



Overland
Storage

SnapServer®

Administrator's Guide

For Appliances Running
GuardianOS® Version 7.6



January 2015
10400589-003



©2008-15 Overland Storage, Inc. All rights reserved.

Overland®, Overland Storage®, ARCVault®, DynamicRAID®, GuardianOS®, NEO®, NEO Series®, PowerLoader®, Protection OS®, RAINcloud®, REO®, REO 4000®, REO Series®, Snap Appliance®, Snap Care® (EU only), SnapSAN®, SnapScale®, SnapScale X2®, SnapServer®, StorAssure®, Ultamus®, VR2®, and XchangeNOW® are registered trademarks of Overland Storage, Inc.

Tandberg Data®, AccuGuard®, AccuVault®, DPS1000 Series®, DPS1100®, DPS1200®, DPS2000®, Magnum®, QuikStation®, QuikStor®, RDX®, RDXPRO®, StorageLibrary®, StorageLoader®, Tandberg SecureService®, Tandberg StorageLibrary®, and VXA® are registered trademarks of Tandberg Data, Inc.

Desktop Cloud Orchestrator® and V3® are registered trademarks of Sphere 3D, Inc.

RapidRebuild™, SnapExpansion XSR™, SnapScale X4™, SnapServer DX Series™, SnapServer XSD Series™, SnapServer XSD 40™, SnapServer XSR Series™, SnapServer XSR 40™, SnapServer XSR 120™, and SnapServer Manager™ are trademarks of Overland Storage, Inc.

BizNAS™, QuadPak™, and RDX+™ are trademarks of Tandberg Data, 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 7.6.122

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 a SnapServer appliance running GuardianOS version 7.6 on their network. It provides information on the installation, configuration, security, and maintenance of the SnapServer appliance and expansion 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 7.6 Web Management Interface
- SnapServer Manager (SSM)

GuardianOS version 7.6 comes preinstalled on all new SnapServer appliances. It can also be upgraded from a previously installed version of GuardianOS version 7.5 or later.

Product Documentation & Software Updates

SnapServer product documentation and additional literature are available online, along with the latest release of the GuardianOS version 7.6.

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, email our technical support staff at:

techsupport@overlandstorage.com.

You can get additional technical support information on the [Contact Support](#) web page at:

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

For a complete list of support times based on your type of coverage, visit our website at:

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

Japanese Voluntary Control Council for Interference (VCCI)

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI- A

(Translation: This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case, the user may be required to take corrective actions.)

Conventions

This document exercises several alerts and typographical conventions.

Convention	Description & Usage
 WARNING	A <i>Warning</i> contains information concerning personal safety. Failure to follow directions in the Warning could result in bodily harm or death.
WARNUNG	Eine <i>Warnung</i> enthält Informationen zur persönlichen Sicherheit. Das Nichtbeachten der Anweisungen in der Warnung kann zu Verletzungen oder zum Tod führen.
AVERTISSEMENT	Un <i>avertissement</i> contient des informations relatives à la sécurité personnelle. Ignorer les instructions dans l'avertissement peut entraîner des lésions corporelles ou la mort.
 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.
 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.
Item_name	Words in this special boldface font indicate the names of buttons or page names 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 > User > Password indicates that you should click the Setup tab, then the User secondary tab, and finally the Password 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.

Contents

Preface

Conventions	4
-------------------	---

Chapter 1: Overview

GuardianOS Specifications	11
What's New in GuardianOS 7.6	13
Using SnapServer Manager with SnapServer	14
SnapServer Manager Installation	14

Chapter 2: Initial Setup and Configuration

Connecting for the First Time	15
Connect Using the Server Name	15
Connect Using SSM	16
Setup a New SnapServer (via Wizard)	17
Step 1 – Enter General Configuration Settings	18
Step 2 – TCP/IP Configuration	19
Step 3 – RAID Type Selection (DynamicRAID/Traditional RAID)	19
Step 4 – Configure Expansion Units	25
Step 5 – Setup Completion	28
Step 6 – Registration Page	29
Step 7 – Scheduling Data Protection Tasks	30
Web Management Interface	32
Alert Messages	34
Site Map	35
Contact, Hardware & Software Information	35

Chapter 3: SnapServer Settings

Server Name	37
Date/Time	38
Secure Shell	39
Disable SSH	40
Connect to the CLI using SSH	40
UPS Protection	41
Printing	44
Procedure to Configure the Printer	44
Procedure to Configure the Client	45
Adding the Network Printer to a Windows Client	45
To Add a Network Printer to a Mac OS X Client	45
To Add a Network Printer to a Linux Client	45
To Monitor Print Jobs Remotely	45

To Delete a Printer	46
---------------------------	----

Chapter 4: Network Settings

Network Information	48
TCP/IP Networking	50
Configuring Port Properties	52
TCP/IP Configuration Considerations	53
Creating a Bond	55
Deleting a Bond	57
Windows/SMB Networking	59
Support for Windows/SMB Networking	60
Support for Windows Network Authentication	61
Configure Windows/SMB Networking	62
Apple Networking (AFP)	65
AFP Configuration Considerations	66
Edit AFP Access	66
NFS Access	67
Assigning Share Access to NFS Users	68
Enable NFS Access to the Server	68
Configure NFSv4 Access	68
LDAP/NIS	70
LDAP vs. NIS Overview	71
Configuring LDAP	71
Configuring NIS	72
FTP/FTPS Access	73
Supported FTP Clients	73
To Configure FTP/FTPS Access	73
To Connect via FTP/FTPS	74
SNMP Configuration	75
Default Traps	75
Supported Network Manager Applications and MIBs	76
Configure SNMP	76
Web Access	77
Configuring HTTP/HTTPS	78
Using Web Root to Configure the SnapServer as a Simple Web Server	78
iSNS Configuration	81

Chapter 5: DynamicRAID Storage

Storage Pools	83
Storage Pool Creation	84
Storage Pool Properties	89
View Disks from Storage Pool Properties Page	91
Storage Pool Deletion	93
Parity Management	94
Volumes	96
Volume Creation	97
Volume Properties	98
Volume Deletion	100

Chapter 6: Traditional RAID Storage

Storage Guides	103
Factors in Choosing a RAID Type	103
Local and Global Spares	104
RAID Sets	105
Create RAID Sets	105
Group RAID Sets	109
Change RAID Settings	113
Manage Global Spares	114
Edit RAID Set Properties	115
Volumes	118
Volumes and the Snapshot Pool	118
Volume Creation	119
Volume Properties	122
Quotas	126
Quotas Page	126
Enable/Disable Quotas	128
Add Quotas Wizard	129
Displaying/Changing Quotas	131

Chapter 7: Other Storage Options

Snapshots	135
Creating Snapshots	136
Schedule Snapshots	139
Snapshot Space	140
Snapshot Properties	142
iSCSI Disks	145
Configuring iSCSI Initiators	146
iSCSI Configuration on the SnapServer	146
Create iSCSI Disks	149
Edit an iSCSI Disk	152
Delete an iSCSI Disk	153
Configuring VSS/VDS for iSCSI Disks	153
Disks	156
Replacing Disk Drives	158
Adding Disk Drives	160
Managing Expansion Unit Storage	165
Integrating Orphaned Expansion Units	166
RDX QuikStor	166
RDX Media Properties	168
Copy Data To/From RDX Media	169
Format RDX Media	170
Eject RDX Media	171
Rename RDX Media Volume	172

Chapter 8: Security Options

Security Considerations	174
Guidelines for Local Authentication	174
User and Group ID Assignments	175
Security Guides	175

Security Guide for Windows Active Directory	176
Security Guide for Entire Volume Access	177
Security Guide for Folder Access on Volume	178
Shares	179
Create Shares	180
Edit Share Properties	182
Delete Shares	185
Configuring Share Access	185
Local Users	193
Create a User	194
Edit User Properties	195
Local User Password Policies	196
Assign User to Group	198
Delete Local User	199
Local Groups	200
Create New Group	201
Edit Group Properties	202
Specify Users in Group	203
Delete Group	204
Security Models	205
Managing Volume Security Models	205
Managing Folder Security Models in Traditional RAID	206
ID Mapping	209
Add Mapping	209
Change Mapping	212
Auto Mapping	215
Remove Mappings	216
Remove Missing ID Mappings	219
Filesystem Updates	221
Home Directories	222
Configure Home Directories	223

Chapter 9: System Monitoring

System Status	226
Active Users	227
Open Files	228
Network Monitor	229
Event Log	233
Filter the Log	233
Tape	234

Chapter 10: Maintenance

Shutdown and Restart	236
Manually Powering SnapServer On and Off	236
Factory Defaults	237
Disaster Recovery	238
Backing Up Server and Volume Settings	239
SnapDRImage File and Volume-Specific Files	239
System Settings Recovery	241
Volume and Storage Pool Security Settings Recovery	242
Replacing or Cloning a Server	244

Data Import	244
Setting Up a Data Import Job	246
Stopping an Import Job	248
Recreating an Import Job	248
Preserving Permissions	249
OS Update	250
Update the GuardianOS	250
Update Notification Option	251
Last OS Update	254
Support	254
Registering Your Server	255
Maintenance Tools	257
Email Notification	258
Host File Editor	260
Checking Filesystems	261

Chapter 11: Misc. Options

Home Pages	265
Home Page	265
Administration Page	267
SnapExtensions	269
BitTorrent Sync	269
Snap EDR	272
Snap Finder	273
Edit Snap Finder Properties	275
Change Password	276
Changing Your Password	276
Management Interface Settings	277

Appendix A: Backup Solutions

Backup and Replication Solutions	279
Snap Enterprise Data Replicator	279
Snap EDR Usage	280
Configuring Snap EDR	280
Scheduling Jobs in Snap EDR	281
Backup via SMB, AFP, or NFS	281
Off-the-Shelf Backup Solutions	281
iSCSI Disk Backups	281
Using Backup Exec for VSS-based Snapshots of SnapServer iSCSI Disks	281

Appendix B: Security and Access

Security Model Rules	283
Security Model Directories	284
Security Model Management	285
Special Share Options	285
Hiding Shares	285
Share Level Permissions	286
Where to Place Shares	286
File and Share Access	286
NFS Share Access	286

Snapshot Access	286
Snapshot Shares and On Demand File Recovery	287
Creating a Snapshot Share	287
File-level Security	288
Security Personalities and Security Models	288
Windows ACLs	288

Appendix C: DynamicRAID Overview

About DynamicRAID	290
Should I use DynamicRAID or Traditional RAID?	291
Setting Up DynamicRAID	293
DynamicRAID Implementation	293
Storage Expansion	293
Snapshots	294
iSCSI Target Volumes	294
Indicators	294
Additional Information on DynamicRAID Sizing	294

Appendix D: GuardianOS Ports

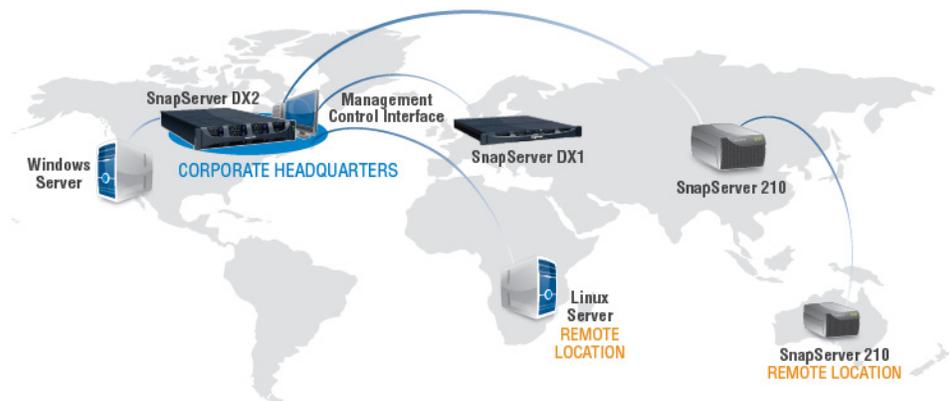
Appendix E: Command Line Interface

SnapCLI Syntax	299
SnapCLI Procedures	301
Scripts in SnapCLI	301
Running a SnapCLI Script	301
Sample Script	302

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 (GOS), 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.



