** dialog box. To open this dialog box, select the **Scanner > Schedule Scan Job > Create** command.

Step 1: Set scan options in the Scan and Selection tabs.

These options are summarized in [“Defining Scan Jobs” on page 11-3](#).

Step 2: Schedule the scan.

The Schedule tab allows you to set a start date and a repeat interval for the scan.

Step 3: Select the directories to scan.

The Directories tab lists all paths that currently exist on the server. You can remove or add new paths as desired. You can also use the Exclude Directories tab to achieve the same result.

Step 4: Click OK.

You can view scheduled scan jobs by clicking the **Scheduled Scan Jobs** folder in the Local Scanner View. To edit a job, right-click it and select **Options**.

Signature Updates

Signature updates contain the latest versions of the signature files that recognize the latest infections. They also contain the latest engine versions, which do the work of looking for infections. Signature updates are made available on a regular basis by Computer Associates.

These updates are cumulative, so they contain everything from all previous file updates, plus the newest information on the latest infections. If you have missed a recent update, you only need to collect the latest signature file to have the most up-to-date protection.

SnapServers are preconfigured to download signature updates from the CA FTP site at <ftp://ftpav.ca.com/pub/inoculan/scaneng>. By default, no signature updates are scheduled. The antivirus software will, however, check for signature updates whenever the server is powered on. To update SnapServers that do not have Internet access, the following methods are available:

Method	Description
FTP	Use FTP to download the update files from the Computer Associates FTP site. You can also use FTP to distribute signature updates from one SnapServer (or any FTP server) to another. NOTE: When using FTP, the user name and password are passed as clear text.
UNC	Use UNC to distribute signature updates from one SnapServer to another (or from any arbitrary SMB or Windows server). Note that for UNC to work, you must have the Enable Guest Account option enabled (Network > Windows/SMB) on the SnapServer on which the signature updates reside. NOTE: Alternatively, you can distribute updates to SnapServers from any Windows/SMB server. If using this method, make sure the guest account on the chosen server exists, is enabled, and has a blank password.
Local Path	As part of the procedure to provide signature updates to the SnapServer with no Internet access, you can connect to a local path relative to the root (for example, /shares/SHARE1/virusdefs). Note that the path to the share is case-sensitive.

Updating SnapServers with Internet Access

If your SnapServers have direct access to the Internet, you only need to schedule the downloads to set up automatic signature updates. If access to the Internet is routed through a proxy server, you may also need to specify the name of the proxy server. Both procedures are explained below:

To Schedule Signature Update Downloads

1. Choose **Scanner > Signature Update Options**.
2. On the Schedule tab, click **Enable Scheduled Download**. Select the initial download date and time, then select how often to repeat the download.
3. Click **OK**.

To Specify a Proxy Server

1. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.

2. Select *FTP* in the list box, then click **Edit**.
3. In the Proxy Name field, enter the IP address of the proxy server, then click **OK**.

Updating a SnapServer without Internet Access

If you have SnapServers that do not have Internet access, use the following procedures to download the signature files to a machine with Internet access and then copy them to the SnapServer.

NOTE: When retrieving signature updates, the antivirus software attempts to connect to all the sites in the site list in the order they are listed. To avoid delays or superfluous error messages, delete the default FTP option from the list on SnapServers that have no Internet access.

1. Using a workstation with Internet access, go to <ftp://ftpav.ca.com/pub/inoculan/scaneng> and download the following files.
 - All **.tar* files containing the word *Linux*, for example, *fi_Linux_i386.tar* and *ii_Linux_i386.tar*
 - All **.txt* files containing the string *Sig*, for example, *Siglist.txt* and *Siglist2.txt*
2. Using a method appropriate to your environment, copy the update files to the SnapServer.
3. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.
4. Click the **Add** button, then select *Local Path* from the Method drop-down list.
5. In the Path field, enter the path to the directory on the server on which the update file resides. If you are using a SnapServer, the path would be similar to the following:
`/shares/SHARE1/sigfiles`
where *SHARE1/sigfiles* is the share path to the directory containing the signature update files.
6. Click **OK**. The path appears in the list box.
7. Click **Download Now**.

Distributing Updates from One SnapServer to Another

When retrieving signature updates, the antivirus software attempts to connect to all the sites in the site list in the order listed. To avoid delays or superfluous error messages, delete the default FTP option from the list on SnapServers without Internet access.

If you have more than one SnapServer with no Internet access, you can perform the previous procedure on just one of them (or any Windows/SMB server), and then configure your other SnapServers to get the update from that server automatically via SMB by specifying the UNC of the server containing the signature files.

The following conditions must be met in order to distribute updates using UNC:

- The correct Signature files must have been downloaded to the root of the share being used for updates.
- The server containing the Signature updates must have the Guest account enabled (**Network > Windows/SMB**) in GuardianOS. For other SMB/CIFS servers, the Guest account must have no password, and there may be additional requirements (for example, Windows servers must allow anonymous connections).
- The share and Signature files must be accessible to the Guest account.

- The server name used in the UNC must be resolvable by the server running CA Antivirus.

To Distribute Files via UNC

1. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.
2. Click the **Add** button, and select **UNC** in the Method list box.
3. Enter the path to the SnapServer (or Windows/SMB server) to which the update files have been downloaded (see previous procedure) using the following format:
`\\server_name\share_name`
where *server_name* is the name of the server, and *share_name* is the name of the share providing access to the files. (On a SnapServer, the update files must reside on the root of the share.)
4. Click **OK**. The path you entered appears in the **Download Sources** list box.
5. Click **Download Now**.

To Distribute Files via FTP

If you have more than one SnapServer with no Internet access, you can perform the FTP download procedure on just one of them (or any FTP server), and then configure your other SnapServers to get the signature updates from that server automatically via FTP.

1. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.
2. Click the **Add** button, and select **FTP** in the **Method** list box.
3. Enter the following information regarding the server on which the update file resides as follows:
 - In the Host Name field, enter the IP address.
 - In the User Name and Password fields, enter the admin user name and password.
 - In the Remote Path field, enter the path to the directory in which the file resides. If you are using a SnapServer, the path would be similar to the following:
`/shares/SHARE1/sigfiles`
where **SHARE1/sigfiles** is the share path to the directory containing the signature update files.
4. Click **OK**. The path you entered appears in the Download Sources list box.
5. Click **Download Now**.

Verifying Download Events

Use the following procedure to verify download and distribution events.

1. Select **View > Log Viewer**.
2. In the left-hand pane, select **Distribution Events**. Distribution events are listed in the upper right-hand pane in chronological order.
3. Select a distribution event. The details of the distribution event display in the lower pane.

Alert Options

Alert options allow you to tailor the notification information that is provided to the Alert Manager, cut down on message traffic, and minimize the dissemination of notifications that are not critical. To set alert options, select **Alert Options** from the Scanner menu. The Alert Options dialog box contains the following tabs:

Tab	Description
Report	Use the Alert Report options to specify where to send notification information, and the Report Criteria options to manage how frequently messages from the General Event Log are reported. NOTE: The Local Alert Manager option is not supported on SnapServers.
Alert Filter	Use the Alert Filter options to manage notification severity levels, and to determine what types of messages should be passed to the Alert Manager. NOTE: In the Custom Notification Module, the <i>Realtime Server</i> and <i>Admin server</i> settings have no effect on SnapServers.

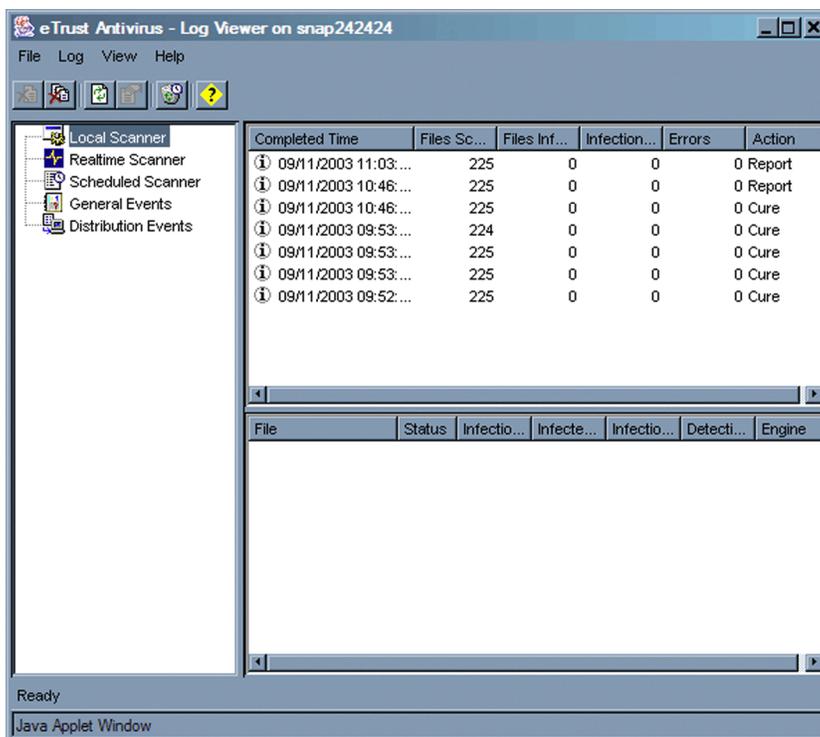
The Move Directory

You can configure scans to move infected files to the move folder (**Scanner > Local Scanner** options). To view infected files, click the **Move** directory on the left-hand pane of the Local Scanner View. To manage a moved file, right-click the file and select from the following options:

Option	Description
Restore	This option removes the file from the Move Folder and restores it to its original location with its original name and type.
Restore as	This option displays a dialog box that allows you to change the directory location and file name. You can rename a file and isolate it safely in a different location. You may want to use this option, for example, if you do not have another source for the data and you need to look at the file. Or you may have a file that you want to analyze. NOTE: To restore a file to a different directory, you must prepend the path to the directory with the string <code>"/shares."</code> For example, to restore a file to the <code>SHARE1/sales</code> directory, enter the path as follows: <code>/shares/SHARE1/sales</code>
Restore and Cure	This option allows you to restore the selected item back to the original folder it was in, and cure it. This option is useful if you update the signature files after items have been put in the Move folder. If a cure is provided that you did not have available, you can get the latest signature update and use this option to restore and cure an infected item.
Delete	This option deletes the infected file; no warning or confirmation message is displayed.

Log View

The Log View provides easy access to detailed information on scan, distribution, and other events. To access this view select **Log View** from the View menu.



Option	Description
Local Scanner	Displays summary information about scan jobs that have run.
RealTime Scanner	Not supported.
Scheduled Scanner	Displays summary information on scheduled scans that have run.
General Events	Displays the Event log for a given day. Click a date to view all events that occurred that day.
Distribution Events	Displays distribution events by date. Click a date to view detailed information on the distribution event in the lower pane.

DynamicRAID Overview

You can configure your SnapServer in either DynamicRAID or Traditional RAID mode. The following section details the benefits of DynamicRAID, as well as providing guidelines to help you choose the mode that is best for your needs.

DynamicRAID is a powerful new feature that simplifies storage management and provides additional configuration options not available in Traditional RAID. A SnapServer can be purchased with any amount of initial storage (or number of drives), and more capacity can be added over time by inserting or replacing drives. Volumes can be added and removed at will, and all volumes share the same underlying pool of storage.



IMPORTANT: A SnapServer head unit or expansion unit supports only one storage pool created from its drives and contained within that enclosure. Multiple volumes can be created on that storage pool.

Topics in DynamicRAID:

- [About DynamicRAID](#)
 - [Should I use DynamicRAID or Traditional RAID?](#)
 - [Features Comparisons: DynamicRAID and Traditional RAID](#)
- [Setting Up DynamicRAID](#)
 - [Implementation](#)
- [Additional Information on DynamicRAID Sizing](#)

About DynamicRAID

- To increase the capacity of the storage pool when the SnapServer is fully populated, you can replace drives, one at a time, with larger-capacity drives. Replacing drives can only be done when the DynamicRAID is healthy.
- DynamicRAID has two forms – one with single parity (protects your data against a single drive failure) and one with dual parity (protects your data against two simultaneous drive failures). The parity model can be changed over time.
- Volumes on DynamicRAID are virtual and may be created almost instantaneously. They all share the same underlying pool of storage, so there is no need to worry about the size of the volume when created. At the administrator's discretion, volumes may be constrained in size so they cannot consume more than a defined limit. This limit can be adjusted or removed as required.

- DynamicRAID is comparable to Traditional RAID for both file-level and block-level access. All of its features apply equally to both file sharing and iSCSI volumes created on the SnapServer NAS system. DynamicRAID uses clear visible indicators on the drive bays to illustrate what can or cannot be done to that bay, thus reducing user error and negating the need for any required skill set or training for operation. Anyone can easily manage and maintain an expandable storage system.

Should I use DynamicRAID or Traditional RAID?

Use the following guidelines and the table on [page A-3](#) to help determine which RAID mode is right for you.

Step 1: Determine how much time and effort you want to spend managing your RAID configuration.

- If you have **little or no time** to manage your RAID solution, choose DynamicRAID.
- If you want to have more **direct control** over your storage configuration, conduct **manual tuning**, and **manually manage the RAID array**, choose Traditional RAID.

Step 2: Determine what kind of storage configuration you need.

- **Will storage requirements in your environment change over time?**
If so, DynamicRAID provides you with the flexibility to respond to these changing needs. For example, you can upgrade smaller drives to larger drives one by one. These drives will be automatically incorporated and will share the same storage pool.
If you plan to add disk drives, you can take advantage of the ability to change parity when you do so. You can optimize parity based on the number of drives inserted into the system. You can either increase parity by adding a new drive, or decrease the parity setting to expand storage space (and sacrifice redundancy).
- **Do you want to aggregate all disk storage on the head unit and attached expansions?**
If you do, choose Traditional RAID.
- **Do you need to use local or global spare drives?**
If you do, choose Traditional RAID.
- **Do you need to use user or group quotas?**
If you do, choose Traditional RAID.

Step 3: Will you need to choose your RAID type?

If it is necessary for you to choose your specific RAID type, select Traditional RAID. For example, you may want maximum speed but no redundancy, a [“RAID Options with Traditional RAID”](#) and thus want RAID 0.

For more information on RAID types, see , as well as [“Factors in Choosing a RAID Type”](#).

The following table summarizes some of the prospective decision factors in choosing DynamicRAID or Traditional RAID. This list is not exhaustive.

Choosing DynamicRAID or Traditional RAID

Feature	DynamicRAID	Traditional RAID
I need a simple, scalable, flexible RAID solution that takes little or no effort to manage.	✓	
I need to easily add more storage capacity as needed.	✓	
I need to be able to change the parity level over time.	✓	
I need to be able to adjust volume size as needed.	✓	
I need to be able to optimize parity based on the number of drives inserted into the system.	✓	
I need the SnapServer to configure and manage the RAID array size and parity.	✓	
I need to manually configure the RAID array.		✓
I need to manually tune my storage system for specific needs, such as RAID levels and/or storage types.		✓
I need user or group quotas.		✓
I need to mix different drives in the same chassis, and then group these different drives together to make a homogenous RAID.		✓
I need the filesystem to span multiple chassis.		✓
I need local and global spares.		✓
I need to do snapshot rollbacks.		✓

Features Comparisons: DynamicRAID and Traditional RAID

The following table compares the features of these two RAID types:

Feature	DynamicRAID	Traditional RAID
RAID Levels	Single- or dual-parity options that can be changed dynamically.	Manually created RAID sets 0, 1, 5, 6, or 10. Must delete and recreate to change.
RAID Creation	Automatic after selection of parity. Snapshot space is configured by the user.	Manual selection of drives, RAID set level, and snapshot space.
RAID Expansion	Can be expanded by adding drives to the SnapServer.	Can be grouped with other RAID sets to increase the space available to volumes.
Mixed Drive Capacities	Additional capacity on larger drives can be utilized within the constraints of single- or dual parity protection. Additional capacity on larger drives can be utilized if there are enough larger drives to satisfy the parity configuration of DynamicRAID.	Only the capacity equivalent to the smallest drive is used on each drive in the RAID set.

Feature	DynamicRAID	Traditional RAID
Mixed Drive Types	All drives in a given Storage Pool must be the same type of drive (for example, SAS 15K).	Different types of drives can be mixed in a head unit or expansion unit (using different RAID sets and volumes).
Volumes	Volumes consume space directly from the storage pool as data is placed on the volume, and allocated as needed.	Volumes allocate from the RAID upon creation of the volume, and volumes must be manually grown to increase space for data as needed.
Snapshots	Snapshots are by Storage Pool and can be mounted for individual file recovery.	Snapshots are by volume and can be mounted for either individual file recovery or volume rollback.
Data Storage Capacity	Data storage capacity for all volumes on a storage pool is limited by the size of the storage pool and/or the maximum size that is set on each volume.	Limited by the storage size on the head unit plus all the expansion units.
Filesystem Spanning	Filesystem is limited to a given Storage Pool.	Filesystem can span multiple RAID sets concatenated together using Instant Capacity Expansion (ICE) and RAID grouping.
Quotas/Size Limits	Volume size limits can be either specified or unlimited.	User and Group quotas can be specified for each volume.

Setting Up DynamicRAID

1. DynamicRAID detects all available disk drives on the SnapServer.
2. Select the parity setting:
 - One disk drive – No parity protection only.
 - Two or three disk drives – Single-parity protection only.
 - Four or more disk drives – Choose either single- or dual-parity protection.
3. The software configures the SnapServer.
4. Use the following options to fine-tune the configuration:
 - **Storage > Storage Pools** (see “[Storage Pools](#)”)
 - **Storage > Volumes** (see “[Volumes](#)”)
 - **Security > Shares** (see “[Shares](#)”)

Implementation

DynamicRAID streamlines the storage management experience. During the initial setup, when making the [RAID Type Selection](#), choose DynamicRAID and the type of parity desired. The SnapServer automatically configures the RAID array, and the user may optimize the parity according to the number of drives inserted into the system. A storage pool is then created that can be divided into volumes for different applications or user groups. These steps are described in detail in the following sections.

Storage Expansion

During the setup process, storage pools are created on the head unit and each expansion unit using all disk drives available in that unit. More capacity can be added to a SnapServer over time by inserting or replacing drives, then adjusting Storage Pool properties. Volumes can be added and removed at will, and all volumes share the same underlying pool of storage.

When adding drives to a storage pool, the Web Management Interface displays the estimated time required until the new drive will be available for storing data, and an estimate of the final capacity that will be available when it is ready. Once the drive has been added to the Storage Pool, any of the following may take place to maximize capacity:

- The filesystem may be expanded to cover the available space (see [“To Edit Volume Properties”](#)).
- The snapshot space may be expanded (see [“Adjusting Snapshot Space Size”](#)).
- Both the filesystem and snapshot space may be expanded within a unit. However, the storage pool on one unit cannot be expanded to a different unit.
- Neither the filesystem nor snapshot space are expanded, but the parity is increased (see [“To Add a Disk Drive to Upgrade Parity”](#)).

When a drive is replaced in the storage pool, DynamicRAID rechecks its size to determine if it is now larger than before the replacement. This way, drives in a DynamicRAID can be replaced with larger drives one at a time, and once enough drives have been replaced with larger drives to support the storage pool's parity setting, the additional space in the larger drives will become available.

Snapshots

DynamicRAID utilizes current GuardianOS technology and snapshots the entire storage pool. Provisioning for snapshots can be increased as the storage pool is grown to ensure the percentage of storage reserved remains consistent. The directories inside the snapshot that represent volumes can be shared individually by the administrator, rather than all at once, to provide a level of access control.

iSCSI Target Volumes

All iSCSI targets use current SnapServer technology. DynamicRAID maintains the iSCSI volumes on the storage pool in a location that is not visible to users.

Indicators

Drives can be inserted into the SnapServer NAS system at any time unless the user is specifically instructed not to do so.

Each drive bay has an associated indicator which can be either red, amber, or green. Indicators show the state of the storage pool. See [“LED Indicators”](#).

Additional Information on DynamicRAID Sizing

All the drives in a chassis are considered part of a single storage pool on that unit and are dynamically configured as such. The first drive detected in the storage pool is used as the basis for the drive-size characteristics of the storage pool. These characteristics center around whether the drive space can be evenly divided into either 300 GB- or 500 GB-sized partitions.

For example, a large capacity drive (such as 1 TB and 2 TB SATA drives) will use 500 GB partitions for the storage pool while a smaller capacity drive (such as 300 GB and 600 GB SAS drives) will use 300 GB partitions. Then, all the other drives in the storage pool (or added later) are configured using the same partition sizing.

Drives of different overall capacity may be added to the same storage pool as long as they have the same partition sizing (such as, 300 GB). However, the extra space on larger drives will only be available to the storage pool if there are enough larger drives to satisfy the storage pool's parity setting using the extra space. Otherwise, the extra space will not be available to the storage pool until more drives are added with the same larger capacity.

For example, adding a 3 TB SATA drive to a group of three 1 TB drives with single parity only adds 1 TB (2x500 GB) of space for a total of 4 TB. The extra 2 TB of space on the 3 TB drive is not available until enough 3 TB drives are added to satisfy parity. Adding two more 3 TB drives opens up the additional 2 TB (4x500 GB) of space on all the 3 TB drives. The total storage pool then expands to 12 TB.

For more information on DynamicRAID, go to <http://docs.overlandstorage.com/dynamicraid>.

Backup Solutions

This appendix provides a brief description of the supported backup solutions and, where applicable, gives instructions on how to install the solutions on the SnapServer.

Topics in Backup Solutions:

- [Backup and Replication Solutions Table](#)
- [Integrated Backup Solutions](#)
- [Off-the-Shelf Backup Solutions](#)
- [iSCSI Disk Backups](#)

Backup and Replication Solutions Table

GuardianOS supports several backup methods, including third-party off-the-shelf backup applications and applications that have been customized and integrated with GuardianOS on the SnapServer.

NOTE: Unicode limits some backup applications' ability to function with the SnapServer.

GuardianOS	Backup and Replication Solutions				
	Snap EDR	CA BrightStor	EMC NetWorker	Symantec Backup Exec	Symantec NetBackup
Snap to Backup Server via installed agent		X	X	X	X
Snap to Backup Server via network protocol		X		X	
SnapServers to SCSI-attached tape drive (disk-to-tape backup)				X	X
SnapServers to SAS-attached tape drive (disk-to-disk backup)				X	X
SnapServers to USB-attached tape drive (disk-to-tape backup)				X	X
Security Meta Data Backup	X				

Integrated Backup Solutions

The [Snap Enterprise Data Replicator \(Snap EDR\)](#) is preinstalled and/or customized for the SnapServer.

Snap Enterprise Data Replicator (Snap EDR)

Snap EDR provides server-to-server synchronization by moving, copying, or replicating the contents of a share from one SnapServer to another share on one or more different SnapServers. It comes preinstalled on some servers with a 45-day free trial, or it can be downloaded from the SnapServer website.

Snap EDR consists of a Management Console and a collection of Agents. The Management Console is installed on a central system. It coordinates and logs the following data transfer activities carried out by the distributed Agents:

- Replicates files between any two systems including Windows, Linux, and Mac Agents.
- Transfers files from one source host to one or more target hosts
- Transfers files from multiple hosts to a single target host, and stores the files on a local disk or locally attached storage device.
- Backs up data from remote hosts to a central host with locally-attached storage.
- Restores data from a central storage location to the remote hosts from which the data was originally retrieved.

Configuring Snap EDR for GuardianOS

To configure the SnapServer as a Snap EDR Management Console or an Agent, do the following:

1. Click the **Snap EDR** link in the Site Map (under **Extras**).
2. Select either the **Configure as the Management Console** or **Configure as the Agent** button.

NOTE: If you are configuring the server as an Agent, you must provide the server name or IP Address of the Management Console.

3. Once the server is configured, a page appears with the following options:

Option	Description
Click here to configure jobs	Opens the Management Console where jobs can be scheduled.
Stop Service	Stops all services.
Restart Service	Restarts all services.  Caution: Use only if you have encountered a problem, and customer support advises you to restart the service. Any jobs currently running will stop and will not resume when you restart the service.
Disable Service on System Boot	By default, when a user reboots the SnapServer, services automatically restart. Select Disable Service on System Boot if you do not want the Snap EDR service to start up automatically. NOTE: When the disable service option is selected, the Enable Service on System Boot button appears.
Uninstall Service	Uninstalls all components of Snap EDR.

Scheduling Jobs in Snap EDR

To schedule jobs, click the **Snap EDR** link in the Site Map (under **Extras**). For complete information on using Snap EDR, see the *Snap EDR Administrator's Guide*, available on the SnapServer website.

Off-the-Shelf Backup Solutions

NOTE: These backup packages do not support the backup of Windows ACLs. If you use one of these packages, Overland Storage strongly recommends you create a SnapServer disaster recovery image (see "Creating the SnapDRImage and Volume Files") before you perform a backup.

In addition to the integrated backup solutions, GuardianOS supports a number of off-the-shelf backup packages that the user can install on the SnapServer, including:

- CA ARCserve 11.5, 12.0
- EMC NetWorker 7.3, 7.4
- Symantec Backup Exec 11d, 12, 12.5
- Symantec NetBackup 6.5

Agent Installation Procedures:

- [Preparing to Install a Backup Solution](#)
- [Preinstallation Tasks](#)
- [Installing the CA ARCserve Agent](#)
- [Installing the Symantec Backup Exec RALUS Agent](#)
- [Installing the Symantec NetBackup 6.5 Client](#)
- [Installing the EMC NetWorker Client](#)

Preparing to Install a Backup Solution

Before performing one of the backup solution installation procedures described here, make sure you have the following information and tools:

- **Backup and media server IP addresses** – Most backup agents need to know the IP addresses of the backup and media servers you plan to use with the SnapServer. Use the **Maintenance > Host File Editor** page in the SnapServer's Web Management Interface to supply a host-name-to-ip-address mapping that persists across system reboots.
- **SnapServer is seen by Backup software as a UNIX/Linux client** – When you configure a backup server to see the agent or client running on the SnapServer, assume the server is a UNIX or Linux client.
- **The agent/client files required by your backup software** – Typically, these files are either provided on your backup software's User CD or are available for download from the manufacturer's website. You will need to copy these files (usually delivered in a compressed format, for example, as *.rpm, *.tgz, or *.tar files) to the SnapServer.
- **A secure shell (SSH) client** – To remotely install any backup solution on the SnapServer, you must have an SSH client installed on a remote workstation. The SnapServer SSH implementation requires SSH v2. If you do not already have an SSH client application installed, you can download one from the Internet.

NOTE: The commands you must enter via SSH to install your backup agent are case-sensitive; pay careful attention to the capitalization of commands, and enter them exactly as shown.

- **Location of the SnapServer backup and restore path** – Backup servers often request the path for backup and restore operations on the SnapServer. When you configure a backup server to see the agent or client running on the SnapServer, use the following path:

```
/shares/sharename
```

where *sharename* is the name of the share to be backed up. If you have accepted the default SnapServer configuration, the correct path is as follows:

```
/shares/SHARE1
```

- **Backup user account is configured to be exempt from password policies (if applicable)** – If the backup application uses a specific local SnapServer user account to perform backups, ensure that the user is exempt from password expiration policies, if enabled (see the Online Help for procedures to set password policy for local users).

Preinstallation Tasks

Perform the following tasks prior to installing any solution:

Step 1: Identify backup and media servers to the SnapServer.

In the Web Management Interface, navigate to the **Maintenance > Host File Editor** page and click **Add**. In the page that opens, enter the IP address of the backup or media server; or enter one or both of the following as required by your backup software:

- **Host name (long form)** – Enter the fully qualified address for the backup server using the *myserver.mydomain.com* format.
- **Host name (short form)** – Enter an abbreviated address for the backup server using the *myserver* format.

Click **OK**. The entry appears on the **Host Editor** page. Repeat this procedure for each backup and media server you plan to use.

Step 2: Make sure SSH is Enabled on the SnapServer.

Navigate to the **Server > SSH** page, make sure the **Enable SSH** box is checked, and then click **OK**. SSH is immediately available.



CAUTION: To maintain security, we recommend disabling SSH when it is not in use.

Step 3: On a client computer connected to the SnapServer, create a directory called *agent*.

You must create a directory to which you will copy the agent files. Create this directory on a client computer connected to the SnapServer. For purposes of illustration, the procedures described here assume that this directory is called *agent*.

Step 4: Copy the agent/backup files to the *agent* directory.

Using a method appropriate to your environment, copy the agent/client files to the directory you just created for this purpose.

Installing the CA ARCserve Agent

This section explains how to install the CA ARCserve Agent versions 11.5 and 12.0. Consider these important points:

- This procedure assumes that you are using the default SnapServer configuration; and you have created a directory called *agent* (to which to copy your agent/client files) on the default share (SHARE1), such that the path to the directory is `/shares/SHARE1/agent`.
- Installing the ARCserve backup agent on a SnapServer requires three agent (**.rpm*) files. These agent files are available from your ARCserve CD, but some ARCserve CDs may not contain all the required files. To obtain the files you need, contact Computer Associates. If you have questions about the agent configuration, refer to your CA ARCserve documentation.

Prepare the SnapServer

1. Connect to the SnapServer via SSH.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.

2. At the prompt, log in as **admin**, using the password you created for this account during the initial setup of the server.

3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press **Enter**:

```
osshell
```

4. To change to superuser, enter the following command, and press **Enter**:

```
su -
```

5. At the prompt, enter the admin user password, and press **Enter**.

6. To change to the agent directory, type the following command and press **Enter**:

```
cd /shares/SHARE1/agent
```

7. To unpack the tar file to get the agent files, type the following command and press **Enter**:

```
tar -zxvf Linux.tar.Z
```

NOTE: If you later delete the volume this directory is on, you will need to reinstall the agent.

8. Determine which volume has the most available space by looking at the **Avail** column in the volume usage table.

```
cd /hd
```

```
ls (lists all volumes)
```

```
df -h (shows volume usage)
```

9. Change directory to the volume with the most available space.

```
cd [volumename]
```

where **[volumename]** is the volume with the most available space.

10. Create a directory **arcserve** on that volume.

11. Create the following symbolic links from the new directories in **arcserve** to the **/opt** directory:

```
ln -s /hd[volumename]/arcserve /opt/CA
```

Install CA ARCserve Agent

1. To install the agent files, enter the following command at the prompt, and press **Enter**:

```
rpm --nodeps -Uvh babagtux.rpm *lic*.rpm
```
2. Once the license is installed, run the Install script by entering the following command at the prompt and pressing **Enter**:

```
./install
```

Answer the prompts using the defaults.

NOTE: You are installing the Linux Client Agent.
3. To change to the agent directory, enter the following command, and press **Enter**:

```
cd /opt/CA/BABuagent/
```
4. To run the setup program, enter the following command, and press **Enter**:

```
./uagentsetup
```

The ARCserve agent is now installed.
5. Enter the following command to run the script that will edit the agent.cfg file:

```
fix-arcserv
```
6. Close the SSH client and return to the Web Management Interface. To start the newly installed backup agent, navigate to the **Maintenance > Shutdown/Restart** page, and click **Restart**.
7. Delete the agent files you copied to the SnapServer because they are no longer needed.
8. To verify the success of the installation, use your backup management software to configure and run a test backup.

Uninstall ARCserve Agent

1. If you still have the tar or install directory that you copied to the SnapServer when you installed the ARCserve Agent, the uninstall script will be in that directory. If you do not have the directory or tar, copy the files again from the ARCserve CD or get them from Computer Associates.
2. Make sure you have the script `uninstall`. Type the following and follow the prompts:

```
./uninstall
```

NOTE: Choose Option 1 to uninstall.
3. Uninstall the license `rpm` by typing the following:

```
rpm -e ca-lic
```
4. Verify that ARCserve Agent has been uninstalled by typing the following and verifying that you do not see the agents:

```
rpm -qa | grep BAB
```

Installing the Symantec Backup Exec RALUS Agent

To install the Backup Exec RALUS agent, follow these procedures.

Prepare the SnapServer

1. Connect to the server over SSH.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.

2. Log in as admin (using the password for the admin account).
3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press **Enter**:

```
osshell
```

4. Change to root by entering the following command:

```
su -
```

Give the root password (same as admin password).

5. Select a volume on which to put a directory called *ralus*.

NOTE: If you later delete the volume the *ralus* directory is on, you will need to reinstall the agent.

```
cd /hd
```

```
ls [lists all the volumes]
```

```
df -h [shows volume usage]
```

6. Determine which volume has the most available space by looking at the **Avail** column in the volume usage table. Change directory (**cd**) to the volume with the most available space.

```
cd [volumename]
```

where **[volumename]** is the name of the volume with the most available space.

7. Create a directory *ralus* on that volume:

```
mkdir ralus
```

8. In the *ralus* directory, create 3 directories called *VRTS*, *VRTSralus*, and *VRTSvxms*.

```
cd ralus
```

```
mkdir VRTS VRTSralus VRTSvxms
```

```
ls [to verify that the directories are there]
```

9. If CA Antivirus has been installed, you will have an */opt* directory. If it has not been installed, create an */opt* directory:

```
mkdir /opt
```

10. Create the following symbolic links from the new directories in *ralus* to the */opt* directory:

```
ln -s /hd/[volumename]/ralus/VRTS /opt
```

```
ln -s /hd/[volumename]/ralus/VRTSralus /opt
```

```
ln -s /hd/[volumename]/ralus/VRTSvxms /opt
```

where **[volumename]** is the name of the volume with the most available space.

11. Use the host file editor (**Maintenance > Host File Editor** page) to add all the Backup Exec servers to **/etc/hosts** on the SnapServer, and verify that the agent server can ping the main Backup Exec server.

NOTE: Do not edit the **/etc/hosts** file directly with a text editor.

Install Backup Exec RALUS Agent

1. From a network client, create a *ralusinstall* directory on SHARE1 of the SnapServer, then copy the RALUS agent tar file or contents of the RALUS agent CD to the directory.

- If you copied the files from the CD, proceed to Step 3. If you downloaded the files from the Symantec website, in SSH, extract the files:

```
cd /shares/SHARE1/ralusinstall
tar -zxvf [filename].tar.gz
```

where [filename] is the name of the Backup Exec tar file.

- Install the agent:

```
cd /shares/SHARE1/ralusinstall
(or other directory containing the CD contents)
./installralus
```

Follow the installation instructions, accepting the default options.

NOTE: During the installation process, you may see an error message about the failure to add root to the *beoper* group. This error will be resolved in the following step.

- Add the user *root* to the group *beoper* manually (or any other local SnapServer user you wish to use to perform backups):

```
cli group member add group-name=beoper user-name=root
```

NOTE: If using a local SnapServer user account other than *root* or *admin*, and if password policies are enabled, configure the user to be exempt from password expiration. See Chapter 7, "Security Options."

- Start the Backup Exec RALUS agent by rebooting the SnapServer either through the Web Management Interface (**Maintenance > Restart**), or by typing:

```
/etc/rc.d/init.d/VRTSralus.init start
```

- Verify that using Backup Exec, you can create a job using the UNIX agent:

- Create a GuardianOS Root login account on the Backup Exec server.
- Connect as *root* (the password will be the same as the admin account password).
- Create a job and choose the UNIX agent representing the SnapServer.
- Verify that you can connect to the agent, configure a job, and run the job.

Uninstall the Backup Exec RALUS AGENT

- To uninstall the RALUS Agent, you will need the tar or install directory that you copied to the SnapServer when you installed the Agent (follow Steps 1 through 3 of "Install Backup Exec RALUS Agent"). Make sure you see the script `uninstallralus`

- Type:

```
./uninstallralus
```

Follow the prompts.

- Verify that the Symantec RALUS agent has been uninstalled by typing the following command:

```
rpm -qa | grep VRTS
```

Installing the Symantec NetBackup 6.5 Client

NOTE: This procedure assumes that you are using the default SnapServer configuration; and you have created a directory called *agent* (to which to copy your agent/client files) on the default share (SHARE1), such that the path to the directory is `/shares/SHARE1/agent`.

To install the Symantec NetBackup 6.5 Client, follow these procedures.

Prepare the SnapServer

1. Connect to the server over SSH.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.

2. Log in as **admin** (using the password for the admin account).
3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press **Enter**:

```
osshell
```

4. Change to root by entering the following command:

```
su -
```

Give the root password (same as admin password).

5. Select a volume on which to put a directory called *openv*.

NOTE: If you later delete the volume the *openv* directory is on, you will need to reinstall the agent.

```
cd /hd
```

```
ls [lists all the volumes]
```

```
df -h [shows volume usage]
```

6. Determine which volume has the most available space by looking at the **Avail** column in the volume usage table. Change directory to the volume with the most available space.

```
cd [volumename]
```

where *[volumename]* is the name of the volume with the most available space.

```
ls [lists what is on that volume]
```

7. Create a directory called *openv* on that volume:

```
mkdir openv
```

8. Create a “symbolic” link to the *openv* directory in the */usr/* directory:

```
ln -s hd/[volumename]/openv /usr/
```

where *[volumename]* is the name of the volume with the most available space.

9. Use the host file editor (**Maintenance > Host File Editor** page) to add the NetBackup servers to */etc/hosts* on the SnapServer. Verify that you can ping the NetBackup server.

Install NetBackup v6.5 Client

1. Using a network client, copy the directory called *NBClient*s from the Client CD to a directory on a share (for example, *SHARE1* or *Agent*) on the SnapServer.

2. In SSH, install the files:

```
cd /shares/SHARE1/NBClient/catalog/anb
```

```
./client.inst
```

Follow the instructions, choosing RedHat Linux (choose 2.6 kernel version, if available) as the type.

3. Once the NetBackup Client is installed, reboot the server using the Web Management Interface (**Maintenance > Restart**) to start the client service.

4. Verify that you can configure the UNIX client:

- a. Create a policy and add the SnapServer as a client.
- b. Look at the client list to verify that the SnapServer client is listed.

Uninstall the NetBackup v6.5 Client

1. Log in to the client system as the root user.
2. Navigate to the volume where you installed the NetBackup directory.


```
cd /hd/vol_mnt[X]/
rm -rf /usr/opensv/
rmdir opensv/
```
3. Remove the NetBackup entries in the client's `/etc/services` file.