Topics in Overview:

- [GuardianOS Specifications](#)
- [What's New in GuardianOS 7.6](#)
- [Using SnapServer Manager with SnapServer](#)

GuardianOS Specifications

These specifications apply to SnapServers running GuardianOS version 7.6:

Feature	Specification
Network Transport Protocols	<ul style="list-style-type: none"> • TCP/IP (Transmission Control Protocol/Internet Protocol) • UDP/IP (User Datagram Protocol/Internet Protocol)
Network Block Protocols	<ul style="list-style-type: none"> • iSCSI (Internet Small Computer System Interface)

Feature	Specification
Network File Protocols	<ul style="list-style-type: none"> • Microsoft Networking (CIFS/SMB1/SMB2) • 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/explicit FTPS such as FTPES or Auth TLS) <p>* AFP v 3.2 ACLs and extended attributes not supported.</p>
Network Security	<ul style="list-style-type: none"> • Microsoft Active Directory Service (ADS) (member server) • Unix Network Information Service (NIS) user/group UID/GID translation • LDAP user/group UID/GID translation • 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)
Network Client Types	<ul style="list-style-type: none"> • Microsoft Windows 2003/2003 R2/2008 SP2/2008 R2 /XP SP3/ Vista SP2/7/8/2012 • Mac OS X 10.5/10.6/10.7/10.8/10.9 • 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
Data Protection	<ul style="list-style-type: none"> • Snapshots for immediate or scheduled point-in-time images of the server 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
RAID Options with Traditional RAID Configuration	<ul style="list-style-type: none"> • RAID 0 (drive striping). • RAID 1 (drive mirroring). • RAID 5 (drive striping with parity) - Available only on systems with four (4) or more drives. • RAID 6 (drive striping with two parity drives) - Available only on systems with four (4) or more drives. • RAID 10 (striped mirroring) - Available only on systems with four (4) or more drives. • Global or local spare support. • Instant Capacity Expansion (ICE).
DHCP Support	<ul style="list-style-type: none"> • Supports Dynamic Host Configuration Protocol (DHCP) for automatic assignment of IP addresses

Feature	Specification
System Management	<ul style="list-style-type: none"> • Browser-based administration tool called the Web Management Interface • SnapServer Manager utility (platform independent) • SnapCLI for volume system deployment • 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)

What's New in GuardianOS 7.6

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

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

Feature	New Functionality
RDX QuikStor	<p>Support for RDX QuikStor devices has been added including media management. The RDX QuikStor can be used to manually copy files for off-site data redundancy and to transfer files between SnapServers and other servers or clients without using the network.</p> <p>Now supports eject locking/unlocking from the Web Management Interface.</p>
New Hardware Support	<p>Full support has been added for:</p> <ul style="list-style-type: none"> • Four-drive bay desktop SnapServer (XSD 40) • Four-drive bay 1U rack-mounted SnapServer (XSR 40) • 12-drive bay 2U rack-mounted SnapServer (XSR 120) • 12-drive bay 2U rack-mounted expansion unit (SnapExpansion XSR).
Improved Security	<p>For increased security, the Web Management Interface no longer supports or allows the use of the SSLv3 protocol when communicating via HTTPS. Only TLSv1.0, TLSv1.1, and TLSv1.2 protocols are supported.</p>
BitTorrent Sync	<p>BitTorrent Sync can be used to replicate data between the SnapServer and other servers or workstations.</p>
Network Monitor	<p>This feature allows for monitoring the network (bandwidth) usage of the server across all network interfaces. Network usage can be monitored both in real-time and historically going back several weeks or months.</p>

Using SnapServer Manager with SnapServer

SnapServer Manager (SSM) is a Java-based application that runs on all major client systems. SSM provides a single screen from which administrators can discover all SnapServer servers, REO appliances, SnapSAN arrays, SnapScale clusters, and SnapScale Uninitialized nodes (that is, nodes that are not part of a SnapScale cluster) on their network.

The screenshot displays the SnapServer Manager application window titled "SnapServer Manager - <All Servers>". The interface includes a menu bar (File, Edit, Administration, Help), a left sidebar for "Servers & Groups" showing a tree view with "All Servers" selected, and a main table of server information. The table has columns for Server, Status, IP Address, Model, OS Version, and Storage Usage. A "Server Group" box highlights the sidebar, a "Server List" box highlights the main table, and a "Status Bar" box highlights the bottom status area. The status bar shows "No operations are currently running or scheduled." and "Online/Offline: 115/18".

Server	Status	IP Address	Model	OS Version	Storage Usage
athos	OK	192.168.192.143 (D...	4400	GOS 6.0.043	8% used
beryl	OK	192.168.193.248 (D...	DX1	GOS 7.0.125	5% used
birthstone	OK	192.168.193.191 (D...	DX1	GOS 7.0.106	0% used
blähfly	OK	192.168.193.147 (D...	N2000	GOS 6.5.026	0% used
bmgala1	OK	192.168.192.146 (D...	VirtualSnap	GOS 7.0.125-kdb	23% used
bmgala2	OK	192.168.192.215 (D...	VirtualSnap	GOS 7.1.0-briansled11	18% used
bobbert	OK	192.168.192.140 (D...	4400	GOS 5.0.133	36% used
CCDragonflyDR	OK	192.168.193.199 (D...	N2000	GOS 6.5.023	34% used
CCEMERALDGRN	OK	192.168.193.138 (D...	DX1	GOS 7.1.007	0% used
CCemeraldRED	OK	192.168.192.144 (D...	DX1	GOS 7.1.007	11% used
CCGalapagos-Q	OK		VirtualSnap	GOS 7.0.089	73% used
CCGalapagos6	OK		VirtualSnap	GOS 7.0.085	0% used
CCStorm520-1	OK		520	GOS 6.5.028	0% used
CCStorm520-3	OK	192.168.193.222 (D...	520	GOS 6.5.028	0% used
CCStorm620-A	OK	192.168.193.227 (D...	520	GOS 6.5.027	5% used
CCStorm650-3	OK	192.168.193.175 (D...	650	GOS 5.2.067	0% used
CCSunDragon	OK	192.168.193.247 (D...	DX2	GOS 7.1.008	?
CCWave210	OK	192.168.192.142 (D...	210	GOS 6.5.023	23% used
CCWAVE410	OK	192.168.192.132 (D...	410	GOS 5.2.067	96% used
daedalus	OK	192.168.192.32 (Stat...	4500	GOS 5.0.133	74% used
elgringo	OK	192.168.192.180 (D...	14000	GOS 5.1.046	32% used
emerald-proto	OK	192.168.193.140 (D...	DX1	GOS 7.1.002	0% used
evomp	OK	192.168.192.171 (D...	18000	GOS 6.0.043	26% used
flis	OK	192.168.192.240 (D...	18000	GOS 5.2.056 SP1	55% used
flyspeck	OK	192.168.192.139 (D...	N2000	GOS 6.5.022	2% used
Glauring	OK	192.168.193.114 (D...	N2000	GOS 6.5.0.buildcrush...	0% used
JY-DSnap2300036	OK	192.168.193.189 (D...	DX1	GOS 7.1.006	1% used
JY-DX2-Boogaloo	Offline	192.168.192.222	?	?	?

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 several client platforms, including Windows, Mac OS X, and Linux.

Refer to the *SnapServer Manager User Guide* for details on discovering and configuring SnapServers.

Initial Setup and Configuration

This section covers the initial setup and configuration of a SnapServer appliance running GuardianOS 7.6.

NOTE: For information concerning the installation and wiring of your SnapServer hardware, refer to the appropriate Quick Start Guide for your product.

Topics in Setup and Configuration:

- [Connecting for the First Time](#)
- [Setup a New SnapServer \(via Wizard\)](#)
- [Web Management Interface](#)

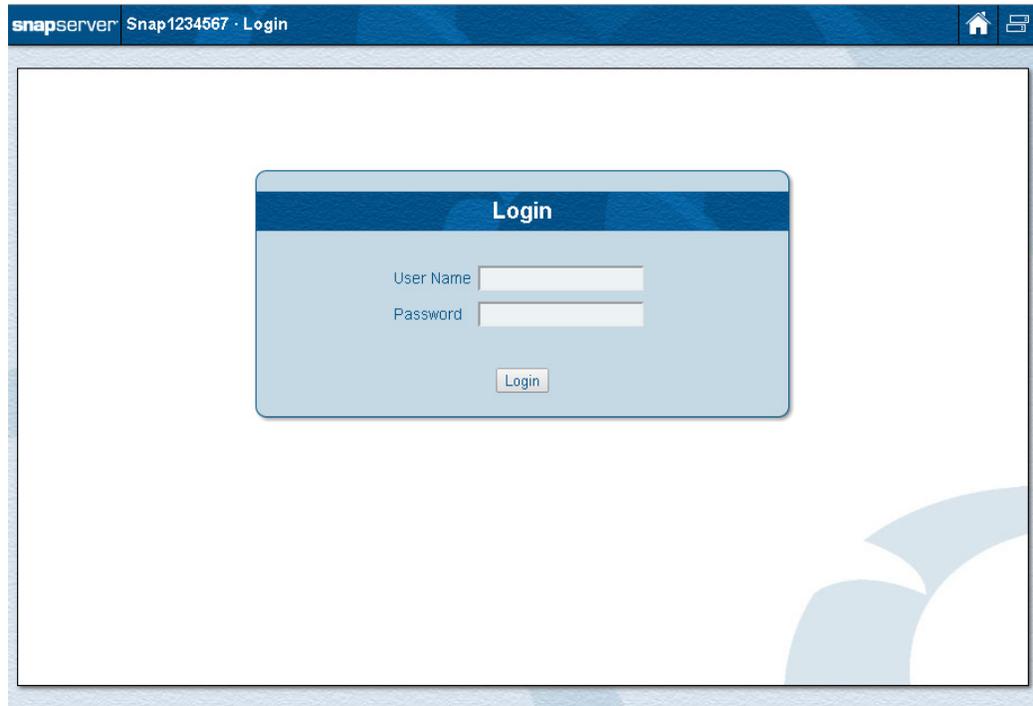
Connecting for the First Time

SnapServers are configured to acquire their IP address from a DHCP server. If no DHCP server is found on the network, the server defaults to an IP address in the range of 169.254.xxx.xxx and is labeled as “ZeroConf” in SnapServer Manager (SSM). You may not be able to see SnapServers on your network until you discover them using either the default server name or the SSM utility and optionally assign them an IP address.

Connect Using the Server Name

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

1. Find the **server name**.
A SnapServer name is of the format “Snapnnnnnnnnn,” where *nnnnnnnnn* is the server number. The number is a unique, numeric-only string that appears on a label affixed to the bottom of the unit.
2. In a web browser, enter the **URL** to connect to the server.
For example, enter “http://Snapnnnnnnnnn” (using the server name).
3. Press **Enter** to connect to the Web Management Interface.



4. In the login dialog box, enter **admin** as the user name and **admin** as the password (the system defaults), then click **Login**.
5. Complete the **Initial Setup Wizard** to setup your server.

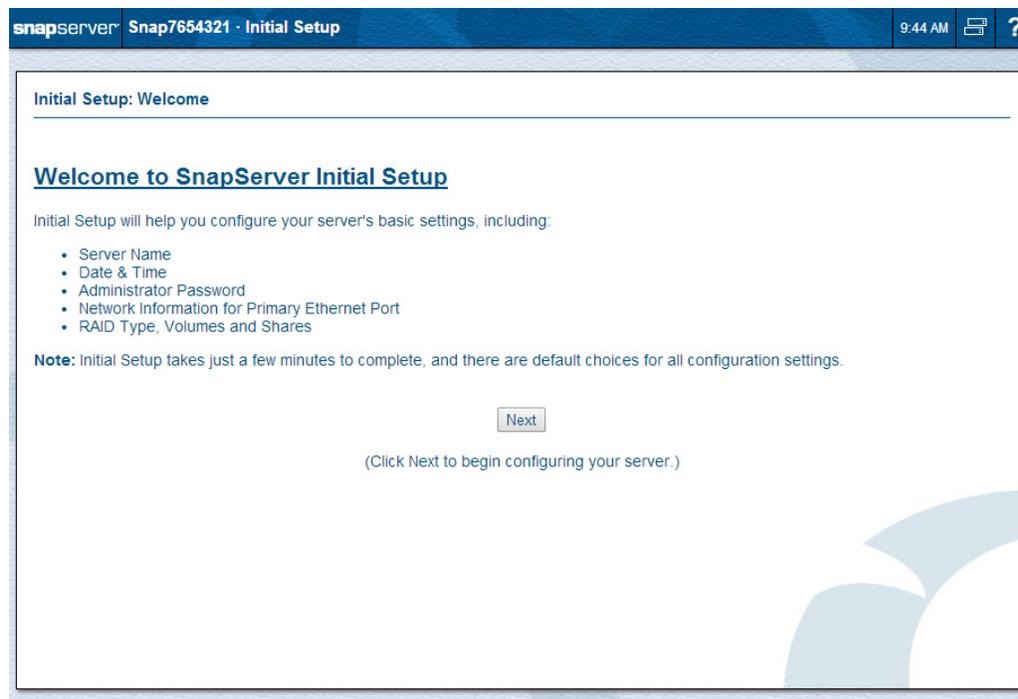
Connect Using SSM

1. Launch **SnapServer Manager (SSM)**.
SSM discovers all SnapServers, SnapScale clusters, and Uninitialized nodes on its local network segment and displays their Server names, IP addresses, and other 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 and SnapScale nodes, you may need to find their default names as explained in [Connect Using the Server Name on page 15](#).
2. If using a DHCP server, proceed to [Step 3](#); otherwise, assign an **IP address** to the 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 (the system defaults), then click **OK**.
5. Complete the **Initial Setup Wizard** to setup your server.

Setup a New SnapServer (via Wizard)

On a new SnapServer, once you log in to the Web Management Interface, the Initial Setup Wizard runs displaying the **Welcome** page:



The Initial Setup Wizard consists of several steps:

[Step 1 – Enter General Configuration Settings](#)

[Step 2 – TCP/IP Configuration](#)

[Step 3 – RAID Type Selection \(DynamicRAID/Traditional RAID\)](#)

[Step 4 – Configure Expansion Units](#)

[Step 5 – Setup Completion](#)

[Step 6 – Registration Page](#)

[Step 7 – Scheduling Data Protection Tasks](#)

Step 1 – Enter General Configuration Settings

Clicking **Next** on the **Welcome** page displays the **General Information** page of the wizard. This page allows you to change the basic information for the SnapServer. It is highly recommended for security that you set your Administrator password to something other than the default setting.