Locate the lines, marked by the following strings and delete them:

```
# NetBackup services#.....# End NetBackup services #
```
4. Remove the NetBackup services by deleting the files for `bpcd`, `vnetd`, `vopied`, and `bpjava-msvc` in the `/etc/xinetd.d/` directory.


```
rm -rf /etc/xinetd.d/bpcd
rm -rf /etc/xinetd.d/vnetd
rm -rf /etc/xinetd.d/vopied
rm -rf /etc/xinetd.d/bpjava-msc
```
5. Restart the SnapServer services by either rebooting or typing:


```
/etc/rc.d/init.d/xinetd reload
```

Installing the EMC NetWorker Client

NOTE: This procedure assumes that you are using the default SnapServer configuration; and you have created a directory called `agent` (to which to copy your agent/client files) on the default share (SHARE1), such that the path to the directory is `/shares/SHARE1/agent`.

This section describes how to install the EMC NetWorker UNIX/Linux client, as well as special procedures EMC NetWorker users must follow in order to perform backup and restore operations on the SnapServer.

Prepare the SnapServer

1. Connect to the server over SSH.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.
2. Log in as `admin` (using the password for the admin account).
3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press **Enter**:


```
osshell
```
4. Change to root by entering the following command:


```
su -
```

Give the root password (same as admin password).
5. Select a volume on which to put a directory called `networker`.

NOTE: If you later delete the volume the *networker* directory is on, you will need to reinstall the agent.

```
cd /hd
ls [lists all the volumes]
df -h [shows volume usage]
```

- Determine which volume has the most available space by looking at the **Avail** column in the volume usage table. Change directory to the volume with the most available space.

```
cd [volumename]
where [volumename] is the name of the volume with the most available space.
```

- Create a directory called *networker* on that volume:

```
mkdir networker
```

- In the *networker* directory, create the following directories called *opt*, *usr*, and *opt/usr*.

```
cd networker
mkdir opt usr opt/usr
ls [to verify that the directories are there]
```

- If CA Antivirus has been installed, you will have an */opt* directory. If it has not been installed, create an */opt* directory:

```
mkdir /opt
```

- Create links from the *networker* working volume to the root filesystem:

```
ln -s /hd/vol_mnt[X]/networker/nsr/
ln -s /hd/vol_mnt[X]/networker/opt/nsr /opt/
ln -s /hd/vol_mnt[X]/networker/usr /usr/
```

where *vol_mnt [X]* is the NetWorker installation target volume.

- Modify the SnapServer environment by editing */etc/profile* as follows:

```
cp /etc/profile /etc/profile.nwbk
echo
PATH=$PATH:/hd/vol_mnt[X]/networker/usr/bin:/hd/vol_mnt{X}/networker/u
sr/sbin:/hd/vol_mnt[X]/networker/usr/lib >> /etc/profile
```

where *vol_mnt [X]* is the NetWorker installation target volume.

NOTE: Be sure to enter `>>` in the command rather than `>` or you will overwrite the file rather than append to the */etc/profile* script. If you need to redo Step 11, copy the backup to the original using the command `cp /etc/profile.nwbk /etc/profile` and then edit the file again.

- To implement the changes, enter the following command:

```
source /etc/profile
```

Install the EMC Networker Client

- Connect to the SnapServer via SSH, and log in as *admin*, using your *admin* user password.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.

2. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press **Enter**:

```
osshell
```
3. To change to superuser, enter the following command, and press **Enter**:

```
su -
```
4. At the prompt, enter the admin user password, and press **Enter**.
5. Use the `cd` command to change to the directory in the share, for example:

```
cd /shares/SHARE1/agent
```
6. To unpack the client files, enter the following commands:

```
tar xvfz nw_linux86.tar.gz
```
7. To install the NetWorker Agent rpm, enter the following command:

```
rpm -Uvh --nodeps --relocate=/usr/=/hd/vol_mnt[X]/NetWorker/usr/  
lgtoclnr-x.x-x.i686.rpm
```

where `vol_mnt [X]` is the NetWorker installation target volume and `x.x-x` is the version number.
8. To start the EMC NetWorker daemon, enter the following command at the console:

```
/etc/rc.d/init.d/networker start
```

The NetWorker client is now installed.
9. Close the SSH client, return to the Web Management Interface. To start the newly installed backup agent, navigate to the **Maintenance > Shutdown/Restart** page, and click **Restart**.
10. Delete the client files you copied to the SnapServer because they are no longer needed.
11. To verify the success of the installation, use your backup management software to configure and run a test backup.

Backup and Restore Operations with the EMC NetWorker Client

This section describes special procedures EMC NetWorker users must use in order to perform backup and restore operations on the SnapServer.

Add the SnapServer as a Root User

For backup operations, NetWorker requires that the SnapServer be configured as a root user. To add the SnapServer root user as one of the administrators, use the following procedure:

1. Open the NetWorker Administrator application.
2. Click the **Configuration** tab.
3. Click the **User Groups** menu item.
4. Click the **Administrators** group.
5. In the Configuration box, add one of the following:

```
user=root@hostname
```

where `hostname` is the host name of the SnapServer for each SnapServer.

Or, enter:

```
user=root
```

to add root for all SnapServers.
6. Click **OK**.

Recover and Retrieve Operations

The EMC NetWorker administrative interface does not support data recovery operations from a remote client for a Linux-based operating system such as GuardianOS. To recover data, you must execute one of the following CLI commands from a SSH client.

- **Recover** – The `recover` command restores data from a normal backup job.
- **Nsrretrieve** – The `retrieve` command restores data from an archive.

Use either the `recover` or the `retrieve` command exactly as described below. For more details on these commands, see the *EMC NetWorker Command Reference*.

1. Connect to the SnapServer via SSH, and log in using the admin user name and password.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.

2. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press **Enter**:

```
osshell
```

3. To change to superuser, enter the following command, and press **Enter**:

```
su -
```

4. At the prompt, enter the admin user password, and press **Enter**.

5. **To recover data from a normal backup operation**, enter one of the following commands, and press **Enter**:

- To recover data to its original location:

```
recover -s backupservername -c SnapServername -f -i  
"/shares/SHARE1/data/" -a  
where /shares/SHARE1/data is the path of the data you are restoring.
```

- To recover data to a different location:

```
recover -s backupservername -c SnapServername -f -i -a R -d  
"/shares/SHARE1/relocated_data/" "/shares/SHARE1/Data/"  
where /shares/SHARE1/relocated_data/ is the path to the new target location  
for the restore operation; and where /shares/SHARE1/Data/ is the path of the  
data you are restoring.
```

6. **To retrieve data from an archival backup operation**, enter one of the following commands, and press **Enter**:

- To retrieve data to its original location:

```
nsrretrieve -f -i -s backupservername -A annotation  
"/shares/SHARE1/data/"  
where /shares/SHARE1/data/ is the path of the data you are restoring.
```

- To retrieve data to different location:

```
nsrretrieve -f -iR -d "/shares/SHARE1/new_dir" -s backupservername -  
A "annotation" "/shares/SHARE1/Data/"  
where /shares/SHARE1/new_dir" is the path to the new target location for the  
restore operation; where annotation is the name of the EMC NetWorker backup;  
and /shares/SHARE1/Data/" is the path of the data you are restoring.
```

iSCSI Disk Backups

iSCSI disks can be backed up from iSCSI clients using any standard backup application on the client operating system. These backups run independently of the SnapServer since the client backs up the contents of the iSCSI disk as if the iSCSI disk were a local hard disk.

Windows clients can make backups of VSS-based snapshots of iSCSI disks using VSS-compatible backup applications. See “iSCSI Disks” on page 6-9 for instructions.

Using Backup Exec for VSS-based Snapshots of SnapServer iSCSI Disks

To configure Backup Exec to take native VSS snapshots of SnapServer iSCSI disks using Backup Exec's *Advanced Open File* or *Advanced Disk-Based Backup* feature, you must first add a Windows registry entry to the systems running the Backup Exec Server and all of the Backup Exec agents backing up iSCSI disks.

After the Backup Exec Server or agent has been installed, modify the registry to add the SnapServer as a Backup Exec VSS provider:

1. Run the following command:

```
regedit
```

2. Navigate to the following key:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Symantec\Backup Exec For Windows\Backup Exec\Engine\Misc\VSSProviders]
```

3. Underneath VSSProviders are other keys numbered sequentially from 0 to some number. Create a new key in VSSProviders named after the highest key value plus 1 (such as, if the highest key value is 9, create a new key value 10). For example:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Symantec\Backup Exec For Windows\Backup Exec\Engine\Misc\VSSProviders]\10
```

4. Inside the new key, create three string values:

VALUE NAME	VALUE DATA
ID	{759c7754-6994-46c9-9cf9-c34ac63a0689}
Name	SnapServer VSS Hardware Provider
Version	5.2

5. Close `regedit`.

The SnapServer VSS Provider should now be available to Backup Exec to use for VSS-based backups.

This appendix provides additional information and configuration options about accessing shares and files on the SnapServer. GuardianOS on the SnapServer supports share-level as well as file- and directory-level permissions for all local and Windows domain users and groups (see [“Windows ACLs” on page C-6](#)).

File and directory security can be configured using either Windows NTFS-style security or classic Unix-style security. The type of security present on a file or directory is its “security personality” (see [“File-level Security” on page C-5](#)).

Files and directories are stored on the server on volumes (or directories underneath) with a configured “security model.” The security model on the volume governs the permitted security personalities, the default personalities, and the ability to change personalities on child files and directories.

Security models can be configured on volumes in either DynamicRAID and Traditional RAID mode. For DynamicRAID, the security models can only be on the top-level volumes. With Traditional RAID, the directories immediately underneath the top-level volume directory can also be configured with a security model and are known as “security model directories.” See [“Security Model Directories” on page C-2](#).

The default security model on newly-created volumes is always Windows/Mixed. It can be changed to a Unix security model if necessary.

Topics in Shares and File Access:

- [Security Model Rules](#)
- [Security Model Directories](#)
- [Security Model Management](#)
- [File and Share Access](#)

Security Model Rules

Files and directories created by clients inside security models will acquire security personality and permissions according to the rules of the chosen security model.

Windows/Mixed Security Model:

- Files and directories created by SMB clients will have the Windows security personality. Permissions will either be inherited according to the ACL of the parent directory (if Windows) or will receive a default ACL that grants the user full access only (if the parent is UNIX or has no inheritable permissions).
- Files and directories created by non-SMB clients will have the UNIX personality. UNIX permissions will be as set by the client (per the user’s local umask on the client).

- The security personality of a file or directory can be changed by any user with sufficient rights to change permissions or ownership. If a client of one security personality changes permissions or ownership of a file or directory of a different personality, the personality will change to match the personality of the client protocol (for example, if an NFS client changes UNIX permissions on a Windows file, the file will change to the UNIX personality).

UNIX Security Model:

- Files and directories created by non-SMB clients will have the UNIX personality. UNIX permissions will be as set by the client (per the user's local umask on the client).
- Files and directories created by SMB clients will have the UNIX personality. UNIX permissions will be set to a default.
- The personality of files and directories cannot be changed on a UNIX security model. All files and directories always have the UNIX personality.

Security Model Directories

With Traditional RAID, a security model can be configured on directories immediately underneath the top-level volume directory.

Default ownership differs according to the method used to create the security model directory:

- **From the client** – For UNIX personality directories, the owner and owning group will be according to the logged-in user. For Windows personality directories, the owner will be the logged-in user, or “Administrators” for directories created by Domain Admins or members of the local admingrp.
- **From the Web Management Interface** – For UNIX personality directories, the user and group owner will be admin and admingrp. For Windows personality directories, the owner will be the local admingrp (“Administrators”).

Security models and permissions differ according to the method used to create the security model directory:

- **From the client:** If SMB, permissions will either be according to ACL inheritance (if the parent volume root directory has the Windows security model) or *Full Access* to the owning user only. Permissions for directories created by all other protocols will be set by the client (per the client's umask).
- **From the Web Management Interface:**
 - If created in a UNIX volume, permissions are *777* (rwxrwxrwx).
 - If created in a Windows/Mixed volume, permissions allows all users to create, delete, and change permissions on files created inside the security model, and grants full control to administrators.

Security Model Management

Changes to a security model can optionally be propagated with the corresponding personality and default permission to all files and directories underneath the security model.

When setting the security model:

- For Traditional RAID, which permits security models to be set on both volumes and directories immediately underneath volumes, you can mix security models on the volume.
- For DynamicRAID, only a single security model can be set on the entire volume at the root level but not the directories immediately underneath the volume.

When changing the security model:

- If changing from Windows to UNIX, all files and directories will be changed to be owned by *admin* and *admingrp*, with UNIX permissions of 777(rwxrwxrwx).
- If changing from UNIX to Windows, files and directories will be changed to default permissions that allow all users the ability to create and manage their own files and directories and to access other users' files and directories.

Special Share Options

The basic setup and configuration of shares on a SnapServer is handled on the Security > Shares page (see [Chapter 7, "Security Options"](#)). This section covers more details about the special options and features of share security on your SnapServer:

- [Hiding Shares](#)
- [Share Level Permissions](#)
- [Where to Place Shares](#)

Hiding Shares

There are three ways a share can be hidden in GuardianOS:

- Name the share with a dollar-sign (\$) at the end. This is the traditional Windows method of hiding shares; however, it does not truly hide the share since Windows clients themselves filter the shares from share lists. Other protocols can still see dollar-sign shares.
- Hide the share from all protocols (except NFS) by navigating to **Security > Shares > Create Share > Advanced Share Properties** and selecting the **Hide this Share** checkbox, or by selecting a share, clicking to expand **Advanced Share Properties**, and selecting the **Hide this Share** checkbox. When a share is hidden this way, the share is invisible to clients, and must be explicitly specified to gain access.

NOTE: Hidden shares are not hidden from NFS, which cannot access invisible shares. To hide shares from NFS, consider disabling NFS access to the hidden shares.

- Disable individual protocol access to certain shares by navigating to **Security > Shares > Create Share > Advanced Share Properties** and enabling/disabling specific protocols, or by selecting a share, clicking to expand **Advanced Share Properties**, and enabling or disabling specific protocols.

Share Level Permissions

Share-level permissions on GuardianOS are applied cumulatively. For example, if the user "j_doe" has Read-Only share access and belongs to the group "sales", which has Read/Write share access, the result is that the user "j_doe" will have Read/Write share access.

NOTE: Share-level permissions only apply to non-NFS protocols. NFS access is configured independently by navigating to the Security > Shares page, selecting from the table the NFS Access level for the share, and modifying the client access as desired. See “NFS Share Access”.

Where to Place Shares

For security and backup purposes, it is recommended that administrators restrict access to shares at the root of a volume to administrators only. After initialization, all SnapServers have a default share named *SHARE1* that points to the root of the default volume *VOL0* (Traditional RAID) or *Volume1* (DynamicRAID). The share to the root of the volume should only be used by administrators as a “door” into the rest of the directory structure so that, in the event that permissions on a child directory are inadvertently altered to disallow administrative access, access from the root share is not affected. This also allows one root share to be targeted when performing backups of the server. If it is necessary to have the root of the volume accessible, using the Hidden option helps ensure only those that need access to that share can access it.

File and Share Access

The shares feature also controls access by other users and groups. This section provides information on setting up the shares options to allow proper access to the files.

NFS Share Access

When controlling share access for NFS clients, administrators limit client access to the shares independently of share level permissions that apply to other protocols. Access is controlled on a per-share basis. To set the NFS access, navigate to **Storage > Shares**. In the Shares table, click in the **NFS Access** column of the share you want to modify. Changes made on this page affect the NFS “exports” file within GuardianOS.



CAUTION: If there are multiple shares to the same directory on the disk, and those shares permit access via NFS, they must all have the same NFS export configuration. This is enforced when configuring NFS access to the overlapping shares.

Snapshot Access

Snapshots are accessed via a snapshot share. Just as a share provides access to a portion of a live volume (or filesystem), a snapshot share provides access to the same portion of the filesystem on all current snapshots of the volume. The snapshot share’s path into snapshots mimics the original share’s path into the live volume.

Snapshot Shares and On Demand File Recovery

A *snapshot share* is a read-only copy of a live share that provides users with direct access to versions of their files archived locally on the SnapServer via a snapshot. Users who wish to view or recover an earlier version of a file can retrieve it on demand without administrator intervention.

Snapshot shares are created during the course of creating a share, or thereafter by navigating to the Snapshots page and clicking the name of a snapshot. For instructions on accessing snapshot shares, see “[Snapshot Access](#)” on page C-4.

Creating a Snapshot Share

You create a snapshot share by selecting the **Create Snapshot Share** option on the **Security > Shares > (share_name) > Share Properties** page, under the **Advanced Share Properties** link.

For example, assume you create a share to a directory called *sales*, and you select the **Create Snapshot Share** option. When you connect to the server via a file browser or use the **Misc. > Web Home** link in the Site Map, two shares display:

```
SALES
SALES_SNAP
```

The first share provides access to the live volume, and the second share provides access to any archived snapshots. Other than read-write settings (snapshots are read-only), a snapshot share inherits access privileges from its associated live-volume share.

NOTE: The same share folders appear on the Web View page when you connect to the SnapServer using a Web browser. However, the snapshot share folder does not provide access to the snapshot; it always appear to be empty. You can prevent the snapshot share from displaying on this Web View page by selecting the Hide Snapshot Share option when creating or editing a share.

Accessing Snapshots Within the Snapshot Share

A snapshot share contains a series of directories. Each directory inside the snapshot share represents a different snapshot. The directory names reflect the date and time the snapshot was created.

For example, assume the snapshot share named *Sales_SNAP* contains the following four directories:

```
latest
2011-09-25.120000
2011-10-01.000100
2011-10-07.020200
```

The *latest* directory always points to the most recent snapshot (in this case, **2011-10-07.020200**, or October 7th, 2011, at 2:02 a.m.). A user may view an individual file as it existed at a previous point in time or even roll back to a previous version of the file by creating a file copy to the current live volume.

NOTE: The latest subdirectory is very useful for setting up backup jobs, as the name of the directory is always the same and always points to the latest available snapshot.

File-level Security

GuardianOS supports two “personalities” of filesystem security on files and directories:

- **Windows ACLs:** Windows NTFS-style filesystem permissions. Windows ACLs fully support the semantics of NTFS ACLs, including configuration, enforcement, and inheritance models (not including the behavior of some built-in Windows users and groups).
- **UNIX:** Traditional UNIX permissions (rwx) for owner, group owner, and other.

By default, volumes are created with the Windows/Mixed security model (Windows-style ACLs for files created by SMB clients and UNIX-style permissions for files created by other protocols and processes), and allow all users to create, delete, and configure permissions on their own files and to access files and directories created by other users.

Security Personalities and Security Models

The security personality of a file or directory is dependent on the security model of the root directory or volume in which the file or directory exists.

Files and directories in a Windows/Mixed security model can have either a Windows or UNIX security personality, depending on the network protocol used to create the file or change permissions on it. Files in a UNIX security model always have the UNIX security personality and can only be set by NFS clients.

Windows ACLs

GuardianOS fully supports Windows NTFS-style filesystem ACLs, including configuration, enforcement, and inheritance models. Inside Windows/Mixed root directories, files created and managed by Windows clients have the Windows security personality and behave just as they would on a Windows server. Clients can use the standard Windows 2000, 2003, XP, Vista, or Windows 7 interface to set directory and file permissions for local and Windows domain users and groups on the SnapServer.

Permissions are enforced for the specified users in the same manner for all client protocols, including non-SMB clients that normally have the UNIX security personality. However, if a non-SMB client changes permissions or ownership on a Windows personality file or directory (or deletes and recreates it), the personality will change to UNIX with the UNIX permissions specified by the client.

NOTE: Group membership of NFS clients is established by configuring the local client's user account or the NIS domain. Group membership of SnapServer local users or users ID-mapped to domain users is not observed by NFS clients. Therefore, ACL permissions applied to groups may not apply as expected to NFS clients.

Default File and Folder Permissions

When a file or directory is created by an SMB client, the owner of the file will be the user who created the file (except for files created by local or domain administrators, in which case the owner will be the **Administrators** group, mapped to the local **admingrp**), and the ACL will be inherited per the inheritance ACEs on the parent directory's ACL. The owner of a file or directory always implicitly has the ability to change permissions, regardless of the permissions established in the ACL. In addition, members of the SnapServer's local admin group, as well as members of Domain Admins (if the server is configured to belong to a domain) always implicitly have *take ownership* and *change ownership* permissions.

Setting File and Directory Access Permissions and Inheritance (Windows)

Access permissions for files and directories with the Windows security personality are set using standard Windows 2003, XP, Vista, 2008, or 7 security tools. GuardianOS supports:

- All standard generic and advanced access permissions that can be assigned by Windows clients.
- All levels of inheritance that can be assigned to an ACE in a directory ACL from a Windows client.
- Automatic inheritance from parent directories, as well as the ability to disable automatic inheritance from parents.
- Special assignment and inheritance of the CREATOR OWNER, CREATOR GROUP, Users, Authenticated Users, and Administrators built-in users and groups.

Procedure to set file and directory access permissions and inheritance in Windows:

1. Using a Windows 2003, XP, Vista, 2008, or 7 client, **map a drive** to the SnapServer, logging in as a user with change permissions for the target file or directory.
2. Right-click the file or directory, choose **Properties**, and then select the **Security** tab.
3. Use the **Windows security tools** to add or delete users and groups, to modify their permissions, and to set inheritance rules.

Troubleshooting SnapServers

Basic techniques for identifying and resolving common hardware and networking issues are described here.

Topics in Troubleshooting SnapServers

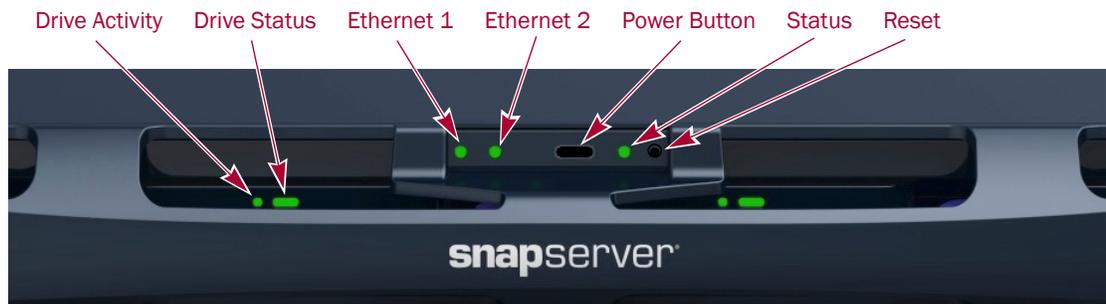
- [LED Indicators](#)
- [System Reset Options](#)
- [Maintenance Mode](#)
- [Networking Issues](#)
- [Miscellaneous Issues](#)
- [Phone Home Support](#)

LED Indicators

LED indicators provide information on the status of basic connectivity, disk drives, fan modules, and power supply modules.

SnapServer DX1

The SnapServer DX1 has two network LEDs (Ethernet 1, left; Ethernet 2, right) and one status LED. Each drive has two disk LEDs (Drive Activity and Drive Status). There is a Power button and a Reset button, as shown in the following illustration:



SnapServer DX2

The SnapServer DX2 has one network LED (Ethernet) and one system status LED on the Power Panel located on the right flange, along with a Power button and a Reset button. Each drive has two disk LEDs (Drive Activity and Drive Status) as shown in the following illustration:



SnapExpansion

The SnapExpansion has one network LED (Ethernet) and one system status LED on the Power Panel located on the right flange, along with a Power button. Each drive has two disk LEDs (Drive Activity and Drive Status) as shown in the following illustration:



The Power button is an override option only as the SnapServer head unit manages the power to all attached expansion units. To use the Power button to shut down the expansion unit, press and hold it for four seconds.

Drive Status LEDs

The following LED behavior can be observed when the SnapServer is ON.

Device State	LED State
No Disk Drive in Bay	Off
Normal Operation	Solid green
Unit Identification Indicator	Flashing amber *
RAID in Degraded Mode/Rebuilding [†]	Flashing green/amber
Failed	Flashing red
Unusable	Solid amber

* When the Unit Identification Indicator is turned on, all drive LEDs and the system/status LED flash amber.

† All member drives of a RAID flash green/amber when the RAID is degraded and throughout the entire rebuild process when the RAID is being repaired with a new drive. Do not remove the drive during this time.

Drive Activity LEDs

Device State	LED State
Powered OFF / No Activity	Off
Drive Activity	Flashing green

Network LEDs

Device State	LED State
Powered OFF	Off
Link Up (SnapServer Powered ON)	Solid green
Link Down	Off

Server Status LEDs

Device State	LED State
Powered OFF	Off
Unit Identification Indicator	Flashing amber *
Booting	Solid amber
Normal Operation	Solid green
Shutting down	Flashing green
Maintenance Mode	Flashing green/amber

* When the Unit Identification Indicator is turned on, all drive LEDs and the system/status LED flash amber.

Power Supply Status LED

Device State	LED State
Normal Operation	Solid green
Standby	Solid red
Power Failure	Off
Fan Failure	Blinking red

System Reset Options

Often the first thing to try in resolving anomalous behavior on a SnapServer is to reset the server to factory defaults using the Web Management Interface. See [“Factory Defaults”](#).

Performing System Resets Without Network Access

Should Web Management Interface access to the server be lost, the **Reset** button can be used to reset server settings and reestablish connectivity. This option is also useful should you forget the admin password.

On the SnapServer DX1, the **Reset** button is accessed via a small hole next to the **Power** button on the front of the server at the top. On the DX2, the hole is located at the bottom of the Power panel on the right flange. Verify that the server is fully booted (as indicated by the system/status LED). Using the end of a straightened paper clip or the fine point of an instrument, press in and hold the **Reset** button for a few seconds.

The system will reboot after about a minute. As a part of the reset and reboot process, the SnapServer does the following:

- Clears user-defined settings such as DHCP configuration.
- Resets the server name to its default setting (**SNAP<server_number>**).
- Resets network speed and bonding settings to their defaults.
- Resets the Administrator password to the default (**admin**).
- Resets the web server to allow http access.

Maintenance Mode

The SnapServer may enter Maintenance Mode (system/status LED blinking amber and green) when GuardianOS has been compromised and is in need of repair or reinstallation. The two functions available in Maintenance Mode should only be used under the direction of Overland Technical Support:

- **Repair** – Reapplies the GuardianOSImage, but preserves system settings.
- **Fresh install** – Reinstalls GuardianOS, overwriting any previous configurations and destroying all disk partitions.



CAUTION: Because of significant changes introduced in GuardianOS 7.2, a fresh install of GuardianOS 7.2 should not be performed on a SnapServer running an older version of GuardianOS. Failure to follow this guideline can result in total failure of the SnapServer to start, even into Maintenance Mode. The Fresh install option should only be performed with the same version of GuardianOS currently installed on the SnapServer, and only under the direction of Overland Technical Support.

NOTE: To install GuardianOS, you must obtain the appropriate GuardianOS image file. This file is available for download by entitled users from the SnapServer support site:
<http://docs.overlandstorage.com/SnapServer>

Networking Issues

These are some of the networking issues you may encounter when using your SnapServer.

The Server Cannot Be Accessed over the Network

Inaccessibility may be caused by a number of reasons. To resolve this issue, use one of the following methods:

- Verify that you have the correct IP address of the server, and try to connect again.
- Verify that the LED for the primary Ethernet port is lit. (This light indicates network connectivity.) If the light is not lit, perform the following:
 - The most likely cause is the physical connection. Check for a loose or damaged cable, or poor connections in the port connector.

- This problem may also be caused by a mismatch between the settings on the switch or hub and the settings on the SnapServer Ethernet port. These settings must match. To resolve the problem, make sure the port settings on the hub or switch match the settings for the primary port as configured on the **Network > TCP/IP** page of the Web Management Interface. Use the autonegotiate setting on both the switch and the server port.

You Have No Access to the SnapServer via HTTP

When trying to access the SnapServer via HTTP, the Web browser times out. The server can be accessed using the ping command or Windows Explorer.

- HTTP and HTTPS are both enabled by default on SnapServers. Try typing HTTPS in the Web address rather than HTTP. If you are able to access the server via HTTPS, you can re-enable HTTP on the **Network > Web** page.
- If you cannot access the server via HTTPS, try resetting the server as described on [“Performing System Resets Without Network Access” on page D-3](#).

An Access Denied Message Appears after Configuring Microsoft Domain Security

Customers who have configured local users and local groups with the same name as their domain users and groups can have security conflicts if they integrate with Microsoft Domain Security. The SnapServer will authenticate the users as local SnapServer users before authenticating through the Domain. However, the Domain users/groups may be the ones that had been granted access to the shares.

Be careful not to add local users or groups that are duplicates of those that are found on the Windows domain controller.

The SnapServer Does Not Operate Properly on a Network Running Gigabit-Full-Duplex

For Gigabit Ethernet to operate properly, both the switch and the SnapServer's primary Ethernet port must be set to *Auto* (autonegotiate). Any other setting will result in unexpected behavior and reduced performance.

The Network Does Not Have a DHCP Server and the SnapServer IP Address Is Unknown

Install SnapServer Manager (available from the SnapServer [support page](#) on the Overland Storage website) onto a client workstation on the same subnet as the SnapServer. You can then use the utility to discover all SnapServers on that network segment, and to assign static IP addresses as necessary.

Problems Occur with Domain Controller Authentication

You are receiving the following errors in your error log:

```
SMB: Domain Controller unavailable
SMB: Username not connected to Domain Controller
```

This means that either your Domain Controller is down, or the SnapServer is unable to reach it. Because it cannot communicate with the Domain Controller, it is not able to authenticate the user. Check to make sure the Domain Controller is online, is consistently reachable via the network, and that users can authenticate to the Domain Controller.

You Start Your SnapServer but Cannot See It on the Network

Ensure that the Ethernet cable is connected securely to both the network port and the server's primary Ethernet port. Also, check to see that the Link light on the front of the SnapServer is lit (solid green). If the Link light is off, this is normally caused by a mismatch between the switch/hub and the Ethernet port on the SnapServer. To resolve this problem, verify that all settings (if using multiple Ethernet ports) on the switch/hub match the setting on the server. When the server is shipped from the factory, both ports are set to autonegotiate. Therefore, the switch/hub *must* be set to autonegotiate to initially connect to the server.

SnapServers are configured by default to acquire an IP address from a DHCP server. If no DHCP server is found on the network, the SnapServer defaults to an IP address in the range of 169.254.xxx.xxx and is labeled ZeroConf in SSM. While you may not be able to see the server on your network, you can discover the SnapServer using either the default server name or the SSM utility (available at our external download site:

<http://www.overlandstorage.com/SSM>). Use the server name method if you are installing one SnapServer on the network. Use SSM if you are installing two or more SnapServers, or if your network does not have IP-to-name resolution services.

You Try to Mount to a Share on Your SnapServer from Your Linux Workstation and You Receive an RPC Timeout Message

Check the firewall configuration to your Linux workstation. Be sure you have not blocked the ability to receive TCP or User Datagram Protocol (UDP) communications. If problems persist, contact Overland Storage Technical Support.

You Receive an Access Denied Message When Attempting to Mount a Share on Your SnapServer from a Linux Workstation

If you are logged in as *root* on your workstation and NFS is enabled on your SnapServer, this message can be misleading, causing you to look for security issues, when in fact it could be a command syntax issue. For example, the common Linux mount command:

```
mount 192.168.32.124:SHARE1 /mnt
```

is missing a forward slash (/) in the command, which will return an Access Denied message. The correct syntax should be the following:

```
mount 192.168.32.124:/SHARE1 /mnt
```

NOTE: The share name is case-sensitive.

You Cannot Log in as Root to the SnapServer

GuardianOS allows you to log in as root over SMB. If this operation has failed or you have trouble logging in, be sure that you have enabled root login in the **Network > Windows/SMB** page. Also note that the root account password is tied to the admin account password. If you cannot log in as root, change the password for the admin account on the **Network > Windows/SMB** page. Use the admin password to log in as root.

You Are Unable to See Your Domain Users When Trying to Set Up Windows Security Permissions on File Folders

The SnapServer (GuardianOS) has joined the Active Directory domain properly, and you can see the domain users when you set Share permissions from the browser-based Web Management Interface.

Make sure the Windows client (PC) you are trying to set permissions from is assigned a valid DNS server. You can check your Windows client using the `ipconfig` command from a command prompt.

Miscellaneous Issues

These are some miscellaneous issues you may encounter when using your SnapServer.

Backup Applications:

You Backed Up Your Snapshot Share, Are Now Attempting to Restore It, and the Operation Fails

A snapshot share is read-only. You can restore the data to a read-write accessible share.

Other Issues:

A Problem Occurred While Booting. The System is Offline and the System/Status LED is Blinking Amber and Green

The SnapServer has booted into Maintenance (Recovery) Mode. This may be due to a boot failure in the previous boot attempt. Try booting again. If the server still returns to Maintenance Mode, call Technical Support.

Power to the SnapServer Is Unexpectedly Cut Off Due to a Power Outage

Overland Storage recommends that you use an uninterruptible power supply (UPS) with the SnapServer. If you did not have a UPS attached to the server at the time of the power outage, do the following:

1. Remove the power cables.
2. Once the power is restored and stabilized, turn the power supplies back on and reboot the server.

Once the SnapServer boots, it begins resynchronizing the RAIDs if necessary. You can use the server during the resynchronization, but performance will be a little slower than normal. Do not remove drives, however, while the server is resynchronizing the RAID.

The Server Is Not Responding to File Requests or Configuration Commands

Call your SnapServer technical support representative.

You Have Problems Seeing the Tape Library Tape Device, Not the Robotic Arm

When you have problems seeing the actual tape device rather than the robotic arm, it is most likely due to the Tape Loader being configured for Sequential Access. Change the Tape Loader to Random or Mixed Mode.

The Admin Password to the Web Management Interface Is Not Available

You can perform a limited reset to defaults, which includes the admin password, then use the Web Management Interface to set a new password. See [“Performing System Resets Without Network Access”](#).

You Cannot Delete Files or Folders From an iSCSI Disk

If an iSCSI disk is mounted to a folder, not a letter drive, in Windows you will not be able to delete files and folders inside that mount point. The Windows Recycle Bin does not understand mount points, so to avoid this problem either mount iSCSI disks to letter drives on your Windows OS, or hold down the shift key while deleting folders or files.

Phone Home Support

Once your SnapServer has been registered, Phone Home Support becomes available for use. Phone Home Support emails system logs and files that contain information useful for troubleshooting purposes to Overland Storage technical support. You can use the **Maintenance > Support** page to open a new case with technical support; or, in the course of working to resolve an issue, a technical support representative may ask you to fill out and submit this page. If a case is already in progress, you will need to enter the case number provided by the technical support representative.

NOTE: Phone Home Support interacts with two fields on the **Maintenance > Tools > Email Notification** page: (1) To use Phone Home Support, you must enter a valid SMTP server IP address on the **Email Notification** page; and (2) the first email address listed in the **Recipients** field populates the **Reply-to Address** field on the **Support** page.

Complete the following fields as appropriate, then click **OK**:

Text Field	Description
Subject	(Required) Enter a concise description that identifies the issue.
Case	(Required) Select <i>New Case</i> if you are emailing technical support for the first time. Select <i>Existing Case</i> if you have previously contacted technical support concerning the issue.
Case Number	If you selected <i>Existing Case</i> above, enter the case number provided by technical support.
Reply-to Address	(Required) This field defaults to the first email address entered as a recipient on the Server > Email Notification page. If necessary, enter at least one email address that will serve as the contact email address for this issue. To receive a copy of the email and system information attachment, check the <i>Cc Reply-to Address</i> box.
Comments	(Required) Enter additional information that will assist in the resolution of the problem.

Command Line Interface

GuardianOS includes a command line interface (SnapCLI) accessible through SSH. Using the CLI, users can access information about most of the SnapServer configuration parameters and perform configuration and maintenance functions without using the GuardianOS Web Management Interface or SSM.

NOTE: Some administrative tasks must still be performed using the Web Management Interface. The CLI is intended as a convenient way to perform some functions; it is not intended as an alternative to using the Web Management Interface.

Before You Begin

Before the storage type is configured to DynamicRAID or Traditional RAID, SnapCLI disables all standard commands and makes only the `system` command available. This command is available *only* before storage is configured, and has the following arguments:

Command	Arguments and Options	Descriptions
system	type	<code>type=DynamicRAID</code> <code>type=Traditional-RAID</code>
	force	<code>yes</code>
		Specify DynamicRAID mode Specify Traditional RAID mode Bypass confirmation prompt

Thus, the following command string:

```
system type=Traditional-RAID force=yes
```

sets the storage type to Traditional RAID and bypasses the confirmation prompt.

Once the `system` command is run and the storage type is chosen, SnapCLI unlocks the rest of the standard commands. A reboot is required if Traditional RAID is chosen as the storage type.

Topics in Command Line Interface

- [SnapCLI Syntax](#)
- [SnapCLI Commands](#)
- [Scripts in SnapCLI](#)

SnapCLI Syntax

SnapCLI command syntax uses three parameters: **COMMANDS**, **ARGUMENTS**, and **OPTIONS**. To generate commands in SnapCLI, use the following syntax:

```
COMMAND [ARGUMENT] [OPTIONS]
```

where **COMMAND** is the name of one of the SnapCLI commands, **ARGUMENT** is an action available for that command, and **OPTIONS** are additional parameters for the command.

Once logged into the CLI, there are several ways of displaying information about available parameters.

Type	To
?	see an overview of the CLI, with a list of available commands and a description of command syntax.
<i>{command}</i> help	see a description of that particular command's function and a list of options available for the command.
tab	finish the command you have started to type (such as, tab-complete).
<i>{command}</i> tab	list any arguments and/or options available for that command.

For example, to see a list of available commands once you have logged into SnapCLI, type “?” at the prompt.