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](#) page above, you may be prompted to log in again before continuing the setup.

Step 2 – 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 Snap7654321 - Initial Setup', a clock showing '10:13 AM', 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 radio button are several input fields: 'IP Address' (10.25.3.37), 'Subnet Mask' (255.255.0.0), 'WINS Servers:' (three empty fields, each labeled '(optional)'), 'Default Gateway' (10.25.1.1, labeled '(optional)'), 'DNS Domain Name' (devnet.myoverland.nl, labeled '(optional)'), and 'Domain Name Servers:' (10.6.8.34 and 10.6.8.35, each labeled '(optional)'). A 'Next' button is located at the bottom right of the form area.

1. If you wish to assign a **static IP** instead of using DHCP, select the radio button for a static IP address and enter the IP address and subnet mask. If needed, the optional information should also be entered.
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.

Step 3 – RAID Type Selection (DynamicRAID/Traditional RAID)

GuardianOS 7.6 offers the 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 [Appendix C - DynamicRAID Overview](#).

1. Choose either **DynamicRAID** or **Traditional RAID**.

Initial Setup: Configure Storage

Select the type of storage environment that you would like to use on this system.
(**Important:** Once selected, the storage environment cannot be changed without first deleting all your data.)

DynamicRAID™ (Recommended)

DynamicRAID is a form of RAID developed for GuardianOS which provides simple and flexible storage provisioning. DynamicRAID offers the following benefits over Traditional RAID:

- Flexible Disk Parity - Switch between single parity and dual parity dynamically.
- Automatic RAID Creation - New RAID groups are created automatically.
- RAID Expansion - When disks are added to the system, the space is easily incorporated into your existing RAID's and your data remains accessible.
- Non-Disruptive Disk Upgrades - Migrate an existing RAID array to higher capacity disks without interrupting data access.
- Dynamic Volumes - The size of a volume is adjustable at anytime and can be limited to a specified level.

Traditional RAID

- You manually manage your storage.

Note: With DynamicRAID, the size of a volume is limited to the capacity available in a single physical chassis. If a volume must span across multiple units, or if you need the ability to stripe or mirror data (such as RAID 0, 1, or 10), then you should select Traditional RAID.

[Tell Me More](#) [Next](#)

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

Initial Setup: Configure Storage

Important: You have selected a DynamicRAID storage environment. Once selected, the storage environment cannot be changed without first performing a system reset which will delete all of your data.

Are you sure you want DynamicRAID storage?

[Back](#) [Yes, I Want DynamicRAID](#)

Initial Setup: Configure Storage

Important: You have selected a Traditional RAID storage environment. Once selected, the storage environment cannot be changed without first performing a system reset which will delete all of your data.

Are you sure you want Traditional RAID storage? (**Note:** The server will be restarted.)

[Back](#) [Yes, I Want Traditional RAID](#)

3. Click the **Yes** option to continue.

4. Available disks are detected and shown on the **Configure Storage - Detected Disks** page.

The screenshot shows the 'Initial Setup: Configure Storage - Detected Disks' page. The page title is 'Initial Setup: Configure Storage - Detected Disks'. Below the title, it says 'Your server's disk slots are fully populated. Click Next to view your storage options.' Below this, there is a table titled 'Detected Disks' with three columns: 'Location', 'Disks Detected', and 'Available Disk Slots'. The table contains one row for the 'Head Unit' with the following details: Model: VirtualSnap, Total disk capacity: 200 GB, Disks Detected: 4, and Available Disk Slots: 4. Below the table are two buttons: 'Re-Detect Disks' and 'Next'.

Location	Disks Detected	Available Disk Slots
Head Unit Model: VirtualSnap Total disk capacity: 200 GB	4	4

If SnapExpansion units are attached, their available disks are also detected and displayed:

The screenshot shows the 'Initial Setup: Configure Storage - Detected Disks' page. The page title is 'Initial Setup: Configure Storage - Detected Disks'. Below the title, it says 'Your server's disk slots are fully populated. Click Next to view your storage options.' Below this, there is a table titled 'Detected Disks' with three columns: 'Location', 'Disks Detected', and 'Available Disk Slots'. The table contains eight rows: one for the 'Head Unit' and seven for 'Expansion Units'. The details for each row are as follows:

Location	Disks Detected	Available Disk Slots
Head Unit Model: DX2 Total disk capacity: 21.83 TB	12	12
Expansion Unit 1 Model: SE DX Total disk capacity: 30.02 TB	11	12
Expansion Unit 2 Model: SE DX Total disk capacity: 12.74 TB	12	12
Expansion Unit 3 Model: SE DX Total disk capacity: 21.83 TB	12	12
Expansion Unit 4 Model: SE DX Total disk capacity: 6 TB	11	12
Expansion Unit 5 Model: SE DX Total disk capacity: 12.74 TB	12	12
Expansion Unit 6 Model: SE DX Total disk capacity: 22.74 TB	12	12
Expansion Unit 7 Model: SE DX Total disk capacity: 43.66 TB	12	12

Below the table are two buttons: 'Re-Detect Disks' and 'Next'.

NOTE: If you are using expansion units and DynamicRAID, the SnapServer first configures the head unit, then expansion units. For Traditional RAID, the expansion units must be configured separately via **Storage > RAID Sets**.

- If empty slots exist in either the head or expansion units, new drives can be physically added and then **Re-Detect Disks** clicked to add them to the configuration.

Continue configuring the SnapServer based on the **RAID mode** you selected:

- If you selected **DynamicRAID**, proceed to the [DynamicRAID Setup subsection below](#).
- 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 subsection below](#).

DynamicRAID Setup

The **Configure Storage - Head Unit** pages let you configure your head unit storage by choosing the parity mode and snapshot pool size under DynamicRAID.

Initial Setup: Configure Storage - Head Unit

Use the settings below to create a Storage Pool on your server.
(Note: A volume and share will be created automatically for this Storage Pool.)

Storage Pool configuration is based on these settings:

Pool	Estimated Available Space	Percent of Storage
Data Pool	85.77 GB	80%
Snapshot Pool	21.44 GB	20%

Storage Pool Name:

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

If you plan on using snapshots, it is recommended that you reserve at least 20% of your Storage Pool for snapshots. You can adjust the snapshot pool percentage at a later time; however to increase it, you will first need to add more capacity to this Storage Pool.

Percentage of this Storage Pool to reserve for snapshots:

The parity mode lets you set the data pool size. The options presented are based on the number and available space of disk drives detected in the previous step. See [Parity Management on page 94](#) for full details.

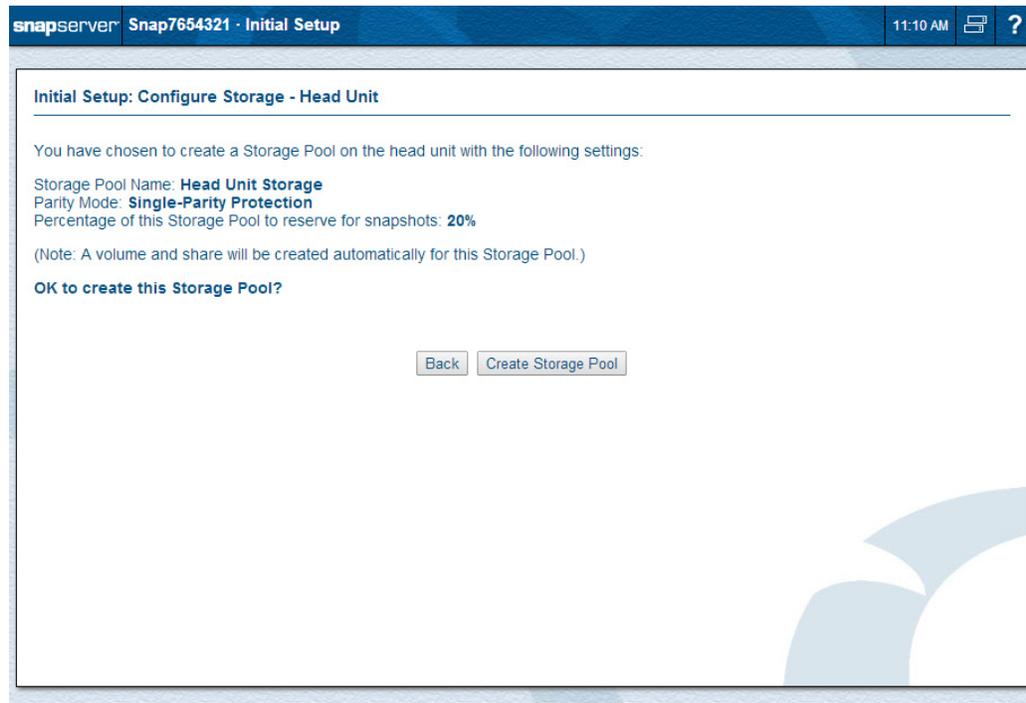
In addition, snapshot space can be reserved using the following guidelines:

- 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 on page 135](#) for more information.

Configure DynamicRAID Storage:

- Select the **parity mode** from the options provided.
- Use the drop-down list to choose the size of the **snapshot pool**.
- After you have made your selections, click **Next**.

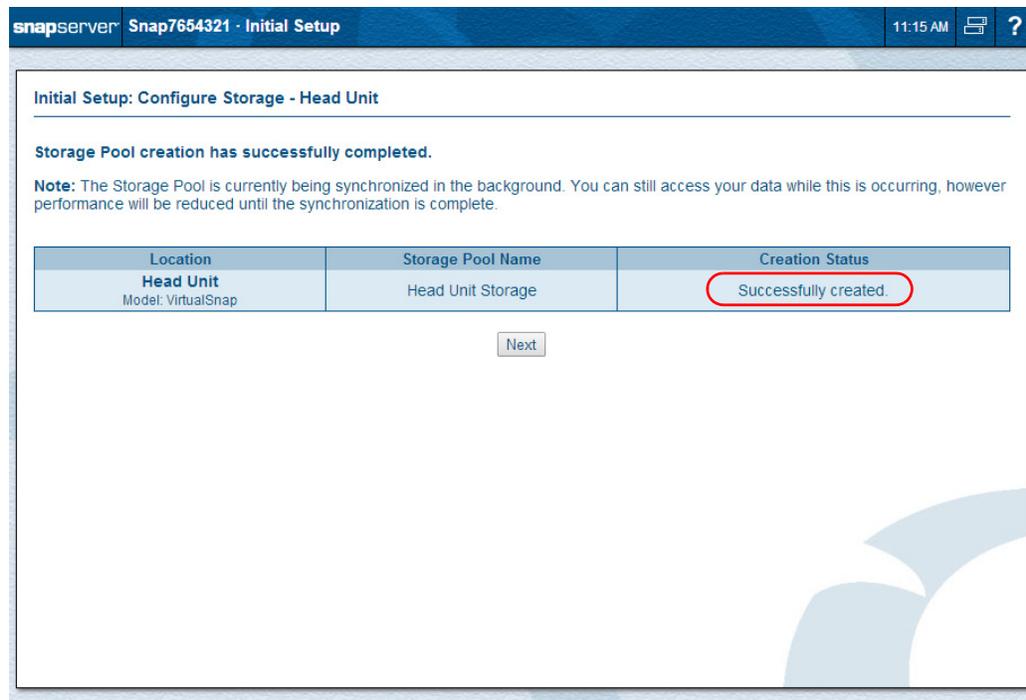
- At the confirmation page, click **Create Storage Pool**.



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.

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



Next steps:

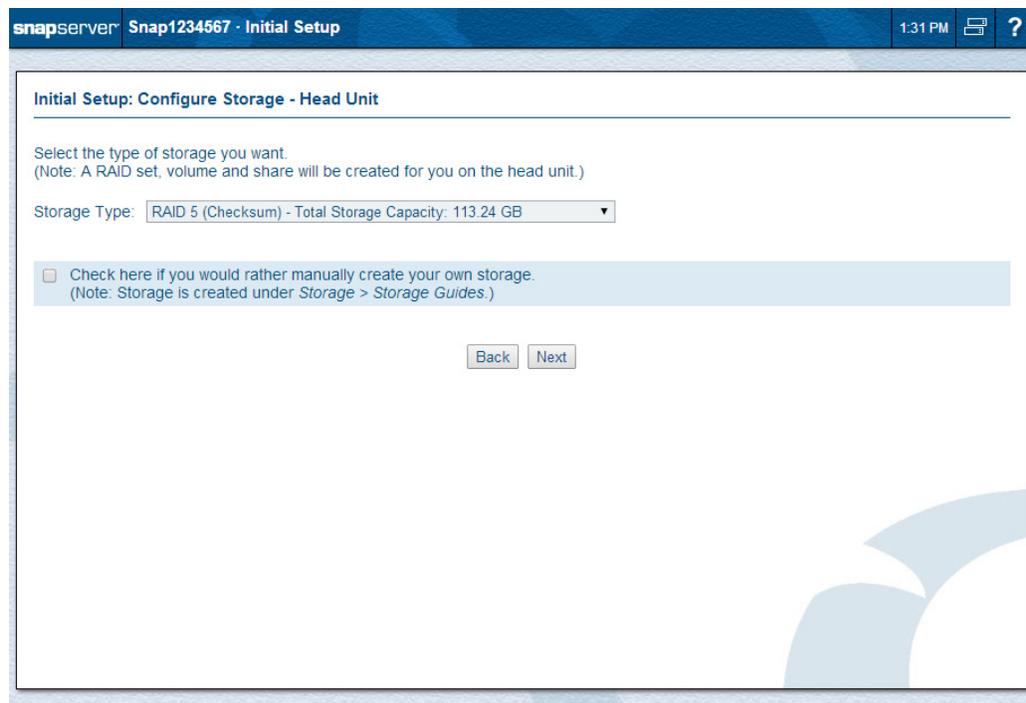
- If you are using SnapExpansion units, their configuration automatically starts next. Continue with [Step 4 – Configure Expansion Units](#).
- Otherwise, continue with [Step 5 – Setup Completion](#).

Traditional RAID Setup

Once Traditional RAID is selected, the wizard continues with the setup of the head unit.

NOTE: In Traditional RAID, only the head unit is configured during the setup process. If you are using expansion units, they are configured manually after the setup wizard is complete ([Storage > RAID Sets](#)).

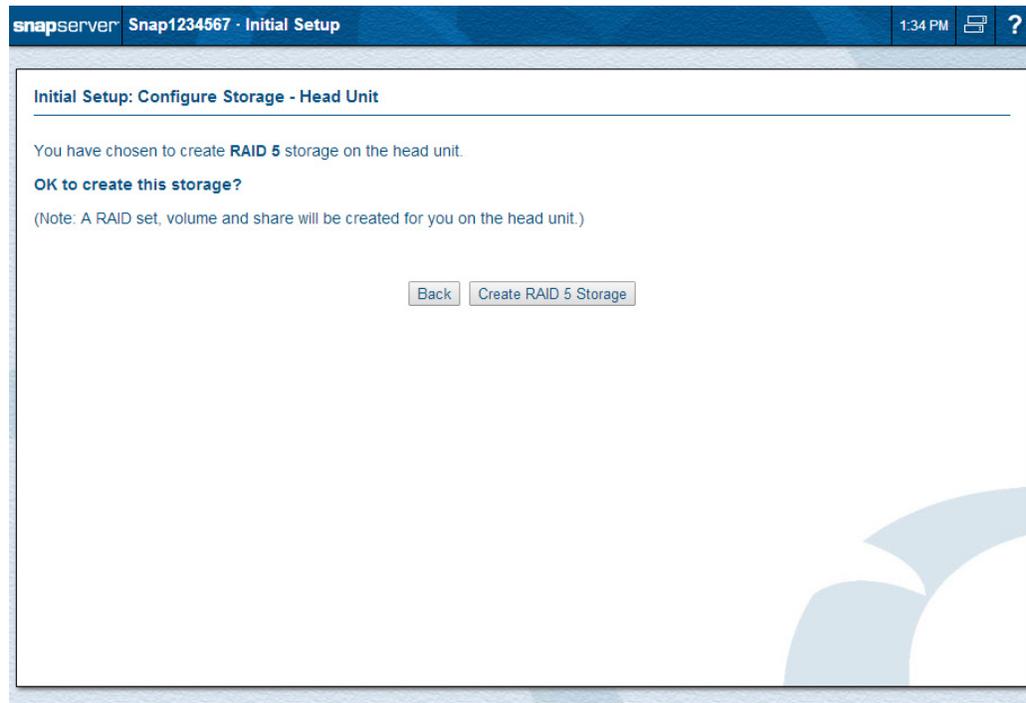
The next head unit configuration page lets you manually configure your head unit storage by selecting the RAID type based on the installed disk drives.



1. Choose **one** of the following:
 - From the drop-down list, select the predefined **storage type** (RAID) available based on units and drives installed.
 - Check the **Check here if you would rather manually create your own storage** box.
2. 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 storage setup process.

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



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

Step 4 – Configure Expansion Units

If you are using expansion units with DynamicRAID, the SnapServer will recognize them during the setup process and you can configure storage pools on them. After the setup is complete, expansion units can be managed via **Storage > Storage Pools**.

NOTE: If you are using Traditional RAID, expansion units can only be configured manually after the setup process is complete (using **Storage > RAID Sets**).

1. During the setup of a DynamicRAID system, after the head unit is configured, you are prompted to create the storage pools on the expansion units. Select the **parity mode** and **snapshot pool size** for each expansion unit.

snapserver Snap7654321 · Initial Setup 9:11 AM ?

Initial Setup: Configure Storage - Expansion Units

Use the settings below to create a Storage Pool on each of your expansion units.
(Note: A Storage Pool, volume and share will be created for each expansion unit.)

Storage Pools for all units: (Note: Units with an existing Storage Pool are displayed in *italics*.)

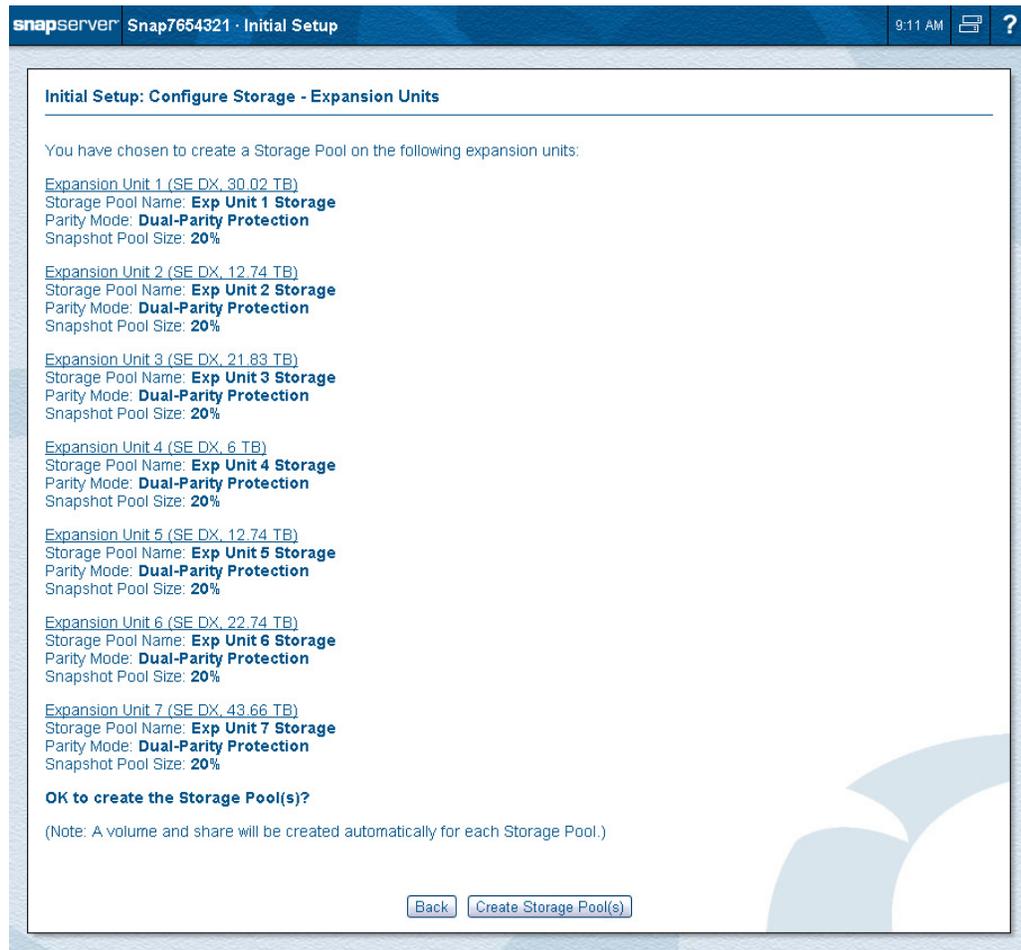
Location	Storage Pool Name	Parity Mode	Snapshot Pool Size
Head Unit Model: DX2 Total disk capacity: 21.83 TB	<i>Head Unit Storage</i>	<i>Dual</i>	<i>20%</i>
Expansion Unit 1 Model: SE DX Total disk capacity: 30.02 TB	Exp Unit 1 Storage	Dual	20%
Expansion Unit 2 Model: SE DX Total disk capacity: 12.74 TB	Exp Unit 2 Storage	Dual	20%
Expansion Unit 3 Model: SE DX Total disk capacity: 21.83 TB	Exp Unit 3 Storage	Dual	20%
Expansion Unit 4 Model: SE DX Total disk capacity: 6 TB	Exp Unit 4 Storage	Dual	20%
Expansion Unit 5 Model: SE DX Total disk capacity: 12.74 TB	Exp Unit 5 Storage	Dual	20%
Expansion Unit 6 Model: SE DX Total disk capacity: 22.74 TB	Exp Unit 6 Storage	Dual	20%
Expansion Unit 7 Model: SE DX Total disk capacity: 43.66 TB	Exp Unit 7 Storage	Dual	20%

Back Next

(Click Next to review and confirm your expansion unit settings.)

2. Click **Next**.

- Review your storage pool configuration. When you are done, click **Create Storage Pools** to create the storage pools on the expansion units (or **Back** to make changes).

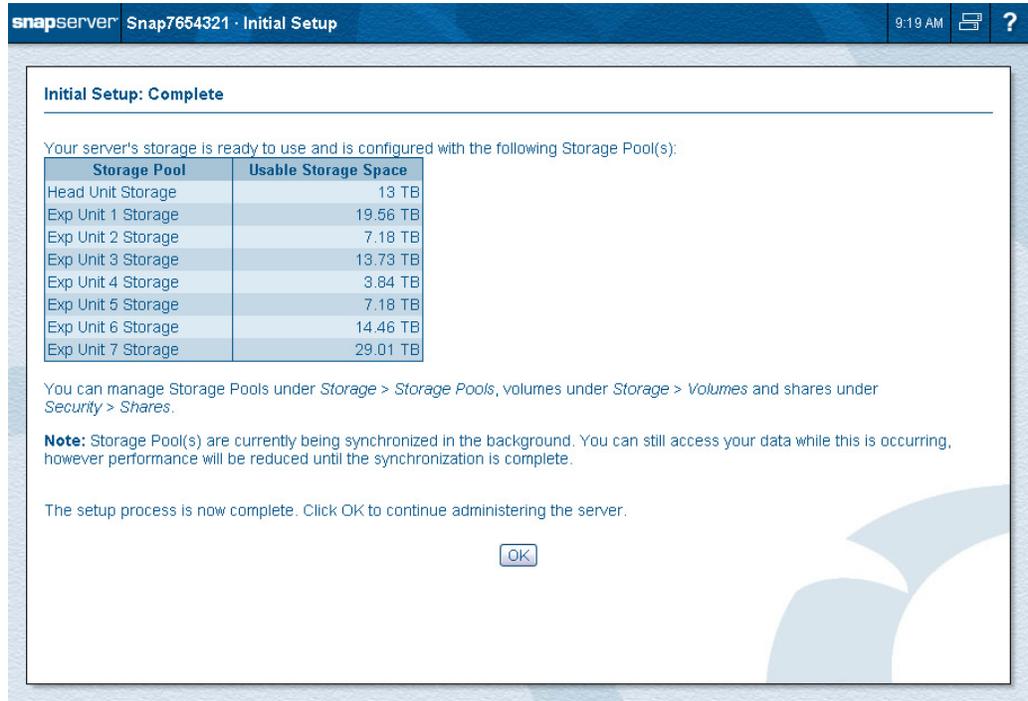


A screen appears showing the creation of the storage pools on the expansion units. The creation of the storage pools may take several minutes.

- When the expansion units' storage pools have been created, click **Next** to continue with the completion steps.

Step 5 – Setup Completion

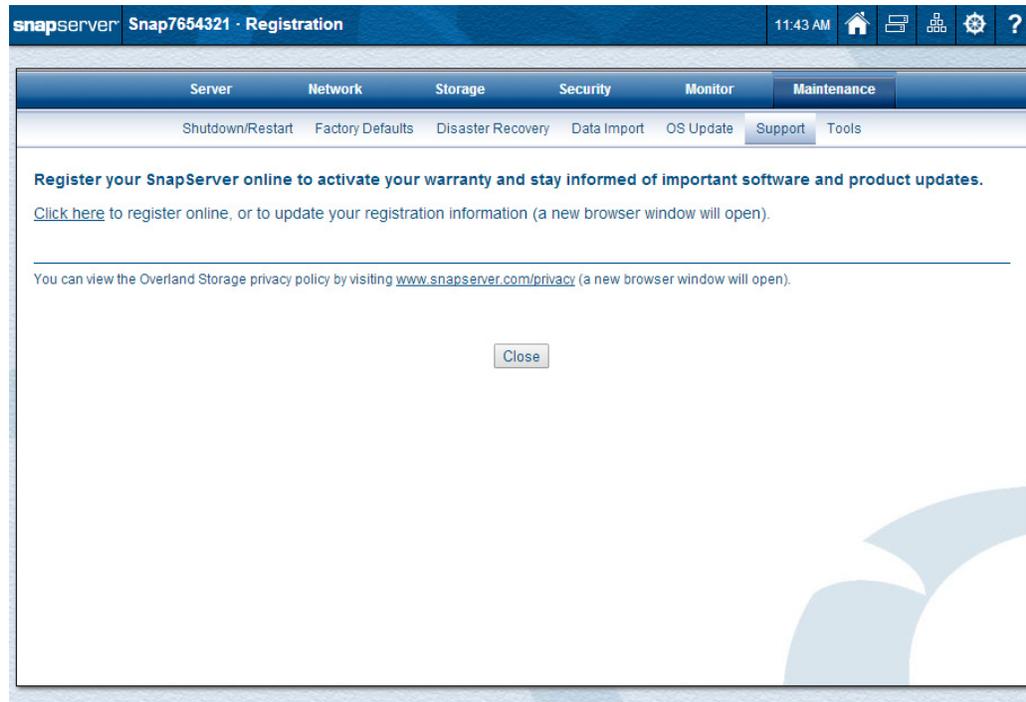
Once the head unit and all the DynamicRAID expansion units are configured, the **Complete** page is shown:



1. Click **OK** to finish the setup process.
Storage Pools will be synchronizing in the background. You can still access your data while this is occurring; however, performance will be reduced until the synchronization is complete.
2. If you have changed the **server name**, you will be prompted to **restart**. Click **Restart** to continue.
The server will restart and your browser will automatically reconnect to the server. **Log in** again when prompted to do so.
3. You are prompted to **register** your system.
It is important that you register your system to activate the warranty coverage. Proceed to [Step 6 – Registration Page](#).