To see a description of a specific command, type the command name (for example, date) + “help” or “?”:

Command	Arguments and Options	Descriptions
date	timezones	- list available time zones
	get	- get server date/time
	set [OPTIONS]	- set server date/time
	- [day=1-31]	- day of month
	- [month=1-12]	- month of year
	- [year=1900-current]	- year
	- [hour=0-23]	- hour
	- [minute=0-59]	- minutes
	- [second=0-59]	- seconds
	- [timezone=1- 40]	- timezone (use the command date timezones to get a list of timezones)

In this instance, to set the date to October 27, 2011, enter:

```
date set day=27 month=10 year=2011
```

NOTE: If, instead of typing the word date, you had typed `d + [tab]`, the word would have been completed for you. If you entered `d + [tab] + [tab]`, the word would have been completed and the available options displayed.

Suppose, instead of **date**, you typed the command **web**. Two arguments would be available, one with options:

Command	Arguments and Options	Descriptions
web	get	- get WEB properties
	set [OPTIONS]	- set WEB properties
	- require-webview-auth= (yes no)	- require HTTP/HTTPS clients to authenticate in order to access the server
	- non-secure-http= (yes no)	- enable/disable non-secure HTTP access

Thus, the following command string:

```
web set require-webview-auth=yes non-secure-http=no
```

sets HTTP/HTTPS properties on the SnapServer to require clients to authenticate in order to access the server and to disable non-secure HTTP access.

SnapCLI Procedures

Use these procedures to access and exit SnapCLI.

Logging into SnapCLI

1. Make sure your client has an SSH v2 client application installed.

NOTE: Free or low-cost SSH applications are available from the Internet.

2. Connect to the server using its name or IP address, and log in as *admin* (or any other member of *admingrp*).

You will automatically be placed in the CLI shell.

NOTE: SSH v2 is required. If you fail to connect to the server, ensure that your SSH client is configured to connect via SSH v2.

Exiting SnapCLI

To exit SnapCLI, type `exit`. The SSH session will close.

SnapCLI Commands

The following table presents a list of the available SnapCLI commands and a brief description of the function of each.

Command	Description
activeusers	Display active users
apple get	Display apple network settings
apple set	Update apple network settings
clear	Clear the page
date get	Get the current date/timezone information
date set	Set the current date/timezone information

Command	Description
date timezones	List the available timezones (used in conjunction with the date set command)
diskunits	Get status information of all the disk units on the server
domain get	Get the domains known to the SnapServer and their properties
domain list	List the domains known to the SnapServer
dri create	Create a Disaster Recovery Image (dri)
dri recover system	Restore a Disaster Recovery Image (dri)
dri recover volume	Restore a Disaster Recovery Volume Image (dri)
email get	Get email notification settings
email set	Set email notification settings
event clear	Clear all events in the System Event Log
event get	Display the System Event Log
exit	Quit the command line, log off, and exit ssh/bash session. NOTE: If user has started another shell, the command 'exit' will return them to the SnapServer command line shell.
factorydefaults	Reset the SnapServer's settings back to the factory defaults, will reboot
fscheck	Check or repair the user or root filesystem
ftp get	Get the current ftp settings, including anonymous user access
ftp set	Set the current ftp settings, including anonymous user access
globalspares list	List global spares
globalspares remove	Remove a disk from the global spares list
globalspares add	Add a disk to the global spares list
group create	Create a local group
group delete	Delete a local group
group get	Get available groups with their associated information
group list	List available groups
group set	Change the properties of a local group
group member add	Add a group member to a local group
group member delete	Delete a group member from a local group
group members get	Get a list of the members of a local group
group members list	List the members of a local group
history	Print the history of commands typed into the SnapServer command line
homedirs get	Get Home Directory configuration information
homedirs set	Set Home Directory configuration information
hostfile add	Add a host file entry

Command	Description
hostfile delete	Delete a host file entry
hostfile get	Get information for a specific host file entry
hostfile set	Set information for a specific host file entry
hostfile list	List all host file entries
idmap auto map	View/Save auto-generated ID mappings
idmap count	Count number of ID mappings
idmap group get	Get ID mapping for a (windows domain) group
idmap group remove	Remove ID mapping for a (windows domain) group
idmap group set	Set ID mapping for a (windows domain) group to a local or NIS group
idmap list	List all ID mappings
idmap remove all	Remove all ID mappings
idmap update files	Update filesystem for ID mapping changes
idmap update status	View status of ID mapping update filesystem operation
idmap user get	Get ID mapping for a (windows domain) user
idmap user remove	Remove ID mapping for a (windows domain) user
idmap user set	Set ID mapping for a (windows domain) user to a local or NIS user
iscsi create	Create an iscsi disk
iscsi delete	Delete an iscsi disk
iscsi get	Get iscsi disk properties
iscsi set	Set iscsi disk properties
isns get	Get configuration settings for iSNS server
isns set	Set configuration settings for iSNS server
jumboframe get	Get jumbo frame settings for all interfaces
jumboframe list	List jumbo frame settings for all interfaces
jumboframe set	Set jumbo frame settings for all interfaces
less	With a file name, this command allows the user to view any file on the system. It should only be used for 'text' files.
name get	Get the name of the SnapServer
name set	Set the name of the SnapServer
netinfo	Get information about the Ethernet interface
nfs get	Get SnapServer NFS Properties
nfs set	Set SnapServer NFS Properties
nis get	Get current NIS settings
nis set	Set current NIS settings
ntp get	Get NTP client settings
ntp set	Set NTP client settings
ntp_server get	Get NTP Server settings
ntp_server set	Set NTP Server settings

Command	Description
openfiles	List the Open Files
osupdate get	Display status of last OS update
osupdate load	Perform an OS update
passwordpolicy get	Display Password Policy settings and status
passwordpolicy set	Update Password Policy settings
phonehome	Send configuration details to SnapServer Technical Support
proxy get	Display the HTTP proxy properties
proxy set	Set the HTTP proxy properties
quit	Quit the command line, log off, and exit the ssh/bash session
quota list	List user or group quotas for a volume
quota get	Get quota settings for a volume
quota set	Set quota settings for a volume
quota group get	Get volume quota limit & usage for a specific group
quota group set	Set volume quota limit & usage for a specific group
quota user get	Get volume quota limit & usage for a specific user
quota user set	Set volume quota limit & usage for a specific user
reboot	Reboot the SnapServer
securitymodel get	Get the security model on a SnapServer Volume
securitymodel set	Set the security model on a SnapServer Volume
share create	Create a share
share delete	Delete a share
share get	View a share
share rename	Rename a share
share set	Modify a share
share list	List available shares
share access get	Get access list for the share
share access set	Set access list for the share
share access delete	Delete access permission of the specified user/group for the share
share nfsaccess get	Get NFS access permission of the host for the specified share
share nfsaccess set	Set NFS access permission of the host for the specified share
share nfsaccess delete	Delete NFS access permission of the host for the specified share
shutdown	Shutdown the SnapServer
slidingwindow get	Get sliding window settings for a specific interface
slidingwindow set	Set sliding window settings for a specific interface
slidingwindow list	List sliding window settings for all interfaces

Command	Description
snapex	Perform a control operation on the snap extension
snapshot create later	Create a new snapshot schedule
snapshot get	Get snapshot properties
snapshot set	Set properties for the specified snapshot
snapshot list	Get list of snapshots
snapshot create now	Create a new one time snapshot to be run immediately
snapshot delete	Delete specified snapshot
snapshot sched delete	Delete specified snapshot schedule
snapshot sched get	Get specified snapshot schedule
snapshot sched set	Set specified snapshot schedule
snapshot sched list	List current snapshot schedules
snapshot pool get	Get snapshot pool properties
snapshot pool set	Set snapshot pool properties
snapshot pool list	List current snapshot pools
snapshot rollback	Start a rollback for the specified snapshot
snmp get	Get SNMP parameters
snmp set	Set SNMP parameters
ssh get	Get current SSH settings
ssh set	Enable and Disable SSH. NOTE: Turning off SSH while running the command line will 'kick' the user off the system and they won't be able to log back into the command line until SSH is re-enabled via the SnapServer Web Management Interface.
syslog all	Create a tar file of syswrapper and all third-party logs
syslog edr	Create a tar file of Snap EDR logs
syslog s2s	Create a tar file of S2Sv2 logs
syslog syswrapper	Create a tar file of syswrapper only
system type	Available only before storage is configured. Specifies storage type.
system type force	Available only before storage is configured. Allows you to bypass the confirmation prompt.
systemstatus	Get system status information for the server
tape list	List the SCSI tape devices
tape settings get	Display current SCSI tape device settings
tape settings set	Update SCSI tape device settings
tcpip get	Get TCP/IP parameters
tcpip set	Set TCP/IP parameters. NOTE: Changing the parameters of the Ethernet interface over which the user is currently running the SSH/command line session may result in the user being disconnected.

Command	Description
tcpip create bond	Create a bond and set TCP/IP properties.
tcpip delete bond	Remove a TCP/IP bond.
updatenotification get	Get update notification properties
updatenotification set	Set update notification properties
updatenotification check	Check to see if updates are available
ups get	Get UPS settings and status
ups set	Set UPS settings
user create	Create a local user
user delete	Delete a local user
user get	Get available users with their associated information
user list	List available users
user set	Change the properties of a local user
user lock	Lock the specified user.
user unlock	Unlock the specified user.
version	Display current version information, including the Server Number. <i>NOTE: This is the same information displayed in the Web Management Interface "About" box</i>
volume list	List of the volumes defined on the SnapServer
volume get	Get a specific volume's properties
volume create	Create a new logical volume
volume edit	Edit an existing logical volume
volume delete	Delete a logical volume
volume write-cache	Enable or disable write cache on a volume.
vxxaccess list	List hostnames with VSS/VDS access
vxxaccess add	Add hostname of VSS/VDS client requiring access to this server
vxxaccess delete	Delete access for a VSS/VDS client hostname
web get	Get current HTTP Web access settings
web set	Enable or Disable HTTP access to Web Management Interface
windows domain-filter	Allows the user to limit the number of domain users and groups that the server imports from the Active Directory domain, by restricting the import to a specific organizational unit as well as to specific domain users and groups.

The following commands are available only in Traditional RAID:

Command	Description
raid list	List available raids
raid create	Create a raid set
raid delete	Delete a raid set
raid get	Get raid set properties
raid add disk	Add a disk to a raid set
raid remove disk	Remove a disk from a raid set
raid repair	Repair a degraded raid set
raidsettings get	Get auto-incorporation and back-round disk settings
raidsettings set	Set the auto-incorporation and background disk properties
raid-speed-limit get	Get the current setting for the RAID sync speed limit.
raid-speed-limit set	Change the maximum RAID sync or resync speed. Use with caution.

The following commands are available only in DynamicRAID:

Command	Description
storagepool create	Create a new storage pool
storagepool edit	Edit an existing storagepool
storagepool get	Get storagepool properties
storagepool list	List available storagepools
storagepool repair	Repair an existing storagepool
storagepool delete	Delete a storage pool

Scripts in SnapCLI

Administrative tasks can be automated with shell scripts that call SnapCLI commands.

Running a SnapCLI Script

1. Create the script and put it in a share on the local server.
 - Be sure to use an application that is compatible with the standard UNIX text file format (for example, vi). Avoid using Windows clients to create or edit scripts.
 - Place the script in a share that will never be part of a delete script.
2. Log in to the SnapCLI (see [Logging into SnapCLI](#) for instructions).
3. Type `osshell` to get a bash prompt (#).

4. At the prompt, make sure the script is executable by typing the following and pressing **Enter**:

```
chmod +x/shares/[sharename]/[scriptname]
```

where **sharename** is the name of the share where you put the script and **scriptname** is the name of the script.

5. To run the script, type the path again, and press **Enter**:

```
/shares/[sharename]/[scriptname]
```

Sample Script

Following is an example script that can be used to create and remove users, groups, and shares:

```
#!/bin/sh

#####
# Copyright 2003-2007 Overland Storage, Inc. All rights reserved. #
# Permission is granted to use this code provided that it #
# retains the above copyright notice. ##
#####
CLI=/bin/cli
USER=myuser
PASSWORD=myuserpass
GROUP=mygroup
SHARE=myshare
VOLUME=VOL0

# usage: 'mkuser <user_name> <password>'
mkuser()
{

Create a User

# if the user does not exist then create it
if ! $CLI user get user-name="$1" > /dev/null 2>&1; then
echo "Creating user '$1' ..."
$CLI user create user-name="$1" password="$2" > /dev/null 2>&1
if [ $? -ne 0 ]; then
echo "Creation of user '$1' failed."
return 1
fi
else
echo "User '$1' already exists."
fi

return 0
}

# usage: 'mgroup <group_name>'
mgroup()
{
```

Create a Group

```
# if the group does not exist then create it
if ! $CLI group get group-name="$1" > /dev/null 2>&1; then
    echo "Creating group '$1' ..."
    $CLI group create group-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Creation of group '$1' failed."
    fi
return 1
else
echo "Group '$1' already exists."
fi
return 0
}
```

```
# usage: 'adduser2group <user_name> <group_name>'
adduser2group()
{
```

Add the User to the Group

```
# if both the user and the group exist add the user as a member of this group
if $CLI user get user-name="$1" > /dev/null 2>&1; then
if $CLI group get group-name="$2" > /dev/null 2>&1; then
echo "Adding user '$1' to group '$2' ..."
$CLI group member add user-name="$1" group-name="$2" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Adding user '$1' to group '$2' failed."
    fi
return 1
fi
fi
return 0
}
```

```
# usage: 'mkshare <share_name> <share_volume>'
mkshare()
{
```

Create a Share

```
# if the share does not exist create it
if ! $CLI share get share-name="$1" > /dev/null 2>&1; then
echo "Creating share '$1' ..."
$CLI share create share-name="$1" share-volume="$2" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Creating share '$1' failed."
    fi
return 1
else
echo "Share '$1' already exists."
fi
return 0
}
```

```
# usage: 'rmuser <user_name>'
rmuser()
{
```

Delete the User

```
# if the user exists then delete it
if $CLI user get user-name="$1" > /dev/null 2>&1; then
echo "Deleting user '$1' ..."
    $CLI user delete user-name="$1" > /dev/null 2>&1
```

```

        if [ $? -ne 0 ]; then
            echo "Deletion of user '$1' failed."
        fi
    return 1
    else
        echo "User '$1' does not exist."
    fi
return 0
}

```

```

# usage: 'rmgroup <group_name>'
rmgroup()
{

```

Delete the Group

```

# if the group exists then delete it
if $CLI group get group-name="$1" > /dev/null 2>&1; then
echo "Deleting group '$1' ..."
    $CLI group delete group-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Deletion of group '$1' failed."
    fi
return 1
else
    echo "Group '$1P' does not exist."
fi
return 0
}

```

```

# usage: 'rmshare <share_name>'
rmshare()
{

```

Delete the Share

```

# if the share exists delete it
if $CLI share get share-name="$1" > /dev/null 2>&1; then
echo "Deleting share '$1' ..."
    $CLI share delete share-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Deletion of share '$1' failed."
    fi
return 1
else
    echo "Share '$1' does not exist."
fi
return 0
}

```

Create a User, Group, and Share; Then Add the User to the Group

```

#####
#   Main   #
#####

# create a user, a group and a share and add the user to the group
mkuser "$USER" "$PASSWORD"
mkgroup "$GROUP"
adduser2group "$USER" "$GROUP"
mkshare "$SHARE" "$VOLUME"

#remove the group, the user and the share
rmgroup "$GROUP"
rmuser "$USER"
rmshare "$SHARE"

```


GuardianOS Ports

The following table outlines the ports used in GuardianOS.

Port #	Layer	GOS Feature	Name	Comment
1	DDP		rtpm	Routing Table Management Protocol
1	TCP & UDP		tcpmux	TCP port service multiplexer
2	DDP		nbp	Name Binding Protocol
21	TCP & UDP	Network > FTP	ftp	File Transfer Protocol (FTP) port; sometimes used by File Service Protocol (FSP)
22	TCP & UDP	Server > SSH	ssh	Secure Shell (SSH) service
25	TCP & UDP	Server > Email Notification	smtp	Simple Mail Transfer Protocol (SMTP)
67	TCP & UDP	Network > TCP/IP	bootps	Bootstrap Protocol (BOOTP) services; also used by Dynamic Host Configuration Protocol (DHCP) services
68	TCP & UDP	Network > TCP/IP	bootpc	Bootstrap (BOOTP) client; also used by Dynamic Host Control Protocol (DHCP) clients
80	TCP & UDP	Web Management Interface	http	HyperText Transfer Protocol (HTTP) for World Wide Web (WWW) services
81	TCP	Web Management Interface	HTTP	Hypertext Transport Protocol
88	TCP & UDP	Network > NFS	Kerberos	Kerberos Security (NFSv4)
111	TCP & UDP	<ul style="list-style-type: none"> • Networking > NFS • Assist • SnapServer Manager 	sunrpc	Remote Procedure Call (RPC) Protocol for remote command execution, used by Network Filesystem (NFS) and SnapServer Manager
123	TCP & UDP	Server > Date/Time > Advanced	ntp	Network Time Protocol (NTP)
137	TCP & UDP	Network > Windows/SMB	netbios-ns	NETBIOS Name Services used in Red Hat Enterprise Linux by Samba
138	TCP & UDP	Network > Windows/SMB	netbios-dgm	NETBIOS Datagram Services used in Red Hat Enterprise Linux by Samba
139	TCP & UDP	Network > Windows/SMB	netbios-ssn	NETBIOS Session Services used in Red Hat Enterprise Linux by Samba
161	TCP & UDP	Network > SNMP	snmp	Simple Network Management Protocol (SNMP)

Port #	Layer	GOS Feature	Name	Comment
162	TCP & UDP	Network > SNMP	snmptrap	Traps for SNMP
389	TCP & UDP	Network > Windows/SMB	ldap	Lightweight Directory Access Protocol (LDAP)
443	TCP & UDP	<ul style="list-style-type: none"> • Web Management Interface • SnapServer Manager • SnapExtensions > Snap EDR 	https	Secure Hypertext Transfer Protocol (HTTP).
445	TCP & UDP	Network > Windows/SMB	microsoft-ds	Server Message Block (SMB) over TCP/IP
515	TCP	Server > Printing		LPD (Linux Printer Daemon)/LPR (Linux Printer Remote)
631	TCP & UDP	Server > Printing		IPP (Internet Printing Protocol)/CUPS (Common UNIX Printing System)
852	TCP	Network > NFS		Used by rpc.mountd
882	UDP	<ul style="list-style-type: none"> • Snap Finder • SnapServer Manager 	Sysbroker	Broadcast Discovery
933	UDP	Network > NFS		Used by rpc.statd
936	UDP	Network > NFS		Used by rpc.statd
939	TCP	Network > NFS		Used by rpc.statd
957	UDP	Assist		Used by assistrecv
959	TCP	Assist		Used by assistrecv
2005	TCP	SnapExtensions	SnapExtensions	Bridge from Servlet to Snap Extension framework
2049	TCP & UDP	Network > NFS	nfs [nfsd]	Network Filesystem (NFS)
2050	UDP	Network > NFS	mountd	
2051	UDP	Network > NFS	lockd	
2599	UDP	<ul style="list-style-type: none"> • Snap Finder • SnapServer Manager 	Sysbroker	Multicast Discovery
3052	TCP	Server > UPS		Port for monitoring UPS status
3205	TCP	Network > iSCSI	iSNS	
3260	TCP	Network > iSCSI	iSCSI	
8001	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8002	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8003	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8005	TCP	Web Management Interface	tomcat	Tomcat Shutdown port
8008	TCP & UDP	Web Management Interface	http-alt	Tomcat - Apache Bridge
9049	TCP	Sysbroker		Sysbroker Shutdown Port

Port #	Layer	GOS Feature	Name	Comment
9050	TCP	Sysbroker		Sysbroker RPC Port
10001	TCP	Snap Extension	Snap Extension	Shutdown Port
12000	TCP & UDP	Network > Apple/AFP	afp2overtcp	Second NIC
12168	TCP	CA Antivirus	inoweb	Admin Interface
16384	UDP		Sysbroker	Random Port
16388	UDP		Sysbroker	Random Port
24066	TCP		poolmgr	Used by /bin/poolmgr
32780	TCP	Web Management Interface	tomcat	Random Port
32781	TCP	Web Management Interface	tomcat	Random Port
49221	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications Port
49229	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications Port
1024 - 65535	TCP & UDP	<ul style="list-style-type: none"> • Network > NFS • Network > FTP 	NFS FTP (passive)	Dynamically allocated in runtime for user connections

Master Glossary & Acronym List

NOTE: This is a general Overland Storage glossary and acronym list. Not all items may be found in this document or be used by this product.