Step 6 – Registration Page

After the setup wizard is done, you will see a **Registration** page where you can register your SnapServer. This page can also be accessed by clicking **Maintenance > Support > Registration**.



NOTE: 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.

Click the **Click Here** link to launch the Overland Storage Support website and register online.

1. At the [Site Login](#), enter your **e-mail address** and **password**, and click **GO**.
If you are not yet a member, follow the **New Member** link to get set up.

E-mail:

Password:

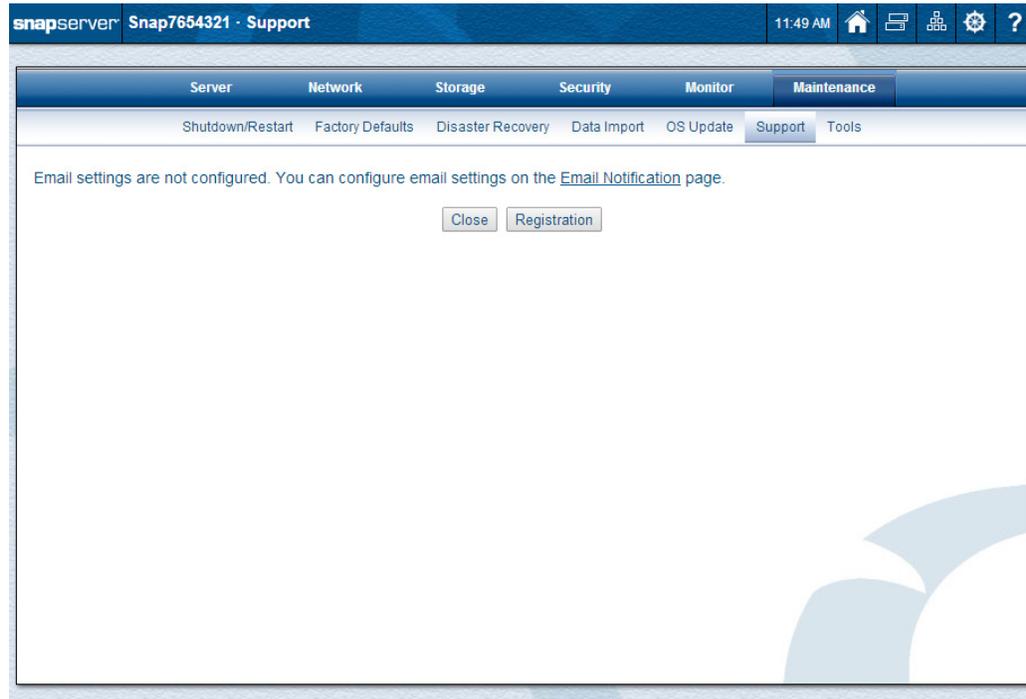
Remember Me

GO >

[Forgot your password?](#) [New member?](#)

2. At the **Confirm Automated Product Registration** page, enter the **date**, **reseller**, and **product site**.
3. Click **Confirm** to complete the process.

Once registration is complete, click **Close** on the **Registration** page to finish your setup. A reminder to configure your **Email Notification** page is displayed:



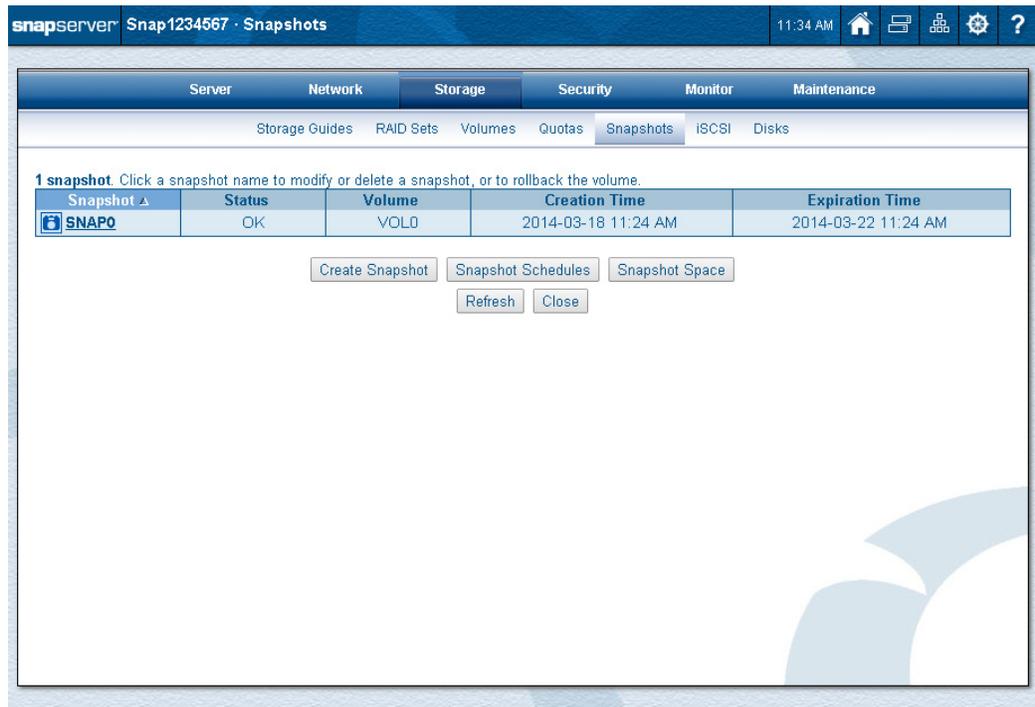
Click the link to go to **Maintenance > Tools > Email Notification** and complete the settings. See [Email Notification](#) on [page 220](#) for details.

Step 7 – Scheduling Data Protection Tasks

Scheduling backups, snapshots, 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.

NOTE: It is recommended completing these tasks before continuing with your configuration.

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:

- For DynamicRAID mode, go to **Storage > Storage Pools** and click the storage pool name.
- For Traditional RAID mode, go to **Storage > Snapshots > Snapshot Space**.

Create a disaster recovery image on the **Maintenance > Disaster Recovery** page. This image should be created after the server configuration is complete and it can be used to recover the server or bring a replacement server to the configured state. See [Disaster Recovery on page 238](#) 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 [Backup and Replication Solutions on page 279](#).

Web Management Interface

snapScale Scale7846949 · Administration

SnapScale Network Storage Security Monitor Maintenance

SnapScale Name: Scale7846949
RAINcloudOS Version: 4.1.091
Uptime: 5:22:03 (D:H:M)
Data Replication Count: 2x
Spare Disks Setting: 2
UPS Support: Enabled
Email Notification: Enabled
Management IP Address: 10.25.11.160
Multicast IP Address: 233.33.0.0

Peer Sets: 5
All peer sets OK.

Nodes: 3
All nodes OK.

Active Spare Disks: 2
All spares OK.

Protocol Manager
All nodes OK.

SnapScale Settings
All settings OK.

UPS Status
All nodes OK.

Total Storage Usage:
27% (18.9 GB / 70.82 GB)

Refresh Close

[Click here to find out what's new in RAINcloudOS 4.1 and this Web Management Interface.](#)

SnapServer appliances use a web-based graphical user interface (GUI), called the Web Management Interface, to administer and monitor the server. It supports most common web browsers. JavaScript must be enabled in the browser for it to work.

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, three options below the shares list, and has special navigation buttons displayed on the right side of the title bar (see the next table).

snapserver Snap7654321 · Home 9:35 AM

Share Name	Description
SHARE1	

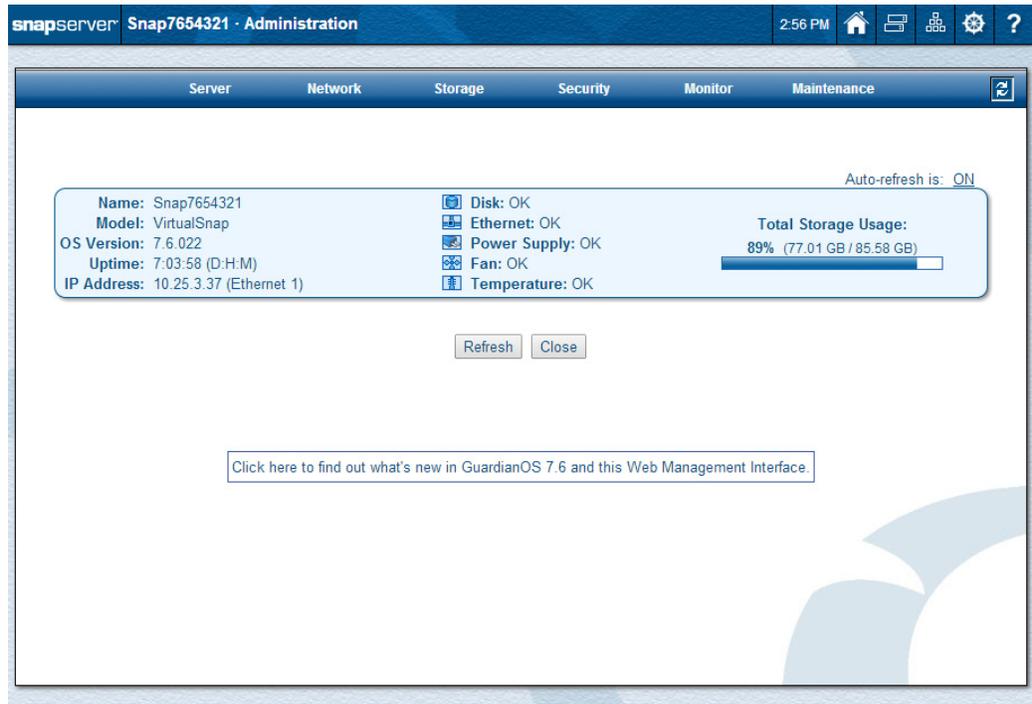
Change Password Switch User (Logout) Administration

NOTE: If you have not 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 this icon to access the password change page. Passwords are case sensitive. Use up to 15 alphanumeric characters.
Switch User 	Click this icon to log out and open the login dialog box to log in as a different user.
Administration 	Click this icon to administer the server. If you are not yet logged in, you are prompted to do so.
Navigation Buttons:	<p>The following navigation buttons are present in the upper right on every Web Management Interface page:</p>
	Home – Click this icon to switch between the Home page and the Admin Home page. If you have not yet logged in to the Admin Home page, only the Home page is available.
	Snap Finder – Click this icon to view a list of all SnapServers, SnapScale clusters, and Uninitialized nodes on your network, and to specify a list of remote servers that can access these servers, clusters, and nodes on other subnets. You can access these servers, clusters, and nodes 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 icon 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 icon to access the web online help for the Web Management Interface 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.

When logged in to the **Administration** page, details about the server's health are shown:



The same icons are available at the top of this page plus an additional refresh icon (🔄) for auto-refreshing pages is located on the tab bar. For more information, see the [Home Page](#) section.

Alert Messages

Alert messages are displayed on Administrator-level Web Management Interface pages that display a menu. Some alerts have clickable options:

- **[Later]** - Hides the alert for 24 hours or until after feature is run, whichever is first.
- **[Hide]** - Suppresses the alert. It will not be shown again until after the feature called out in the alert is run and a new alert for that feature is generated.

Site Map

The GuardianOS site map (⚙️) provides links to all the web pages that make up the Web Management Interface. All the pages are each covered in detail in the following chapters.

snapserver						
Server	Network	Storage	Security	Monitor	Maintenance	Misc.
Server Name	Information	Storage Pools	Security Guides	System Status	Shutdown/Restart	Administration
Date/Time	TCP/IP	Volumes	Shares	Active Users	Factory Defaults	Home
SSH	Windows/SMB	> Create Volume	> Create Share	Open Files	Disaster Recovery	SnapExtensions
UPS	Apple/AFP	Snapshots	Local Users	Network Monitor	Data Import	Snap Finder
Printing	NFS	> Create Snapshot	> Create Local User	Event Log	OS Update	> Snap Finder Properties
	LDAP/NIS	> Snapshot Schedules	> Password Policy	Tape	> Update Notification	BitTorrent Sync
	FTP	iSCSI	Local Groups		> Check for Updates	Change Password
	SNMP	> Create iSCSI Disk	> Create Local Group		> OS Update Status	Mgmt. Interface Settings
	Web	> VSS/VDS Access Control	Security Models		Support	
	ISNS	Disks	ID Mapping		> Registration	
		RDX QuikStor	Home Directories		Tools	
					> Email Notification	
					> Host File Editor	
					> Add Host	
					> Check Filesystem	
					> Check Root Filesystem	

Close

To close the site map, click either **Close** or outside the map.

Contact, Hardware & Software Information

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

The screenshot shows the SnapServer Web Management Interface. A red box highlights the SnapServer logo in the top left corner, with a red arrow pointing to the information page that opens. The information page includes the following details:

© 2003-2014 Overland Storage, Inc. All rights reserved.
 TEL 1.888.343.SNAP
www.snapserver.com

Model	Software	Hardware	Server #	BIOS
VirtualSnap	GOS 7.6.022	04.05.00	9771821	r6.0

Serial #	JVM	Storage Environment
VMware-56 4d f7 72 70 53 70 a0-	1.7.0_45	DynamicRAID™

For technical support visit: www.snapserver.com/support

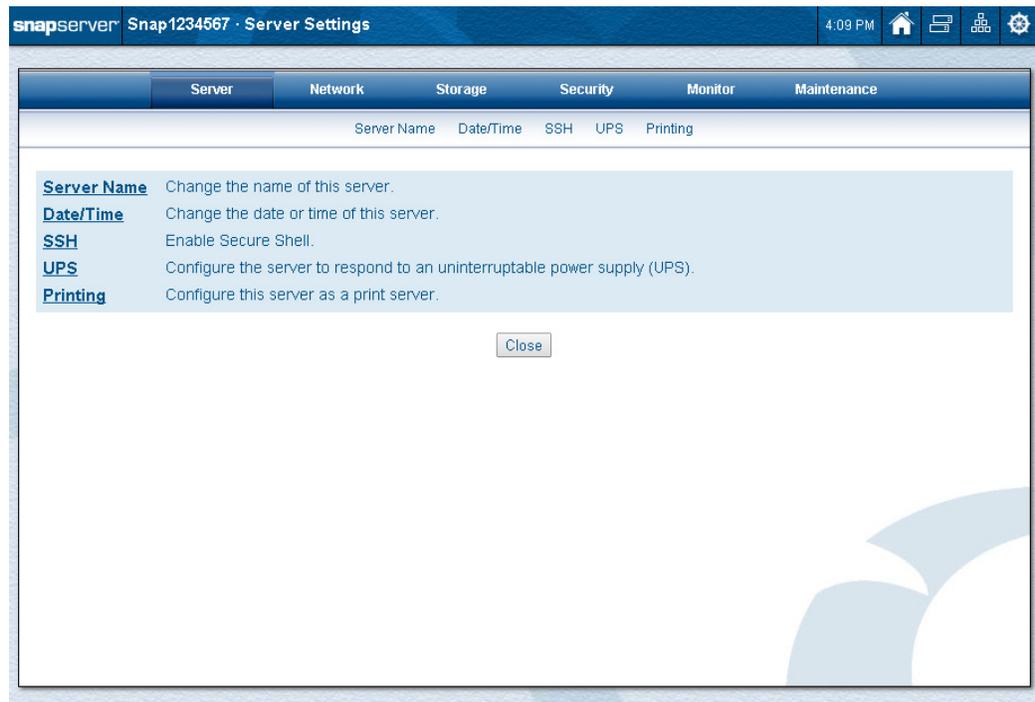
Close

Click here to find out what's new in GuardianOS 7.6 and this Web Management Interface.

Scroll down to view additional contact information. Click either **Close** or outside the box to dismiss.

SnapServer Settings

This section covers the configuration options for a SnapServer appliance. The five options for server settings are found under the Server tab. They can also be accessed using the site map icon (⚙️).

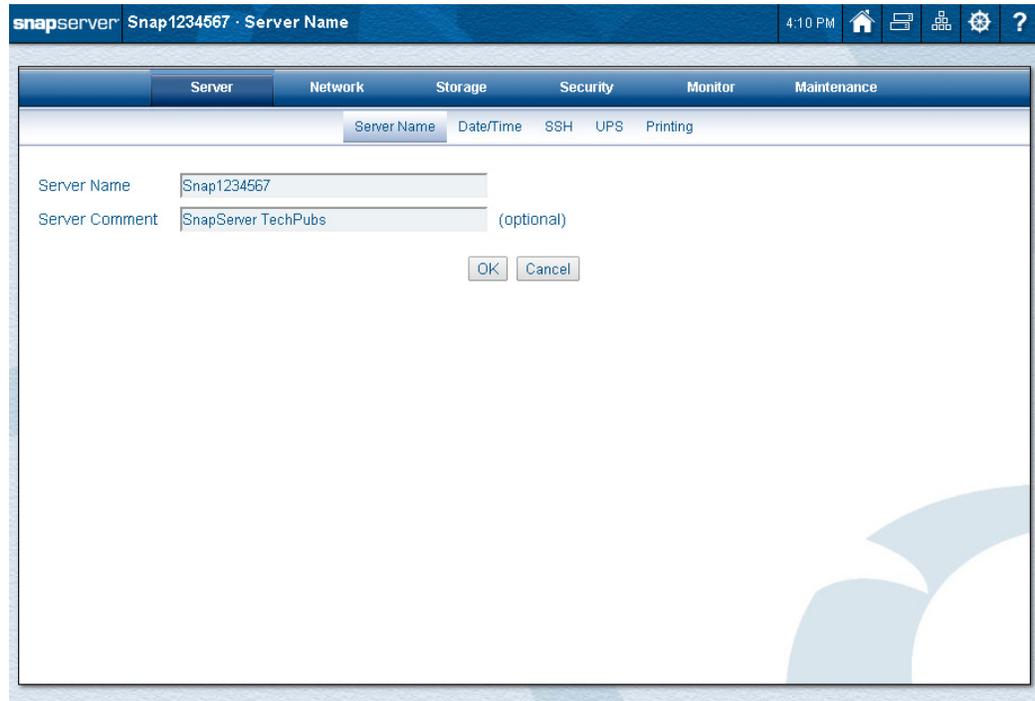


Topics in Server Settings:

- [Server Name](#)
- [Date/Time](#)
- [Secure Shell](#)
- [UPS Protection](#)
- [Printing](#)

Server Name

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



1. Edit the following **fields**:

Option	Description
Server Name	<p>The default server name is Snapnnnnnnnn, where nnnnnnnn is your server number. For example, the default name for a SnapServer with the server number 12345678 would be Snap12345678.</p> <p>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 or by clicking the SnapServer logo at the upper left of the Web Management Interface.</p>
Server Comment	<p>Optionally, add a comment specific to the server (for example, the server location).</p>

2. Click **OK** to save the changes.

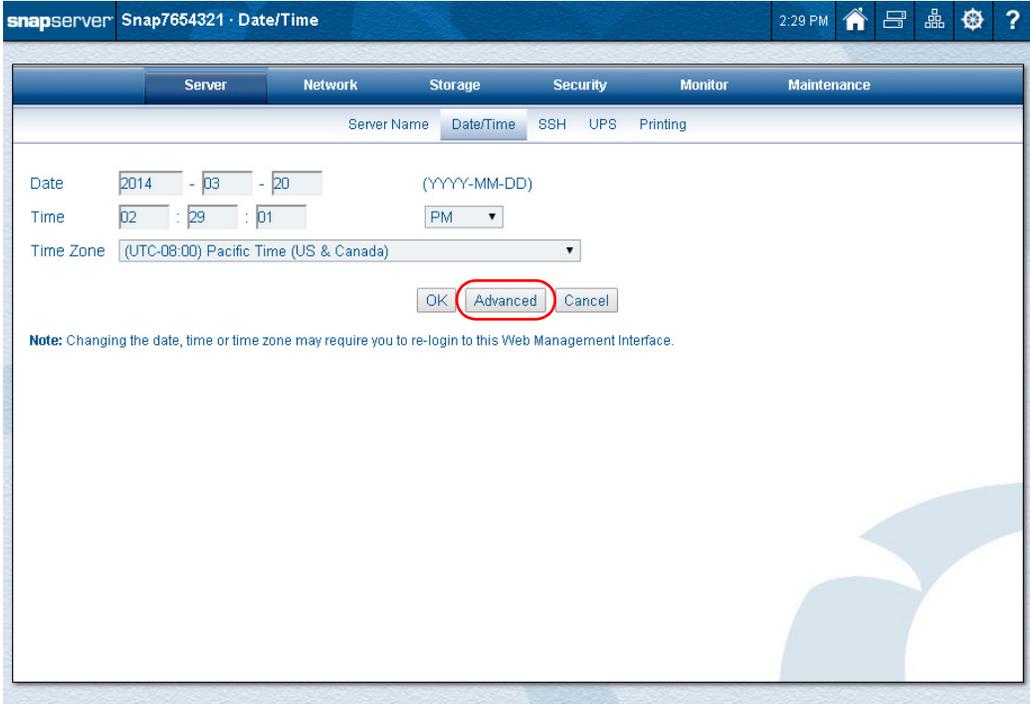
Date/Time

Use this page to configure date and time settings in ISO 8601 formatting. You can set the server date and time manually, or have it set automatically via NTP or Windows Active Directory domain membership.

The time stamp is applied when recording server activity in the Event Log (**Monitor** tab), when creating or modifying files and when scheduling snapshot operations.

 **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.

If NTP was not selected during the setup process, only the manual configuration is shown:



The screenshot displays the SnapServer web management interface for configuring the Date/Time settings. The page title is "SnapServer Snap7654321 · Date/Time". The top navigation bar includes 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:

- Date: 2014 - 03 - 20 (YYYY-MM-DD)
- Time: 02 : 29 : 01 PM
- Time Zone: (UTC-08:00) Pacific Time (US & Canada)

At the bottom of the configuration area, there are three buttons: OK, Advanced (circled in red), and Cancel. A note below the buttons states: "Note: Changing the date, time or time zone may require you to re-login to this Web Management Interface."

To view all options including using NTP servers, click **Advanced**:

The screenshot shows the 'Date/Time' configuration page in the SnapServer web management interface. The page title is 'SnapServer Snap1234567 - Date/Time'. The navigation bar includes tabs for Server, Network, Storage, Security, Monitor, and Maintenance. The 'Date/Time' tab is selected, and sub-tabs for Server Name, Date/Time, SSH, UPS, and Printing are visible. The main content area contains the following text: 'You can set this server's date and time to specific values, or you can use NTP* (Network Time Protocol) servers to automatically synchronize this server's date and time. Visit www.ntp.org for a list of public NTP primary and secondary servers.' Below this, there are two radio button options. The first option, 'Set this server's date and time to the following:', is selected. It includes fields for Date (2014 - 07 - 28) and Time (04 : 12 : 17 PM). The second option, 'Automatically synchronize this server's date and time to the following NTP servers.', is unselected. It includes two text input fields for NTP Servers, with '(optional)' text below the second field. A note below the NTP fields states: '*Note: Due to security restrictions, NTP cannot be used when the server is joined to an Active Directory domain.' Below the NTP options, there is a checkbox labeled 'Enable this server as an NTP server' which is unselected. At the bottom, there is a dropdown menu for 'Time Zone' set to '(UTC-08:00) Pacific Time (US & Canada)'. 'OK' and 'Cancel' buttons are located at the bottom center. A final note at the bottom reads: 'Note: Changing the date, time or time zone may require you to re-login to this Web Management Interface.'

1. Choose to either manually enter or automatically synchronize (using NTP servers) the **date and time**:

NOTE: For security reasons, NTP cannot be used with Active Directory domains.

- **Manually** – Select the first option, enter the correct date and time in the appropriate fields, and use the drop-down list to choose either AM or PM. Once you join a Windows domain, the settings are automatically adjusted to synchronize with the domain settings.
- **Automatically** – Select the second option and enter a valid NTP server IP address or host name. Optionally, enter a second address or name for a second server. In some cases, this change may require you to log back in to the Web Management Interface when done.

2. To use this SnapServer as an NTP server, check the **enable box**.
3. From the drop-down list, select the **time zone**.

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

4. Click **OK**.

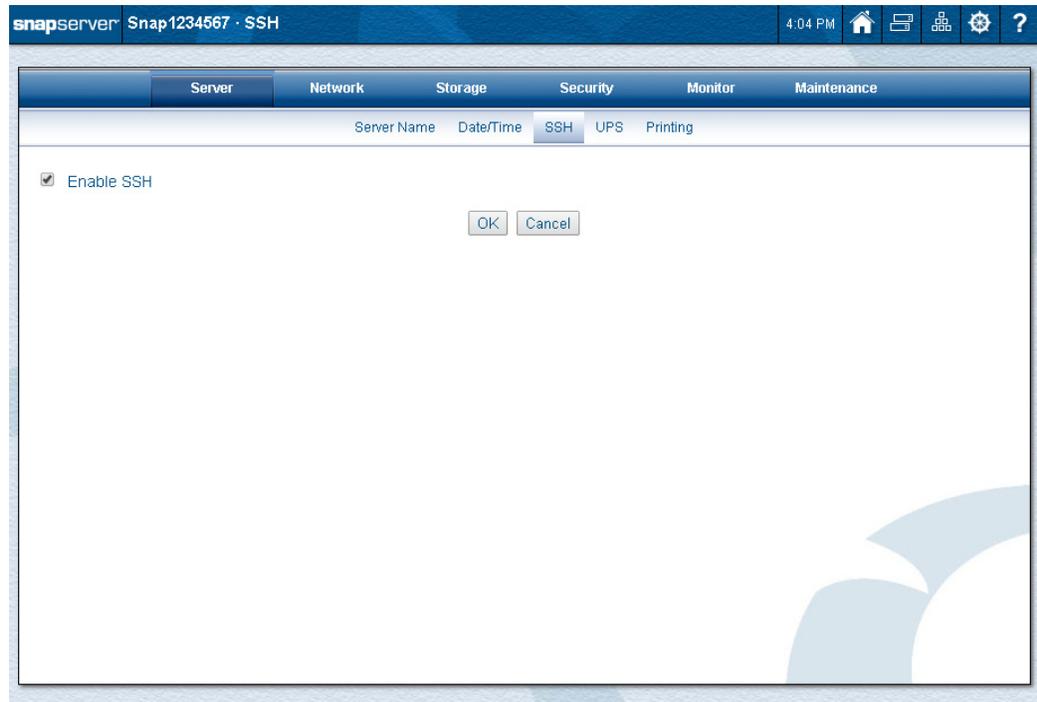
Secure Shell

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 Web Management Interface. See [Command Line Interface on page 299](#) 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, at the **SSH** page, uncheck the **Enable SSH** 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.
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.