1000BASE-T

1000BASE-T (also known as IEEE 802.3ab) is a standard for gigabit Ethernet over copper wiring. It requires, at a minimum, Category 5 cable (the same as 100BASE-TX), but Category 5e (Category 5 enhanced) and Category 6 cable may also be used and are often recommended. 1000BASE-T requires all four pairs to be present and is far less tolerant of poorly installed wiring than 100BASE-TX.

Access Permissions

A rule associated with a share, a file, or a directory on a disk drive to regulate which users can have access to the share and in what manner.

Address

An address is a data structure or logical convention used to identify a unique entity, such as a particular process or network device.

ACL

Short for *Access Control List*. A mechanism for restricting access to disk drive directories and files. It is a list of initiator IQNs, along with type of access (read/write or read only) granted to each initiator together with any information required for authentication.

ADS

Short for *Active Directory Service*. The preferred authentication method for Windows XP, Windows 2000, Windows 2000 Advanced Server, and Windows 3000 network users. This authentication allows Active Directory users to connect to shares on the SnapServer. The SnapServer supports the Microsoft Windows 2000 family of servers that run in native ADS mode.

Agent

A program that performs some information-gathering or processing task in the background. SnapServers support Data Protection Agents and can be configured as SNMP agents.

Algorithm

A sequence of steps designed to solve a problem or execute a process.

AllLocalUsers Group

The default group for all local users on SnapServers. Local users are set up by the SnapServer administrator. Network users or Windows domain users are not part of the AllLocalUsers group.

AllUsers Group

A collection of all users. The SnapServer automatically maintains the AllUsers group.

Array

A group of disk drives that are combined together to create a single large storage area. Up to 64 arrays are supported, each containing up to 16 drives per array. There is no capacity limit for the arrays. In a server context, an array refers to the grouping of hard drives into a RAID set.

ATA

Short for *Advanced Technology Attachment*. A standard interface for connecting storage devices to a PC.

Auto Balance

A feature that automatically balances preferred paths evenly among all available host ports and controller ports. Auto balancing spreads I/O load by utilizing as many host ports and controller ports as possible.

Authentication

The validation of a user's identity by requiring the user to provide a registered login name and corresponding password.

Autonegotiation

An Ethernet feature that automatically negotiates the fastest Ethernet speed and duplex setting between a port and a hub or switch. This is the default setting and is recommended.

Autosensing

An Ethernet feature that automatically senses the current Ethernet speed setting.

Back-end

Front-end and back-end are terms used to characterize program interfaces and services relative to the initial user, human or program, of these interfaces and services. A "front-end" application is one that application users interact with directly. A "back-end" application or program serves indirectly in support of the front-end services, usually by being closer to the required resource or having the capability to communicate with the required resource. The back-end application may interact directly with the front-end or, perhaps more typically, is a program called from an intermediate program that mediates front-end and back-end activities.

Back-off Percent

In order to allow drives from a different family or manufacturer to be used as a replacement for a drive in an array, it is recommended that a small percentage of the drive's capacity be reserved when creating the array. This is user selectable, from 0 to 10 percent. This is sometimes known as Reserved Capacity.

Bar Code

The machine-readable representation of a product code. Bar codes are read by a scanner that passes over the code and registers the product code. The width of black lines and white spaces between varies. Combinations of lines and spaces represent characters. Overland uses 3-of-9 code (Code 39) where each character is represented by 9 bars, 3 of which are wide.

Bonding

A technology that treats two ports as a single channel, with the network using one IP address for the server. SnapServer appliances support load balancing and failover bonding modes.

Bridging

Devices that connect and pass packets between two network segments that use different communications protocol.

Bus or Channel

A common physical path composed of wires or other media, across which signals are sent from one part of a computer to another. A channel is a means of transferring data between modules and adapters, or between an adapter and SCSI devices. A channel topology network consists of a single cable trunk that connects one workstation to the next in a daisy-chain configuration. All nodes share the same medium, and only one node can broadcast messages at a time.

CA

Short for *Certificate Authority*. A trusted third-party in a network that issues and manages security credentials.

CA Antivirus

The antivirus software bundled with the SnapServer as a SnapExtension.

Cache Flush Array

This is the array that is used to automatically flush cache data in a situation where power has failed to some of the drives.

Cat 5 Cable

Short for *Category 5*, it is network cabling that consists of four twisted pairs of copper wire terminated by 8P8C modular connectors. CAT 5 cabling supports frequencies up to 100 MHz and speeds up to 100 Mbps. (CAT 5e cabling supports frequencies up to 1000 MHz and speeds up to 1000 Mbps.) It can be used for ATM, token ring, 1000BASE-T, 100BASE-T, and 10BASE-T networking.

Cat 5 is based on the EIA/TIA 568 Commercial Building Telecommunications Wiring Standard developed by the Electronics Industries Association as requested by the Computer Communications Industry Association in 1985.

Cat 6 Cable

Short for *Category 6*, it is network cabling that consists of four twisted pairs of copper wire terminated by 8P8C modular connectors made to higher standards that help reduce noise caused by crosstalk and system noise. The ANSI/TIA-568-B.2-1 specification states the cable may be made with 22 to 24 AWG gauge wire, so long as the cable meets the specified testing standards.

It is designed for Gigabit Ethernet that is backward compatible with the Category 5/5e and Category 3 cable standards. Cat 6 features more stringent specifications for crosstalk and system noise. The cable standard provides performance of up to 250 MHz and is suitable for 10BASE-T / 100BASE-TX and 1000BASE-T (Gigabit Ethernet).

Chaining

A native SnapServer technology in which all snapshots of a volume depend on successive snapshots for part of their content.

Channel

A communications path between two computers or devices.

CHAP

Short for *Challenge Handshake Authentication Protocol*. A three-way handshake scheme used to verify the identity of remote clients in a network.

If there are security concerns, it is possible to set up authentication of targets and initiators, using the CHAP authentication protocol. With CHAP authentication, an initiator can only connect to a target if it knows the target's password or secret. To set up CHAP, the same secret must be known by both the initiator and target.

Checksum

The result of adding a group of data items that are used for checking the group. The data items can be either numerals or other character strings treated as numerals during the checksum calculation. The checksum value verifies that communication between two devices is successful.

Chunk Size

This is the amount of data that is written on a single drive before the controller moves to the next drive in the stripe.

CIFS

Short for *Common Internet Filesystem*. Also known as [SMB](#). The default Windows protocol for communication between computers. A specification for an Internet file access protocol that complements HTTP and FTP and reduces access time.

daemon

A process that runs in the background.

default gateway

The router used when there is otherwise no known route to a given subnet.

degraded

A RAID state caused by the failure or removal of a disk drive in which data is consistent, but there is no redundancy.

DHCP

Short for *Dynamic Host Configuration Protocol*. A communications protocol that lets network administrators centrally manage and automate the assignment of IP addresses on a computer network. Each system that connects to the Internet/intranet needs a unique IP address. A SnapServer can be configured to perform as a DHCP server and assign IP addresses with a single subnet.

Disaster Recovery

A strategy that allows a company to return to normal activities after a catastrophic interruption. Through failover to a parallel system or by restoration of the failed system, disaster recovery restores the system to its normal operating mode.

Discovery

Discovery is the process by which an initiator 'discovers' a target. Discovery uses a special type of session, called a Discovery Session, where an initiator connects to a RAID storage controller and asks it to send a list of the targets present on the controller. The target will respond with a list of all the targets to which the initiator has access.

Disk Roaming

This is the process of removing a disk drive from a controller and putting it back later, either on the same controller, or a different one, and having it recognized as the same disk drive. The disks may be can be attached to different ports than they were originally attached to, without harm to the data. The disks may be attached to the same ports or different ports on the controller.

DNS

Short for *Domain Name Service*. A network service that translates domain names into IP addresses using a server that maintains a mapping of all host names and IP addresses. Normally, this mapping is maintained by the system administrator, but some servers support dynamic mappings.

Domain

A set of network resources in Windows 2000/2003/2008, such as users and groups of users. A domain may also include multiple servers on the network. To gain access to these network resources, the user logs into the domain.

Domain Name

The ASCII name that identifies the domain for a group of computers within a network.

DSM

Short for *Device Specific Module*, it is a software module that allows RAID storage array hardware to use Microsoft's MPIO.

DynamicRAID™

DynamicRAID is a powerful SnapServer feature that simplifies management of disk additions and replacements in a RAID environment. All RAID and filesystem capacity management is entirely automated. More capacity can be added over time by just inserting or replacing drives. Filesystem volumes can be added and removed at will because all volumes share the same underlying pool of storage.

Ethernet

The most widely installed LAN technology. 100BASE-T Ethernet provides transmission speeds of up to 100 Mbps. Fast Ethernet or 1000BASE-T provides transmission speeds up to 1000 Mbps and is typically used for LAN backbone systems, supporting workstations with 100BASE-T cards. Gigabit Ethernet (GbE) provides an even higher level of backbone support at 1000 Mbps (one Gigabit or one billion bits per second).

Ethernet Address

The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet interface.

Ethernet Port

The port on a network card to provide Ethernet access to the computer.

Event

Any significant occurrence or error in the system that may require notifying a system administrator or adding an entry to a log.

Expansion Slot

Area in a computer that accepts additional input/output boards to increase the capability of the computer.

F_port

A *Fabric* port within a Fibre Channel switch that provides a point-to-point link attachment to a single N_Port. F_Ports are intermediate ports in virtual point-to-point links between end ports, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

Failback

Failback occurs when a path with a higher priority than the currently active path is restored. In this case, I/O will “fail back” to the higher priority path once it is available again.

Failover

A strategy that enables one Ethernet port to assume the role of another port if the first port fails. If a port fails on a SnapServer, the second port assumes its network identity (if the two Ethernet cards have been configured for failover). When the port comes back online, the original identities are restored. Failover is possible only in a multi-Ethernet configuration.

Failover/Failback

A combination of Failover and Failback. When a preferred path becomes unavailable, another path is used to route I/O until the preferred path is restored. In this case I/O will “fail back” to the preferred path once it is available again.

FC-AL

Short for *Fibre Channel Arbitrated Loop*. An FC-AL is a Fibre Channel network in which up to 126 systems and devices are connected in a loop topology, with each transmitter connecting to the receiver of the device on its logical right. The Fibre Channel Arbitrated Loop protocol used for transmission is different from Fibre Channel switched and point-to-point protocols. Multiple FC-AL loops can be connected via a fabric switch to extend the network.

Fibre Channel

Fibre Channel (FC) is a gigabit-speed network technology which transports SCSI commands over Fibre Channel networks. Fibre Channel was primarily concerned with simplifying the connections and increasing distances, but later designers added the goals of connecting SCSI disk storage, providing higher speeds and far greater numbers of connected devices.

Filesystem

See [NFS](#).

Firmware

Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first switched on.

FL_port

A *Fabric Loop* port within a Fibre Channel switch that is capable of Fibre Channel Arbitrated Loop operations and is connected to one or more NL_Ports via a Fibre Channel Arbitrated Loop. An FL_Port becomes a shared entry point for public NL_Port devices to a Fibre Channel fabric. FL_Ports are intermediate ports in virtual point-to-point links between end ports that do not reside on the same loop, for example NL_Port to FL_Port to F_Port to N_Port through a single Fibre Channel fabric switch.

Front-end

See [Back-end](#).

FTP

Short for *File Transfer Protocol*. A standard Internet protocol that provides a way to exchange files between computers on the Internet. By default, a SnapServer is set up to be an FTP server.

Full-duplex

A type of transmission that allows communicating systems to both transmit and receive data simultaneously.

Gateway

The hardware or software that bridges the gap between two network subnets. It allows data to be transferred among computers that are on different subnets.

Gigabit Ethernet

Also known as GigE or GbE, this Ethernet standard uses a one Gigahertz (1000 Hz) clock rate to move data.

GID

Short for *Group Identification*. On a SnapServer, the unique ID assigned to each group of users for security purposes.

GuardianOSImage.gsu

An image file used to upgrade the GuardianOS.

HBA

Short for *Host Bus Adapter*. An HBA is an I/O adapter that sits between the host computer's bus and the Fibre Channel loop and manages the transfer of information between the two channels. In order to minimize the impact on host processor performance, the HBA performs many low-level interface functions automatically or with minimal processor involvement.

Half-duplex

A type of transmission that transfers data in one way at a time.

Hidden Share

A share that restricts the display of the share via the Windows (SMB), Web View (HTTP/HTTPS), FTP, and AFP protocols. See also [SMB](#).

Host Bus Adapter

Connects a host system (such as a SnapServer) to other network and storage devices.

Host Name

The unique name by which a computer is known on a network. It is used to identify the computer in electronic information interchange.

Hot Spare

A disk drive that can automatically replace a damaged drive in a RAID 1, 5, 6, 10, 50 or 60. If one disk drive in a RAID fails or is not operating properly, the RAID automatically uses the spare to rebuild itself without administrator intervention. A *local* spare is associated with and available only to a single RAID. A *global* spare is associated with a single RAID, but may be used for any RAID in the system.

Hot Swapping

The ability to remove and add disk drives to a system without the need to power down or interrupt client access to filesystems. Not all components are hot-swappable. Please read installation and maintenance instructions carefully.

HTTP

Short for *Hypertext Transfer Protocol*. An application protocol for transferring files (text, graphic images, sound, video, and other multimedia files) over TCP/IP on the World Wide Web.

HTTPS

Short for *Hypertext Transfer Protocol Secure*. The HTTP protocol using a Secure Sockets Layer (SSL). SSL provides data encryption, server authentication, message integrity, and client authentication for any TCP/IP connection.

IDE

Short for *Integrated Drive Electronics*. A standard interface for connecting storage devices to a PC.

Inheritance

In Windows permissions, inheritance is the concept that when permissions for a folder are defined, any subfolders within the defined folder inherit its permissions. This means an administrator need not assign permissions for subfolders as long as identical permissions are desired. Inheritance greatly reduces administrative overhead and also results in greater consistency in access permission management.

Initialization

RAID 5, 6, 50, and 60 disk arrays must have consistent parity before they can be used to protect data. Initialization writes a known pattern to all drives in the array. If you choose not to initialize an array, the array will be trusted. Any drive failure results in data corruption in a trusted array. (It is possible to later perform a parity rewrite that recalculates the parity based on the current data, thus ensuring the data and parity are consistent.)

Initiator Device

A system component that originates an I/O command over an I/O bus or network. An initiator issues the commands; a *target* receives them.

An initiator normally runs on a host computer. It may be either a software driver or a hardware plug-in card, often called a Host Bus Adapter (HBA). A software initiator uses one of the computer's Ethernet ports for its physical connection, whereas the HBA will have its own dedicated port.

Software initiators are readily available for most host operating systems. Hardware initiators are not widely used, although they may be useful in very high performance applications or if 10 Gigabit Ethernet support is required.

Internal Logical Drive

An internal logical drive is identical to a regular logical drive, except it is NOT made visible to a host adapter as a LUN. Instead, internal logical drives are used for setting up snapshot ODAs that are only accessed internally by the RAID controller.

Internet

A global network of networks used to exchange information using the TCP/IP protocol. It allows for electronic mail and the accessing and retrieval of information from remote sources.

I/O (Input/Output)

The operation of transferring data to or from a device, typically through an interface protocol like CIFS, NFS, or HTTP. The SnapServer presents a filesystem to the user and handles block I/O internally to a RAID array.

IP

Short for *Internet Protocol*. The unique 32-bit value that identifies the location of the server. This address consists of a network address, optional subnetwork address, and host address. It displays as four addresses ranging from 1 to 255 separated by periods.

IQN

Short for *iSCSI Qualified Name*. A name format used in the iSCSI protocol. Initiators and targets have IP addresses, just like any other network entity. They are also identified using an iSCSI name, called the iSCSI Qualified Name (IQN). The IQN should be unique worldwide. It is made up of a number of components, specifying the date, identifying the vendor in reverse format, and then uniquely identifying the initiator or target. An example of an IQN is:

```
iqn.2001-04.com.example:storage:diskarray-sn-123456789
```

Since these IQNs are rather unwieldy, initiators and targets also use short, user friendly names (sometimes called alias names or just aliases).

iSCSI

Short for *Internet SCSI*. iSCSI is an IP-based storage networking standard for linking data storage facilities. iSCSI is a standard that defines the encapsulation of SCSI packets in TCP and then routing it using IP. It allows block-level storage data to be transported over widely used IP networks.

iSNS Server

Short for *Internet Storage Name Service Server*. A protocol enabling the automatic discovery, configuration, and management of iSCSI devices on a TCP/IP network.

Kerberos

A secure method for authenticating a request for a service used by ADS. Kerberos lets a user request an encrypted “ticket” from an authentication process that can then be used to request a service from a server. The user credentials are always encrypted before they are transmitted over the network.

In Windows 2000/XP, the domain controller is the Kerberos server. The Kerberos key distribution center (KDC) and the origin of group policies are applied to the domain.

LACP

Link Aggregation Control Protocol provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).

LAN

Short for *Local Area Network*. A network connecting computers in a relatively small area such as a building.

LCD

Short for *Liquid Crystal Display*. An electronic device that uses liquid crystal to display messages.

LED

Short for *Light-Emitting Diode*. An LED is a type of diode that emits light when current passes through it. Visible LEDs are used as indicator lights on electronic devices.

Linux

A UNIX-like operating system that was designed to provide personal computer users a free or very low-cost operating system comparable to traditional and usually more expensive UNIX systems. GuardianOS is based on the Linux operating system.

Load Balancing

A process available only in multi-Ethernet configurations. The Ethernet port transmission load is distributed among two or more network ports (assuming the cards are configured for load balancing). An intelligent software adaptive agent repeatedly analyzes the traffic flow from the server and distributes the packets based on destination addresses.

Local Group/Local User

A group/user defined locally on a SnapServer using the Web Management Interface. The local user is defined by the server administrator. Windows domain, ADS, and NIS users are not considered local.

Logical Drive

A drive that is defined or created from regions of an array, a whole array, or a combination of regions of different arrays. The logical drive appears as a single disk to one or more host systems.