UPS Protection

SnapServer supports automatic shutdown when receiving a low-power warning from an APC uninterruptible power supply (UPS). Use **Server > UPS** to manage this feature:

The screenshot displays the SnapServer web interface for configuring UPS protection. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'UPS' tab is active, showing a sub-menu with 'Server Name', 'Date/Time', 'SSH', 'UPS', and 'Printing'. The main content area contains the following text and form elements:

You can configure this server to automatically shut down when a low power warning is sent from a USB- or network-based uninterruptible power supply (UPS). Currently, only APC® brand UPS devices* are supported.

Enable UPS Support

Automatically restart server when power is restored or UPS comes back online

Use a single USB-connected UPS device

Use the following network-connected UPS device(s):

Primary UPS Device

IP Address	10.25.12.10
APC User Name (for authentication)	apc1
APC Authentication Phrase	•••••

Secondary UPS Device (optional)

IP Address	10.25.12.11
APC User Name (for authentication)	apc2
APC Authentication Phrase	•••••

Shut down server if low battery message received from: Either UPS device

Buttons: OK, Refresh, Cancel

*UPS technology copyright © 2001 American Power Conversion Corporation.

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 one or more APC network-enabled or USB-based UPS devices (some serial-only APC UPS devices are also supported by using the IOGear GUC232A USB to Serial Adapter Cable). To do this, you must enable UPS support on the server (as described in this section) to listen to the IP address of one or more APC UPS devices and you must supply the proper authentication phrase configured on the UPS devices.

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

To set up APC UPS support on the server:

1. Verify that the **SnapServer** is plugged into an APC UPS.
2. Go to **Server > UPS**.
3. Check **Enable UPS Support**.
4. Select **one** of the following:
 - Select **Use a single USB-connected UPS device** if your UPS is USB-based.
 - Select **Use the following network-connected UPS device(s)** if your UPS is network-enabled.

5. Enter the **UPS Device** data in the appropriate **fields**:

Option	Description
Enable UPS Support	Check the Enable UPS Support box to enable; leave the 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 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 server 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)	If your server has two power supplies, this option is shown for connecting a second network-connected UPS device. Check the Secondary UPS device check box to enable; leave the check box blank to disable secondary UPS support.
Low Battery Alert	If your server has two power supplies, this option is shown. 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 network-connected UPS devices. • Both UPS Devices: Select this option to allow shutdown only upon receipt of one message from each of the two specified network-connected UPS devices.

6. Click OK.

The system activates the UPS connection. Allow a few minutes for the system to update and show the **Online** status.

The screenshot displays the SnapServer web interface for configuring UPS settings. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'UPS' tab is selected, showing configuration options for two network-connected UPS devices. The primary device is configured with IP 10.25.12.10 and is currently 'Online'. The secondary device is at IP 10.25.12.11 and is 'On Battery'. A red arrow points to the 'Online' status of the primary device. The interface also includes checkboxes for enabling UPS support and automatically restarting the server, and a dropdown menu for selecting the device to shut down the server if a low battery message is received.

Server Name Date/Time SSH UPS Printing

You can configure this server to automatically shut down when a low power warning is sent from a USB- or network-based uninterruptable power supply (UPS). Currently, only APC® brand UPS devices* are supported.

Enable UPS Support

Automatically restart server when power is restored or UPS comes back online

Use a single USB-connected UPS device

Use the following network-connected UPS device(s):

Primary UPS Device

IP Address	10.25.12.10
APC User Name (for authentication)	apc1
APC Authentication Phrase	•••••
APC Status:	→ Online

Secondary UPS Device (optional)

IP Address	10.25.12.11
APC User Name (for authentication)	apc2
APC Authentication Phrase	•••••
APC Status:	→ On Battery

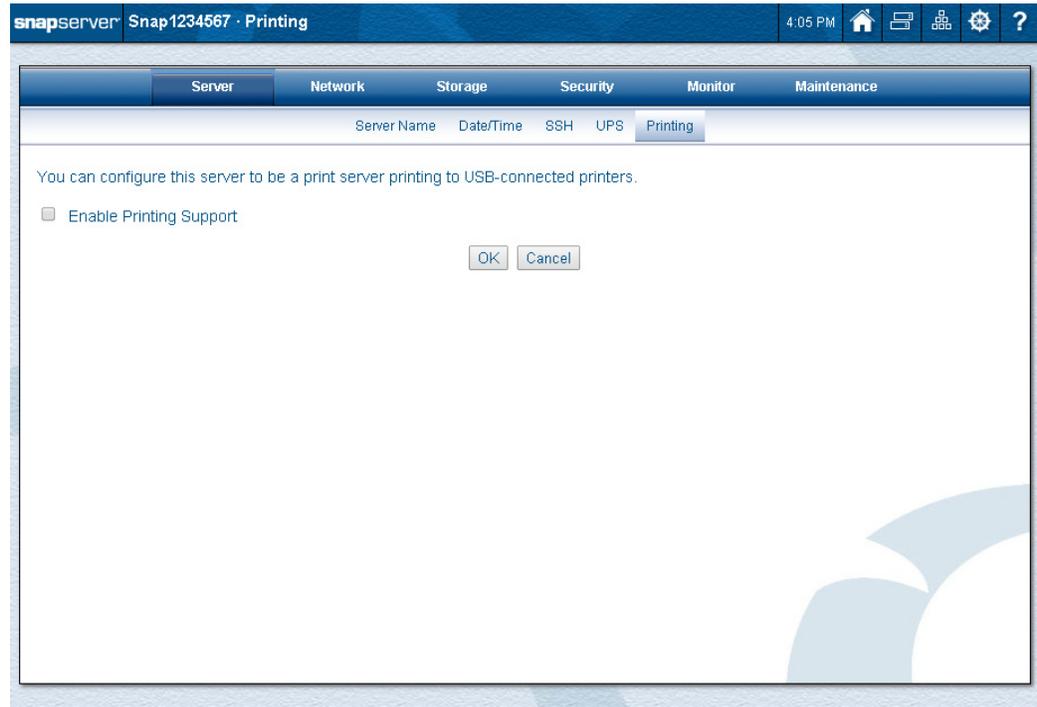
Shut down server if low battery message received from: Either UPS device

OK Refresh Cancel

*UPS technology copyright © 2001 American Power Conversion Corporation.

Printing

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:** Configure the **printer** on the SnapServer.
- Step 2:** Configure 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 Web Management Interface, navigate to **Server > Printing**.
A list of currently defined USB printers is displayed.
5. To add the new printer, click **Add Local Printer**.
6. The SnapServer will detect the new printer and show it as an option in the **Local Printer Device** drop-down list. Select that **printer**.
7. **Name** the printer and, if desired, complete Description and Location information.
8. 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 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 **Pause Printer** to pause all print jobs.
When the printer is paused, the button becomes **Resume Printer**, 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 **Cancel All Jobs**. Click **Refresh** to update the page with the current list of print jobs.

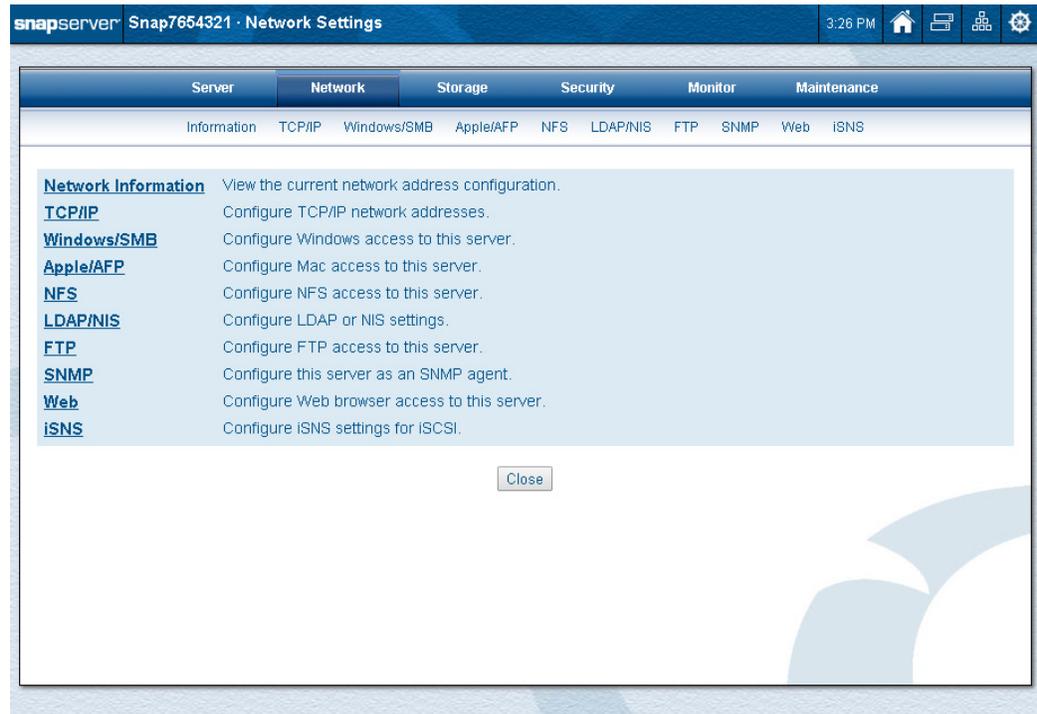
To Delete a Printer

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 the printer you just removed appears as **Offline**.
3. Click the printer link to open the **Edit Printer** page, then click **Delete** to remove the printer.

This chapter addresses the network options for configuring TCP/IP addressing, network bonding, and access protocols.

SnapServers are preconfigured to use DHCP to autonegotiate network settings and to allow access to the server for Windows (CIFS/SMB1/SMB2), Unix (NFS), Mac (AFP), FTP/FTPS, and HTTP/HTTPS clients. 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 Settings:

- [Network Information](#)
- [TCP/IP Networking](#)
- [Windows/SMB Networking](#)
- [Apple Networking \(AFP\)](#)
- [NFS Access](#)
- [LDAP/NIS](#)
- [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.

Network Information

Browse to **Network > (Network) Information** to access the **Network Information** page that displays the current network settings. One column appears for each Ethernet port in use.

The screenshot shows the SnapServer web interface for 'Snap1234567 - Network Information'. The 'Network' tab is selected, and the 'Information' sub-tab is active. The page displays configuration for two Ethernet interfaces: Ethernet 1 (Primary Interface) and Ethernet 2. Below the interface settings, there is a 'Gateway Information' section showing a default gateway of 10.25.1.1, and a 'DNS Information' section showing a domain name of devnet.myoverland.net and primary/secondary DNS servers at 10.6.8.34 and 10.6.8.35 respectively.

Ethernet Interface Information		
Port Name	Ethernet 1 (Primary Interface)	Ethernet 2
Enabled	Yes	Yes
TCP/IP Mode	Static	DHCP
IP Address	10.25.3.35	10.25.3.10
Subnet Mask	255.255.0.0	255.255.0.0
Primary WINS Server	-	-
Secondary WINS Server #1	-	-
Secondary WINS Server #2	-	-
Secondary WINS Server #3	-	-
Ethernet Address	00:0C:29:8B:BF:3D	00:0C:29:8B:BF:47
Speed Status	1000 Mbps (Auto)	1000 Mbps (Auto)
Duplex Status	Full Duplex (Auto)	Full Duplex (Auto)
Bonding Status	Standalone	Standalone

Gateway Information	
Default Gateway	10.25.1.1

DNS Information	
Domain Name	devnet.myoverland.net
Primary DNS	10.6.8.34
Secondary DNS #1	10.6.8.35
Secondary DNS #2	-

Field definitions for the **Network Information** page are given in the following table:

Ethernet Interface Information	
Port Name	The names of the Ethernet interfaces.
Enabled	Yes or No.
TCP/IP Mode	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.

Ethernet Interface Information	
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 (#1, #2, and #3)	Secondary Windows Internet Naming Service servers. Up to three secondary servers can be used.
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 (ALB), 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 (#1 and #2)	Up to two secondary Domain Name System servers can be used.

TCP/IP Networking

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

The screenshot shows the SnapServer web interface for TCP/IP Networking. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under 'Network', there are sub-tabs for 'Information', 'TCP/IP', 'Windows/SMB', 'Apple/AFP', 'NFS', 'LDAP/NIS', 'FTP', 'SNMP', 'Web', and 'iSNS'. The main content area contains a table with the following data:

Port/Bond	Status	IP Address	Bond Type	Modified
Ethernet 1	OK (Primary Interface)	10.25.3.35 (static)	Standalone	No
Ethernet 2	OK	10.25.3.10 (DHCP)	Standalone	No

Below the table are three buttons: 'OK', 'Create Bond', and 'Cancel'.