Logical Drive Availability

To accommodate hosts with multiple ports and multiple host systems, it is possible to restrict a logical drive's availability to a particular HBA or controller port. Access can be enabled or disabled for each host port of each controller.

LUN

Short for *Logical Unit Number*. A SCSI or Fibre Channel device identifier. LUN is a subdivision of a SCSI target.

MAC Address

Short for *Media Access Control address*, a hardware address that uniquely identifies each node of a network. In the Open Systems Interconnection (OSI) model, one of two sublayers of the Data Link Control layer concerned with sharing the physical connection to the network among several computers. Each Ethernet port has a unique MAC address. SnapServer appliances with dual-Ethernet ports can respond to a request with either port and have two unique MAC addresses.

Maintenance Mode

A series of HTML pages in the GuardianOS Web Management Interface that allows you to perform repair, upgrade, or reinstall GuardianOS in a disaster recovery situation.

Mapped LUN Number

Each logical drive is presented to the host system with a unique LUN. In certain cases (such as after deleting another logical drive) it may be desirable to change the number that a logical drive is presented as. This can be done at any time, bearing in mind that any attached host systems may need to be rebooted or reconfigured to maintain access to the logical drive.

Mapping table

A table indexed by sequential LUN values, indicating the selected BUS:TARGET:LUN devices. Mapping tables are used by routers and bridges like the GEOi to perform Ethernet-to-SCSI pathing.

MD5 Algorithm

MD5 is a way to verify data integrity, and is much more reliable than checksum and many other commonly used methods.

MIB

Short for *Management Information Base*. A formal description of a set of network objects that can be managed using the Simple Network Management Protocol (SNMP). The format of the MIB is defined as part of SNMP.

Mirroring

Used in RAID 1 and 10, a process of storing data on one disk and copying it to one or more disks, creating a redundant storage solution. RAID 1 is the most secure method of storing mission-critical data.

Mounted

A filesystem that is available.

MPIO

Short for *Multipath Input/Output*. A multipath solution built into Microsoft server-grade operating systems. It requires the DSM to work with RAID storage array hardware.

MTU

Short for *Maximum Transfer Unit*. It is the largest size packet or frame, specified in octets (eight-bit bytes), that can be sent in a packet- or frame-based network.

Multihomed

A SnapServer that is connected to two or more networks or has two or more network addresses.

N_port

A *Node* port connects via a point-to-point link to either a single N_Port or a single F_Port. N_Ports handle creation, detection, and flow of message units to and from the connected systems. N_Ports are end ports in virtual point-to-point links through a fabric, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

NAS

Short for *Network Attached Storage*. Hard disk storage that is set up with its own network address as opposed to being attached to the department computer that is serving applications to a network's workstation users. By removing storage access and its management from the department server, both application programming and files can be served faster because they are not competing for the same processor resources. The NAS device is attached to a local area network (typically an Ethernet network) and assigned an IP address.

NAT

Short for *Network Address Translation*. A technique for passing network traffic through a router whereby one set of IP addresses is used on one side of the router and another set of addresses is used on the other side. This is done to avoid address conflicts and to increase the address space of the internal network.

NDMP

Short for *Network Data Management Protocol*. A protocol standard used by some Network Attached Storage systems to provide an industry standard means to do backup and restores of the NAS system without the need for 3rd party agents to be installed on the NAS device. Also see NDMP.org for further details.

NFS

Short for *Network Filesystem*. A client/server application that allows a computer user to view and optionally store and update files on a remote computer as though they were on the user's own computer. The user's system needs to have an NFS client and the other computer needs the NFS server. The SnapServer is configured as an NFS server by default.

NIC

Short for *Network Interface Card*. A board that provides network communication capabilities to and from a computer.

NIS

Short for *Network Information Service*. A network naming and administration system for smaller networks that was developed by Sun Microsystems. NIS+ is a later version that provides additional security and other facilities. The SnapServer accepts NIS users and groups.

NL_port

A *Node Loop* port is capable of arbitrated loop functions and protocols. An NL_Port connects via an arbitrated loop to other NL_Port and at most a single FL_Port. NL_Ports handle creation, detection, and flow of message units to and from the connected systems. NL_Ports are end ports in virtual point-to-point links through a fabric, for example NL_Port to F_Port

to F_Port to N_Port using a single Fibre Channel fabric switch. In the absence of a fabric switch FL_Port, NL_Ports can communicate with other NL_Ports in virtual point-to-point links through a FC-AL open loop circuit often through FC-AL (Arbitrated Loop) hub or loop switch devices.

Node

Any device, including servers, workstations, or tape devices, that are connected to a network; also the point where devices are connected.

Node Name

This is an eight-byte, 16-character hexadecimal number, uniquely identifying a single fibre device. It incorporates the World Wide Name and two additional bytes that are used to specify the format. In a host system with multiple FC ports, all adapters typically use the same Node Name, but unique Port Names.

NTFS

Short for *New Technology File System*. The standard file system used by Windows NT and later versions of the Windows operating system.

NTP

Short for *Network Time Protocol*. A protocol for synchronizing the system clocks of computers over a packet-switched network.

NVRAM

Abbreviation of *Non-Volatile Random Access Memory*, a type of memory that retains its contents when power is turned off.

ODA

The *Overwrite Data Area* is an internal storage area on an array that is dedicated to storing data from a snapshot logical drive. The data stored on the ODA is the data from the logical drive that needed to be overwritten after a snapshot was initiated. The ODAs are mapped on top of internal logical drives. An ODA cannot be accessed externally through a host LUN; it is only accessed internally.

ODA Stripe Size

The read/write block size that the system will use when copying data from the original logical drive to the ODA.

Orphan

A disk drive that has become disconnected from its RAID either by accidental removal of the drive or the intermittent failure of the drive.

Parity

Error correction data. RAID5, RAID6, RAID50, and RAID60 store equal portions of each file on each disk and distributes parity information for each file across all disks in the group. This distributed parity allows the system to recover from a single disk drive failure.

Permissions

A security category, such as no access, read-only, or read-write, that determines what operations a user or group can perform on folders or files.

Pool

A pool is a collection of RAID disks, grouped together by the RAID storage controller. iSCSI volumes are created from these pools. New volumes can be created and existing volumes can be extended, provided there is spare capacity in the pool from which the volume was created.

PoP

Short for *Proof of Purchase*. The number used to obtain a license key for an upgrade to third-party applications.

Port Name

This is an eight-byte hexadecimal number, uniquely identifying a single host **HBA** port. It incorporates the World Wide Name and two additional bytes that are used to specify the format and indicate the port number.

Portal

A target's IP address together with its TCP port number.

POSIX

Short for *Portable Operating System Interface*. A set of standard operating system interfaces based on the UNIX operating system. The need for standardization arose because enterprises using computers wanted to develop programs that could run on multiple platforms without the need to recode.

Preferred Path

The preferred path is the default path. When the path selection policy is set to Failover/Failback, the preferred path is always used if it is available. If the preferred path fails, I/O switches to another path. If it is later restored, I/O switches back to the preferred path.

Protocol

A standardized set of rules that specifies the format, timing, sequencing, and/or error checking for data transmissions.

PTP

Short for *Point-to-Point*. PTP is the common mode of attachment to a single host. PTP is sometimes used to attach to a Fibre Channel switch for **SAN** connectivity.

Public Access Share

A share that allows all users read/write access to the filesystem.

Quota

A limit on the amount of storage space on a volume that a specific user or NIS group can consume.

RAID

Short for *Redundant Array of Independent Disks*. A data storage scheme where multiple hard drives are combined to form a single logical unit which is highly reliable and gives good performance. Reliability is achieved by mirroring (the copying of data to more than one disk), striping (the splitting of data across more than one disk) and error correction (redundant data is stored to enable faults to be detected and corrected).

RAID 0 (Striped)

RAID 0 is ideal for environments in which performance (read and write) is more important than fault tolerance, or you need the maximum amount of available drive capacity in one volume.

Data is striped across multiple disks so that it can be read and written in parallel. It provides higher performance than a single disk, especially when reading or writing large files, but it is vulnerable to a disk failure. If any disk in the pool fails, the entire pool is effectively lost. For this reason, RAID 0 pools should only be used in cases where the loss of the data is unimportant, for example, because it can easily be recreated from another data source. The capacity of a RAID 0 pool is equal to the total capacity of all the disks making up the pool¹. For example, a RAID 0 pool made up of 4 x 1 TB disks will have a capacity of 4 TB.

RAID 1 (Mirrored)

RAID 1 is useful for building a fault-tolerant system or data volume, providing excellent availability without sacrificing performance. However, you lose 50 percent of the assigned disk capacity.

RAID 1 is also called disk mirroring: data is stored on two identical disks, so that if one disk fails, the other can still be used to access the data. Write operations are performed in parallel to both disks, so write performance is identical to that of a single disk; read operations can be done to either disk, so effectively read performance is doubled.

If one of the disks fails, it should be replaced. When it is replaced, the RAID pool will automatically be rebuilt by copying all the data from the surviving disk to the new disk. While the rebuild is occurring, there will be a degradation in performance.

Because disks are mirrored, the usable capacity of a pair of RAID 1 disks is only equal to the capacity of a single disk, so that a RAID 1 pool made of 2 x 500 GB disks will have a capacity of 500 GB.

RAID 5 (Striping with Parity)

With a RAID 5 pool, because data is read from many disks in parallel, as for RAID 0, read performance is good. Write performance is slightly lower because, in addition to writing the data, parity data has to be calculated and written. If a hardware RAID controller is used, this will be done using dedicated hardware; if software RAID is used, the work will be done on the main processor of the storage controller.

The capacity of a RAID 5 pool is reduced by exactly one disks worth of capacity, which is required to store the parity data. For example, a RAID 5 pool made up of 3 x 500 GB disks will have a capacity of 1 TB.

In principle, a RAID 5 pool could have a very large number of disks. However, the more disks there are, the greater the chance of a double disk failure. If a single disk fails, the data is no longer protected until the disk has been replaced and the pool has been rebuilt by reconstructing all the data from the failed disk and writing it to the new disk. If the disk capacities are very large, it may take many hours to rebuild the pool. If a second disk fails before the rebuild has completed, all the data in the pool will be lost. That is to say, large capacity disks increase the time taken to rebuild the pool, during which time the pool is vulnerable to a second disk failure. Moreover, the chance of a second disk failure increases as the number of disks in the pool increases.

RAID 5 is similar to RAID 0 in that data is striped across multiple disks. However, one disks worth of space is reserved to store parity data, which can be used to reconstruct the pool in the event of one of its disks failing. With RAID 5, the parity data is distributed across all the

¹Capacity is usually very slightly less because a small but insignificant amount of space is reserved by the RAID controller to store internal metadata.

disks in the pool. If a single disk fails, each block of data stored on that disk can be reconstructed using the corresponding data block from all the other disks along with the parity block. This means that if a single disk fails, data can still be read, albeit at a rather slower rate (because it needs to be reconstructed, rather than read directly). For this reason, a RAID 5 pool with a disk failure is referred to as a degraded pool.

RAID 6 (Striping with Dual Parity)

RAID 6 is similar to RAID 5 but instead of storing a single disk's worth of parity data, two disk's worth are stored, making the pool capable of withstanding the failure of two disks. However, there is an additional write overhead involved in calculating the double parity data. Since RAID 6 works best with dedicated hardware, RAID 6 is only offered on systems with a hardware RAID controller. Read performance is similar to that of RAID 0 or 5. Since two disks are used for storing parity data, the capacity of a RAID 6 pool made up of 8 x 500 GB disks will be 3 TB.

RAID 10 (Striped Mirroring)

RAID 10 is defined as mirrored stripe sets or also known as RAID 0+1. You can build RAID 10 either directly through the RAID controller (depending on the controller) or by combining software mirroring and controller striping, or vice versa (called RAID 01).

RAID 50

A RAID 50 combines the straight block-level striping of RAID 0 with the distributed single parity of RAID 5. That is, a RAID 0 array striped across RAID 5 elements. It requires at least 6 disks. This can increase the performance by allowing the controller to more efficiently cluster commands together. Fault tolerance is also increased, as one drive can fail in each individual array.

RAID 60

A RAID 60 combines the straight block-level striping of RAID 0 with the distributed double parity of RAID 6. That is, a RAID 0 array striped across RAID 6 elements. It requires at least 8 disks. This can increase the performance by allowing the controller to more efficiently cluster commands together. Fault tolerance is also increased, as two drives can fail in each individual array.

Recurring Snapshot

A snapshot that runs at an administrator-specified time and interval.

Restrict Anonymous

A Windows feature in which anonymous users cannot list domain user names and enumerate share names. Microsoft has provided a mechanism in the Registry called restrict anonymous for administrators to restrict the ability for anonymous logon users (also known as NULL session connections) to list account names and enumerate share names.

The implementation of the restrict anonymous mechanism may prevent the SnapServer from obtaining the list of account names it needs to authenticate Windows domain users.

Resynchronization

A RAID state that describes the process of integrating a new drive into the RAID.

RETMA

Short for *Radio-Electronics-Television Manufacturers' Association*. It is the common name given for a 19-inch distribution frame rack for mounting components.

Rollback

A snapshot feature that allows the administrator to restore a volume to a previous state as archived in a snapshot without resorting to tape.

Round Robin

The Round Robin path selection policy causes all healthy paths to be used for I/O. Paths are used in a round-robin order.

Router

A router is a device that enables connectivity between Ethernet network segments.

SAN

Short for *Storage Area Network*. Data storage connected to a network that provides network clients access to data using block level protocols. To the clients, the data storage devices appear local rather than remote. An iSCSI SAN is sometimes referred to as an IP-SAN.

SAS

Short for *Serial Attached SCSI*. It is a point-to-point serial protocol that replaces parallel SCSI bus technology (multidrop) and uses the standard SCSI command set. It has no termination issues, supports up to 16,384 devices (using expanders), and eliminates clock skew. It consists of an Initiator that originates device service requests, a Target containing logical units that receives device service requests, and a Service Delivery Subsystem that transmits information between the Initiator and the Target.

SCSI

Short for *Small Computer System Interface*. SCSI is an industry standard for connecting peripheral devices and their controllers to an initiator. Storage devices are daisy-chained together and connected to a host adapter. The host adapter provides a shared bus that attached peripherals use to pass data to and from the host system. Examples of devices attached to the adapter include disk drives, CD-ROM discs, optical disks, and tape drives. In theory, any SCSI device can be plugged into any SCSI controller.

SCSI addressing

Each device supported by a SCSI adapter has its own unique SCSI address, which dictates the device's priority when arbitrating for access to the SCSI bus. A SCSI address of 7 has the highest priority. For a fast/wide SCSI adapter that supports up to 16 devices, the next highest priority address is 6, then 5, 4, 3, 2, 1, 0, 15, 14, 13, 12, 11, 10, 9, and 8. The narrow SCSI adapter supports up to eight devices, including itself. The SCSI address 7 has the highest priority, followed by 6, 5, 4, 3, 2, 1, and 0.

SCSI bus

A SCSI bus provides a means of transferring data between SCSI devices. A SCSI bus is either an 8- or 16-bit bus that supports up to 8 or 16 devices, including itself. The bus can consist of any mix of initiators and targets, with the requirement that at least one initiator and one target must be present.

SCSI device

A SCSI device is a single unit on a SCSI bus that originates or services SCSI commands. A SCSI device is identified by a unique SCSI address. SCSI devices can act as initiators or targets.

SCSI port

A SCSI port is an opening at the back of a router that provides connection between the SCSI adapter and SCSI bus.

Serial Number

The ten-character alphanumeric number assigned by the manufacturer at the factory.

Server Number

A numeric derived from the MAC address of your SnapServer's primary Ethernet port that is used to uniquely identify a SnapServer.

Session

When an initiator wants to establish a connection with a target, it establishes what is known as an iSCSI session. A session consists of one or more TCP/IP connections between an initiator and a target. Sessions are normally established (or re-established) automatically when the host computer starts up, although they also can be established (and broken) manually.

Share

A virtual folder that maps to the root of a volume or a directory on the volume. Permissions are assigned to a share that determine access for specific users and groups.

Share Access

Permissions granted or denied to users and groups that control user and group access to the files.

S.M.A.R.T.

Short for *Self Monitoring, Analysis and Reporting Technology*. A standard mechanism for querying disk drives to monitor performance and reliability attributes, such as temperature, read error rates and seek times. S.M.A.R.T. systems are built into most modern disk drives.

SMB

Short for *Server Message Block*. A protocol for Windows clients. SMB uses the TCP/IP protocol. It is viewed as a complement to the existing Internet application protocols such as FTP and HTTP. With SMB, you can access local server files, obtain read-write privileges to local server files, share files with other clients, and restore connections automatically if the network fails.

SMS

Short for *Short Message Service*. Is a means of sending short text messages to a mobile phone.

SMTP

Short for *Simple Mail Transfer Protocol*. A TCP/IP protocol used for sending and receiving email.

Snap EDR

A SnapExtension that copies the contents of a share from one SnapServer to another share on one or more SnapServers. Snap EDR is designed to work with SnapServers and other SnapServer Storage Solutions.

Snapback

The process of restoring a logical drive from a selected snapshot. This process takes place internally in the RAID controller firmware and needs no support from any backup utility.

SnapDRImage

The SnapServer disaster recovery image that saves server-specific settings such as server name, network, RAID, volume and share configuration, local user and group lists, and snapshot schedules.

SnapExtension

A Java application that extends a SnapServer's functionality. SnapExtensions are produced both by SnapServer and third-party vendors.

SnapServer Manager

The SnapServer Manager (SSM) is a Java-based utility for discovering and monitoring SnapServers.

Snapshot

A method for producing a point-in-time image of a logical drive that results in a consistent, stable, point-in-time image of a volume (filesystem) used for backup purposes. In the process of initiating a snapshot, no data is actually copied from the snapshot logical drive. However as new writes are made to a snapshot logical drive, existing data blocks are copied to the [ODA](#) before the new data is written to the logical drive.

Snapshot LUN

A special LUN created from a combination of the snapshot logical drives' data and the data contained in the [ODA](#).

Snapshot Number

Identifier that references one of several snapshots of the same logical drive.

Snapshot Pool

Disk space reserved within a RAID for the storage of snapshot data. In the default storage configuration of many SnapServers, twenty percent of the RAID capacity is allocated to the snapshot pool.

Snapshot Share

A virtual folder that allows access to all current snapshots at the same directory level as the original share on which it is based.

SNMP

Short for *Simple Network Management Protocol*. A system to monitor and manage network devices such as computers, routers, bridges, and hubs. SNMP views a network as a collection of cooperating, communicating devices, consisting of managers and agents.

SSH

Short for *Secure Shell*. A service that provides a remote console for special system administration and customer support access to the server. SSH is similar to telnet but more secure, providing strong encryption so that no passwords cross the network in clear text.

SSL

Short for *Secure Sockets Layer*. A protocol for managing the security of a message sent on the Internet. It is a type of technology that provides data encryption, server authentication, message integrity, and client authentication for any TCP/IP connection.

Standalone

A network bonding mode which treats each port as a separate interface. This configuration should be used only in multihomed environments in which network storage resources must reside on two separate subnets.

Static IP Address

An IP address defined by the system administrator rather than by an automated system, such as DHCP. The SnapServer allows administrators to use DHCP-assigned or statically assigned IP addresses.

Storage Area Network

See [SAN](#).

Stripe

The process of separating data for storage on more than one disk. For example, bit striping stores bits 0 and 4 of all bytes on disk 1, bits 1 and 5 on disk 2, etc.

Stripe Size

This is the number of data drives multiplied by the chunk size.

Sub-array

In RAID 50 applications, this is the name given to the individual RAID 5 arrays that are striped together. Each sub-array has one parity drive.

Subnet Mask

A portion of a network that shares a common address component. On TCP/IP networks, subnets are all devices with IP addresses that have the same prefix.

Target

A target is a device (peripheral) that responds to an operation requested by an initiator (host system). Although peripherals are generally targets, a peripheral may be required to act temporarily as an initiator for some commands (for example, SCSI COPY command).

Targets are embedded in iSCSI storage controllers. They are the software that makes the RAID storage available to host computers, making it appear just like any other sort of disk drive.

TCP/IP

Short for *Transmission Control Protocol/Internet Protocol*. The basic protocol used for data transmission over the Internet.

Telco

Short for *Telephone Company*. When used in reference to a rack, it refers to the two-posted, light-weight rack for center-mounted appliances.

Telnet

A terminal emulation program for TCP/IP networks such as the Internet. The Telnet program runs on a computer and connects it to a server on the network. You enter commands through the Telnet program and they will be executed as if you were entering them directly on the server console. This enables you to control the server and communicate with other servers on the network. To start a Telnet session, you must log in to a server by entering a valid user name and password. Telnet is a common way to remotely control Web servers.

Terminator

A terminator refers to the electrical connection at each end of a SCSI bus. The terminator is composed of a set of resistors, or possibly other components. The function of a terminator is to provide a pull-up for open collector drivers on the bus, and also impedance matching to prevent signal reflections at the ends of the cable. SCSI buses require that a terminator be placed on the SCSI connector on the last SCSI peripheral. Data errors may occur in a SCSI bus that is not terminated.

TOE (TCP Offload Engine)

Short for *TCP Offload Engine*. TOE is a technology used in network interface cards to offload processing of the entire TCP/IP stack to the network controller. It is primarily used with high-speed network interfaces, such as gigabit Ethernet and 10 gigabit Ethernet, where processing overhead of the network stack becomes significant.

Topology

Logical layout of the parts of a computer system or network and their interconnections. There are two types of topology: physical and logical. The physical topology of a network refers to the configuration of cables, computers, and other peripherals. Logical topology is the method used to pass the information between workstations.

Trap

A signal from a device informing an SNMP management program that an event has occurred.

U

A standard unit of measure for designating the height in computer enclosures and rack cabinets. One U equals 1.75 inches. For example, a 3U server chassis is 5.25 inches high.

UDP

Short for *User Datagram Protocol*. A communications protocol for sending messages between computers in a network that uses the Internet Protocol (IP). UDP is an alternative to the Transmission Control Protocol but, unlike TCP, does not guarantee reliability or ordering of data packets.

UID

Short for *User Identification*. A unique ID assigned to each user on a SnapServer for security purposes.

Unassigned

The state of a disk drive that is seated in a bay but has not been incorporated into a RAID.

Unassigned Free Space

The controller keeps a map of all the space that is not assigned to any logical drive. This space is available for creation or expansion. Each unassigned region is individually listed.

UNC

Short for *Universal Naming Convention*. In a network, a way to identify a shared file in a computer without having to specify (or know) the storage device it is on. In the Windows OS, the UNC name format is as follows:

```
\\server_name\share_name\path\file_name
```

UPS

Short for *Uninterruptible Power Supply*. A device that allows a computer to keep running for a short time when the primary power source is lost. It also provides protection from power surges. A UPS device contains a battery that starts when the device senses a loss of power from the primary source.

URL

Short for *Uniform Resource Locator*. A Web address.

USB Port

USB is short for *Universal Serial Bus*. A USB port is a hardware interface for low-speed peripherals such as the keyboard, mouse, joystick, scanner, printer, and telephony devices.

VDS

Short for *Virtual Disk Service*. VDS is a feature of Microsoft Windows (from Windows Server 2003 onwards). It provides a consistent interface for managing storage devices and creating volumes. Each vendor of a storage solution can write their own hardware provider module that enables the standard set of VDS commands to be used with different enclosures. Thus, multiple storage systems by different vendors can be controlled using the same set of VDS commands.

Virtual LUN

See [Snapshot LUN](#).

VLAN

Short for *Virtual LAN*. It consists of a network of computers that behave as if they are connected to the same wire - even though they may actually be physically connected to different segments of a LAN.

Volumes

A logical partition of a RAID's storage space that contains a filesystem. Volumes are created from storage pools, using unused capacity in a pool. They can be extended in size, so long as there is free capacity in the pool.

A volume appear to a host computer just like a regular, physical disk except that it is attached by means of iSCSI instead of traditional disk interconnects such as IDE, SCSI or SATA.

Each volume has an iSCSI target associated with it. The volume is mapped to Logical Unit Number (LUN) 0 of the iSCSI/SCSI target, just like a regular physical disk. Associated with the target is an Access Control List (ACL) that defines which host systems are allowed to access the volume.

When the iSCSI initiator on the host computer connects to the iSCSI target, the iSCSI volume becomes available for use.

VSS

Short for *Volume Shadow Copy Service*. A low level communications interface that enables volumes to be backed up without having to halt all applications that are reading or writing the volumes. Microsoft VSS provides a mechanism for creating consistent point-in-time copies of data known as shadow copies.

Web Management Interface

A Web-based utility used for configuration and ongoing maintenance, such as monitoring server conditions, configuring email alerts for key events, or for SNMP management.

Web View

The Web-browser page that opens when users access a SnapServer using their Web browsers, and displays a list of all shares.

Windows Domain Authentication

Windows-based networks use a domain controller to store user credentials. The domain controller can validate all authentication requests on behalf of other systems in the domain. The domain controller can also generate encrypted challenges to test the validity of user credentials. Other systems use encrypted challenges to respond to CIFS/SMB clients that request access to a share.

WINS

Short for *Windows Internet Naming Service*. The server that locates network resources in a TCP/IP-based Windows network by automatically configuring and maintaining the name and IP address mapping tables.

Workgroup

A collection of computers that are grouped for sharing resources such as data and peripherals over a LAN. Each workgroup is identified by a unique name.

Write-Back Cache

A caching method in which modifications to data in the cache aren't copied to the cache source until absolutely necessary. Write-back caching yields somewhat better performance than write-through caching because it reduces the number of write operations to main memory. With this performance improvement comes a slight risk that data may be lost if the system crashes.

Symbols

> (menu flow indicator) **PR-iv**

A

access

problems with **D-4**

Access Denied message **D-5**

ACLs

backing up **9-5**

resetting to defaults **9-4**

setting file-level permissions (Windows) **C-6**

Active Directory

and name resolution servers **3-9**

joining AD domain **3-12**

SnapServer interoperability with **3-10**

admin password

default **7-2**

resetting forgotten **9-3, D-7**

Administration page **10-4**

AFP terminology **3-15**

AFP, see *Mac OS*

alert definitions **PR-iv**

Antivirus

dependencies on other software components **11-1**

distributing updates **11-6**

excluding snapshots from **11-3**

HTTP requirement **11-1**

launching configuration GUI **11-2**

scan job configuration **11-3**

using logs **11-9**

Authentication

default settings **7-2**

HTTPS/HTTP **3-24**

Kerberos **3-10**

NIS domain **3-20**

automatic incorporation, and previously configured drives **5-11**

automatic shutdown **2-18**

B

background disk scan **5-11**

backup

coordinating with snapshots **6-5**

identifying backup/media servers to the SnapServer **B-4**

inability to back up iSCSI Disks **6-4, 6-17**

iSCSI Disks **6-17**

of server and volume settings **9-5**

off-the-shelf solutions **B-3**

backup.acl **9-5**

backup.qta.groups **9-5**

backup.qta.users **9-5**

C

CA Antivirus, see *Antivirus*

CA ARCserve, installing agent **B-5**

CA Unicenter TNg **3-23**

change password **10-9**

changing server name **2-14**

CLI connection via SSH **2-16**

client access, configuring

Apple **3-14**

FTP **3-21**

HTTPS/HTTP **3-24**

NFS **3-16**

cloning a server **9-8**

Command Line Interface **E-1**

running scripts **E-9**

syntax **E-1**

connecting

a Mac OS X client **3-10**

- from a Windows client **3-10**
- to SnapServers **1-7**
- conventions, typographical **PR-iv**
- customer support **PR-iii**

D

- data import **9-9**
- data protection tasks **2-12**
- date and time settings **2-15**
- defaults
 - admin password **7-2**
 - TCP/IP **3-3**
- directories, home **7-36**
- Disaster Recovery
 - backing up server and volume settings **9-5**
 - creating recovery files **9-5**
- disk drives
 - adding
 - in DynamicRAID **4-2, 4-9, A-5**
 - in Traditional RAID **5-13**
 - automatic incorporation **5-11**
 - and previously configured drives **5-11**
 - detecting **2-6, 2-8**
 - hot swap **6-26**
 - incompatible **4-2**
 - LED indicators **D-2**
 - previously configured **4-2, A-5**
 - in different system **2-7, 4-2**
 - in Traditional RAID **5-14**
 - reintegrating orphaned **6-32**
 - replacing **6-26**
- documents, related to SnapServers **PR-v**
- domain search
 - authentication required **5-22, 7-13, 7-27, 7-31**
- domains
 - joining NIS **3-20**
- download website link **PR-v**
- drive-size characteristics, DynamicRAID **A-5**
- dual parity **6-26**
- dynamic volumes **A-4**
- DynamicRAID **4-1**
 - compared to Traditional RAID **A-3**
 - drive indicators **A-5**
 - how it works **A-4**
 - implementation **A-4**
 - storage pools **4-2**
 - volumes **4-10**

E

- electrostatic discharge information **PR-v**
- email notification of server events **9-15**
- EMC NetWorker
 - installing agent **B-10**
 - special backup and restore operations **B-12**
- Ethernet, see *Gigabit Ethernet*
- Expand Volume button **5-18**
- expansion units **6-32**
 - adding to Traditional RAID **5-10**
 - DynamicRAID usage **4-6**
 - integrating orphans **6-34**
 - SnapExpansion usage **6-32**
 - Traditional RAID usage **5-10**
- exports file, NFS **7-6**

F

- failover, see *Network bonding*
- features, changes from earlier releases **1-4**
- files, setting permissions for **C-6**
- FTP
 - configuring access **3-20**
 - connecting via **3-22**

G

- GID **7-3**
- Gigabit Ethernet
 - autonegotiation required **3-6**
- global hot spares **5-3**
- Global Spare **6-31**
- Groups
 - creating local **7-23**
 - file-level access for **C-6**
 - joining NIS domain **3-20**
- GuardianOS
 - ports **F-1**
 - specifications **1-1**
 - updating **9-13**
 - what's new in **1-4**

H

- hardware information screen **2-12**
- Head Unit properties page **4-5**
- home directories **7-36**

- Home page **10-2**
- Host File Editor **B-4**
- hot spares **5-3**
- hot swap
 - and automatic incorporation of disks **5-11**
 - disk drive **6-26**
- HP Open View **3-23**
- HTTPS/HTTP
 - configuring **3-24**

I

- ID mapping **7-26**
- increase capacity **6-32**
- Initial Setup
 - configuring storage
 - in DynamicRAID **2-6**
 - in Traditional RAID **2-8**
 - wizard **2-1**
- integrate expansion units **6-34**
- internal temperature, e-mail notification of **9-16**
- IP address
 - setting **3-5**
 - using SSM to discover **1-8**
- IPP port number
 - Linux **2-21**
 - Windows **2-21**
- iSCSI disks **6-9**
 - and DynamicRAID **A-5**
 - backing up **6-22**
 - configuring iSNS **3-28**
 - creating **6-19**
 - LUNs **6-24**
 - multi-initiator support **6-18**
 - write cache options **6-18**
- iSNS **3-28**

K

- Kerberos **3-10**

L

- LEDs
 - disk drive indicators **D-2**
 - in DynamicRAID **A-5**
 - network **D-3**
 - power/unit status **D-3**

- system/status **D-3**
- understanding **D-1**
- load balancing, configuring server for **3-5**
- local groups **7-22**
- local hot spares **5-3**
- login, to antivirus GUI **11-2**

M

- Mac OS
 - configuring client access **3-14**
 - Finder **3-15**
 - launching SnapServer Manager on **1-6**
- Macintosh, supported OS versions **1-2**
- maintenance
 - data import **9-9**
 - disaster recovery **9-4**
 - factory defaults **9-3**
 - OS update **9-13**
 - restore **9-6**
 - shutdown and restart **9-2**
 - support **9-14**
 - tools **9-15**
- managing snapshots **6-2**
- mapping, ID **7-26**
- menu flow indicator **PR-iv**
- mixed drive capacities **A-3**
- mixed drive types **A-4**
- monitoring
 - system **8-1**
 - tape **8-5**
- Multihomed configurations **3-6**

N

- network
 - access **3-1**
 - bonding, cabling requirements for **3-6**
 - current settings **3-2**
 - LED indicators **D-3**
 - problems with access **D-4**
 - reset to factory defaults **9-3**
- Network Time Protocol (NTP) **2-16**
- NFS
 - access **3-16**
 - and share-level permissions **7-14**
 - configuring **3-16**
 - exports file **7-6**

read-only share access **3-17**
 NIS domains **3-20**

O

orphaned disk drive **6-32**
 orphaned disk drives **6-32**
 OS update **9-13**
 Overland technical support **PR-iii**

P

parity
 adding disk drives to upgrade **4-9**
 configuring DynamicRAID **A-4**
 disk drive failure
 dual **4-9**
 single **4-9**
 increasing protection **4-9**
 management **4-8**
 options **A-3**
 Parity Mode **4-3, 4-6**
 Parity Mode change **6-30**
 removing a drive **6-29**
 parity, dual **6-26**
 password
 changing **10-9**
 default for admin account **7-2**
 unlock **7-20**
 paths
 connecting via web browser **3-25**
 for backing up snapshots **6-5**
 for distributing antivirus updates **11-6, 11-7**
 for restoring a “cured” file **11-8**
 permissions
 share- and file-level interaction **7-12**
 file-level, default behavior **C-6**
 Phone home support **D-8**
 power/status LED **D-3**
 Print Server **2-20**
 adding a printer **2-20**
 canceling print jobs **2-22**
 configure the printer **2-20**
 deleting a printer **2-22**
 IPP port number, Linux **2-21**
 IPP port number, Windows **2-21**
 monitoring print jobs **2-21**
 pausing the printer **2-21**

product documentation **PR-iii**

Q

Quota
 Sort drop-down list **5-23**
 View drop-down list **5-23**
 Quotas
 adding **5-23**
 backing up configuration **9-5**
 defaults **5-19**
 displaying **5-23**
 properties **5-21**
 usage calculation **5-20**

R

RAID
 types defined **G-17**
 adding disk drives to **5-13**
 choosing **5-2**
 creating new **5-4**
 effect of deleting on antivirus software **11-1**
 grouped
 deleting **5-10**
 grouping **5-7**
 multiple RAIDs **5-10**
 with other grouped RAIDs **5-10**
 scrubber **5-11**
 sets
 creating new **5-4**
 grouping **5-7**
 screen **5-4**
 settings **5-11**
 Traditional RAID and replacement disks **5-13**
 type selection **2-4**
 reboot, setting up alert for **9-16**
 reduced parity with Traditional RAID **5-13**
 rejoining servers to a Windows domain **9-7**
 remote SnapServer discovery **10-7**
 replacing disk drives **6-26**
 replication **B-2**
 reset options **D-3**
 restart **9-2**
 restore **9-6**
 resynchronization, setting alert for completion of **9-16**

SSecure Shell (SSH) **2-16**

security

guides **7-3**models **7-25**resetting default ACLs for volumes **9-4**shares **7-6**Windows ACLs **C-6**

server

and volume settings, backing up **9-5**

name

changes **2-14**discovering **1-7**options **2-13**registration, via Web Management Interface **9-14**status **2-11**Shares **7-6**backing up configuration **9-5**delete **7-10**edit properties **7-9**shutdown **9-2**Simple Network Management Protocol, see *SNMP*Single-subnet configuration **3-6**site map **2-11, 10-1**server links **2-13**SMB **3-8**Snap EDR **B-2**Snap Finder **10-6**SnapCLI **E-1**running scripts **E-9**syntax **E-1**SnapDRImage **9-5**SnapExpansion **6-32**SnapExtensions **10-5**SnapServer Manager **1-4**

SnapServers

backup and restore path **B-4**configuring email notification of server events **9-15**connecting to **1-7**setting e-mail alerts for **9-16**snapshot **6-2**access **C-4**autobackup of volume settings **9-5**combined pools **5-10**coordinating with backup jobs **6-5**estimating storage requirements for **6-6**excluding from antivirus scans **11-3**excluding iSCSI Disks from **6-17**shares **C-5**ways to adjust pool size **6-6**SNMP configuration **3-23**software update **PR-iii, PR-v**specifications, GuardianOS **1-1**speed/duplex options **3-6**SSH enable **B-4**standalone **3-6**

storage

guides **5-2**pools **4-2**and incompatible disk drives **4-2**creating initial **2-6**RAID Sets screen **5-4**Volumes screen **5-14**

storage pools

and expansion units **4-6**support **9-14**system monitor **8-1**system reset **D-3**system/status LED **D-3****T**tape monitor **8-5**

TCP/IP

configuring **3-5**initial configuration **2-4**options **3-3**technical support **PR-iii**terminology for AFP **3-15**Tivoli NetView **3-23**tools **9-15**Traditional RAID **5-1, 7-1**adding disk drives **5-13**and expansion units **5-10**compared to DynamicRAID **A-3**quotas **5-19**RAID sets **5-2**shares **5-2**storage guides **5-2**volumes **5-2**troubleshooting **D-1**typographical conventions **PR-iv**

U

UID **7-3**

Uninterruptable Power Supplies (UPS) **2-18**

unlock a user password **7-20**

updates to GuardianOS **9-13**

UPS

configuring **2-18**

enabling support for **2-18**

low-power warning **2-18**

users

creating local **7-17**

file-level access for **C-6**

setup **7-4**

Shares **7-6**

write cache **5-17, 6-18**

V

volumes

and DynamicRAID **4-10**

and Traditional RAID **5-14**

backing up configuration **9-5**

capacity reached alert **9-16**

dynamic **A-4**

effect of deleting on antivirus software **11-1**

expanding capacity of **5-17**

Properties screen **5-17**

quotas **A-4**

screen **5-14**

size limits **A-4**

W

Wake-on-LAN Support **1-9**

Web Management Interface, overview **2-11**

Web Root **3-25**

Web Server **3-25**

Web View **3-27**

Windows

connecting from a client **3-10**

enabling guest account access **3-11, 3-13**

guest account access **3-10**

issues with PDC **D-5**

name resolution server support **3-9**

networking (SMB) **3-8**

security, joining

active directory domain **3-12**

see also *Active Directory*

see also *Authentication*

Windows Active Directory **3-9**