The following table for the **TCP/IP Networking** page 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 Port Properties on page 52 .
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.

Column	Description
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. Both ports of the bond need to be connected to the same physical or logical switch, and the switch ports must be configured for static link aggregation. • Link Aggregation (802.3ad) – This method of combining or aggregating multiple network connections in parallel is used to increase throughput beyond what a single connection could handle. It also provides a level of redundancy in case one of the links fails. It uses Link Aggregation Control Protocol (LACP), also called dynamic link aggregation, to autonegotiate trunk settings. Both ports of the bond need to be connected to the same switch or logical switch.
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 Port Properties

To configure the TCP/IP properties of a specific port or bond, click the name in the table on the **TCP/IP Networking** page. A **TCP/IP Port Properties** page displays the configuration options for the Ethernet port selected.

The following table for the **TCP/IP Port Properties** page describes these options.

Option	Setting	Description
Enable Ethernet <i>n</i>	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 n</i>) can be disabled by selecting the port and unchecking the Enable Ethernet <i>n</i> box. 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 (DHCP or Static)	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 Speed & Duplex	Using the drop-down list, the SnapServer may also be set to one of five fixed speed (Mbps)/duplex settings: <ul style="list-style-type: none"> • 10 Half Duplex • 10 Full Duplex • 100 Half Duplex • 100 Full Duplex • 1000 Full Duplex. NOTE: To prevent connectivity problems when changing to a fixed setting, see Changing from Auto to a Fixed Link Setting on page 54 .
Primary Interface	Checked or Unchecked	By default, the primary Ethernet port is Ethernet n 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.

TCP/IP Configuration Considerations

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 Switch Trunking or Link Aggregation (802.3ad) network bonding configuration, be sure the switch is configured correctly for that bonding method **after** configuring the bond on the server. 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 Networking** 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.

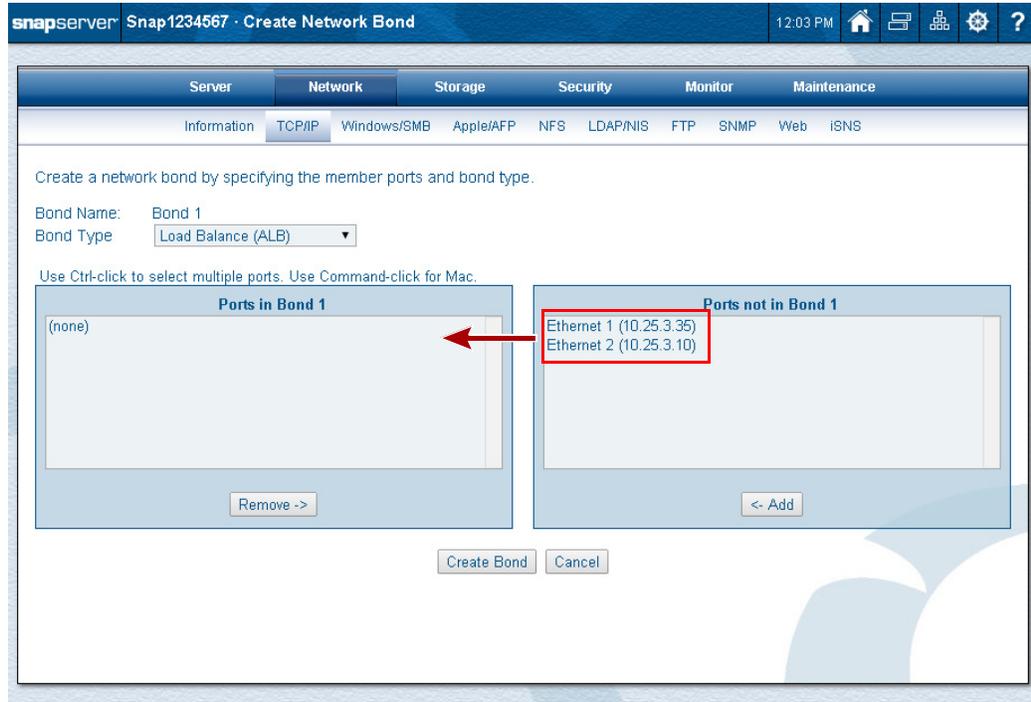


IMPORTANT: If you change the switch setting **before** you change the setting in the Web Management Interface, the SnapServer may not connect to the network. The **Link** LED on the SnapServer front panel will be off or amber if the server is not connected to the network.

Creating a Bond

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

1. At the **TCP/IP Networking** page, click **Create Bond**.
2. Using the **Bond Type** drop-down list, select a **bonding type**:



- **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 or Link Aggregation (802.3ad)** – use either of these two options to 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).

3. 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.
4. Click **Create Bond**.

The **TCP/IP Networking** page is displayed showing the bond details:

The screenshot shows the SnapServer web interface for TCP/IP Networking. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'Network' section is active, with sub-tabs for 'Information', 'TCP/IP', 'Windows/SMB', 'Apple/AFP', 'NFS', 'LDAP/NIS', 'FTP', 'SNMP', 'Web', and 'iSNS'. Below the navigation, a table displays bond details:

Port/Bond	Status	IP Address	Bond Type	Modified
Bond 1	Ethernet 1 - OK Ethernet 2 - OK	10.25.3.35 (static)	Load Balance (ALB)	Created

Below the table are buttons for 'OK', 'Create Bond', and 'Cancel'. A red arrow points to the 'Created' text in the 'Modified' column. Below the buttons, a note reads: '(Important: Click OK to save your changes.)'

5. Click **OK** to save the changes.



CAUTION: The changes made require restarting the server's network. Restarting the server's network will disconnect all connected clients.

6. At the confirmation/restart page, click **Save Changes**.

The screenshot shows the SnapServer web interface for TCP/IP Networking. The top navigation bar is the same as in the previous screenshot. Below the navigation, a warning message is displayed:

Warning: The changes you have made may require restarting the server's network. Restarting the server's network will disconnect all connected clients.

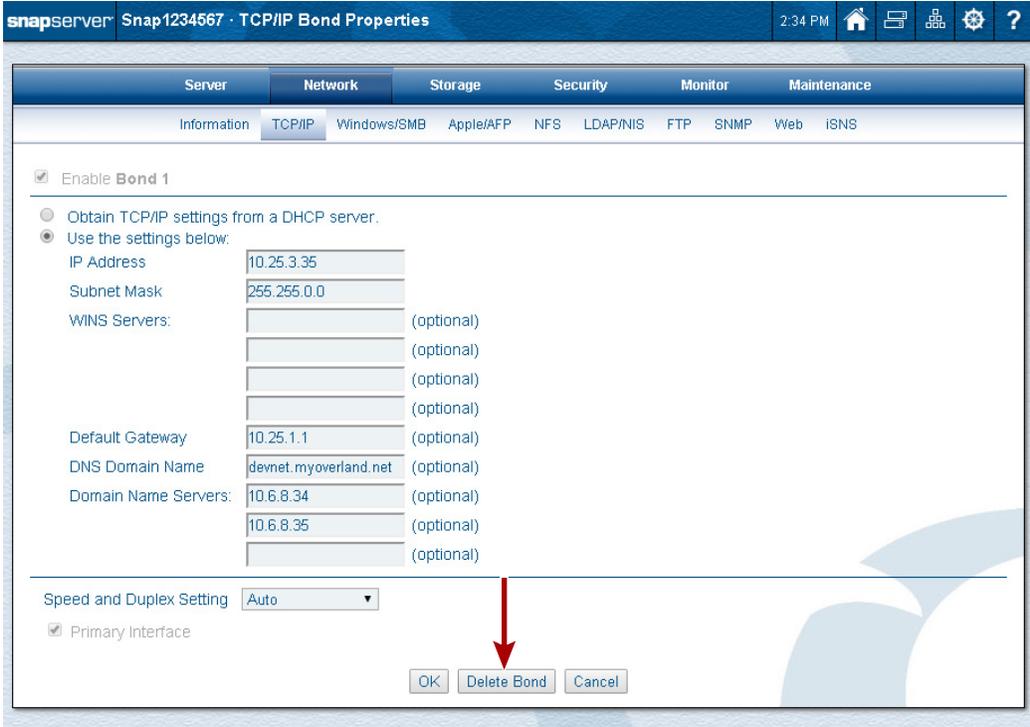
Below the warning, the text reads: 'OK to save your TCP/IP changes?'. At the bottom, there are buttons for 'Save Changes' and 'Cancel'.

 **IMPORTANT:** You must reconfigure the network switch accordingly if using Switch Trunking or Link Aggregation (802.3ad).

Deleting a Bond

On a SnapServer with an existing bond, the bond can be deleted as follows:

1. At the **TCP/IP Networking** page, click the **bond name** in the table to view the properties page.



The screenshot displays the 'TCP/IP Bond Properties' configuration window for 'Bond 1'. The window has a title bar with 'snapserver Snap1234567 - TCP/IP Bond Properties' and a system tray showing '2:34 PM' and various icons. Below the title bar are tabs for 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under the 'Network' tab, there are sub-tabs for 'Information', 'TCP/IP', 'Windows/SMB', 'Apple/AFP', 'NFS', 'LDAP/NIS', 'FTP', 'SNMP', 'Web', and 'iSNS'. The 'TCP/IP' sub-tab is active, showing a form with the following fields and values:

- Enable Bond 1
- Obtain TCP/IP settings from a DHCP server.
- Use the settings below:
- IP Address: 10.25.3.35
- Subnet Mask: 255.255.0.0
- WINS Servers: (optional)
- Default Gateway: 10.25.1.1 (optional)
- DNS Domain Name: devnet.myoverland.net (optional)
- Domain Name Servers: 10.6.8.34 (optional), 10.6.8.35 (optional)
- Speed and Duplex Setting: Auto
- Primary Interface

At the bottom right of the form, there are three buttons: 'OK', 'Delete Bond', and 'Cancel'. A red arrow points to the 'Delete Bond' button.

2. Click **Delete Bond**.

The **TCP/IP Networking** page is displayed showing the details of the unbonded ports.

Click a Port/Bond name to edit the TCP/IP settings. Click a bond's port members to edit the members.

Port/Bond ▲	Status	IP Address	Bond Type	Modified
Ethernet 1	OK (Primary Interface)	10.25.3.35 (static)	Standalone	Unbonded
Ethernet 2	OK	(not available)	Standalone	Unbonded

OK Create Bond Cancel

(Important: Click OK to save your changes.)

3. Click **OK** to save the changes.



CAUTION: The changes made require restarting the server's network. Restarting the server's network will disconnect all connected clients.

4. At the confirmation page, click **Save Changes**.

Warning: The changes you have made may require restarting the server's network. Restarting the server's network will disconnect all connected clients.

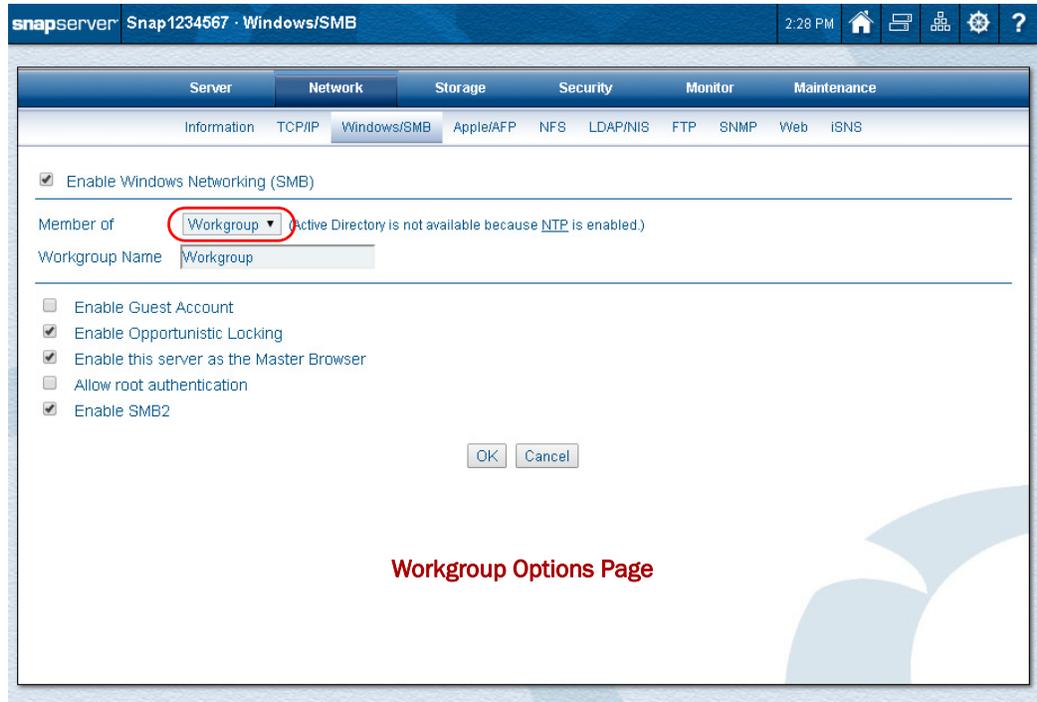
OK to save your TCP/IP changes?

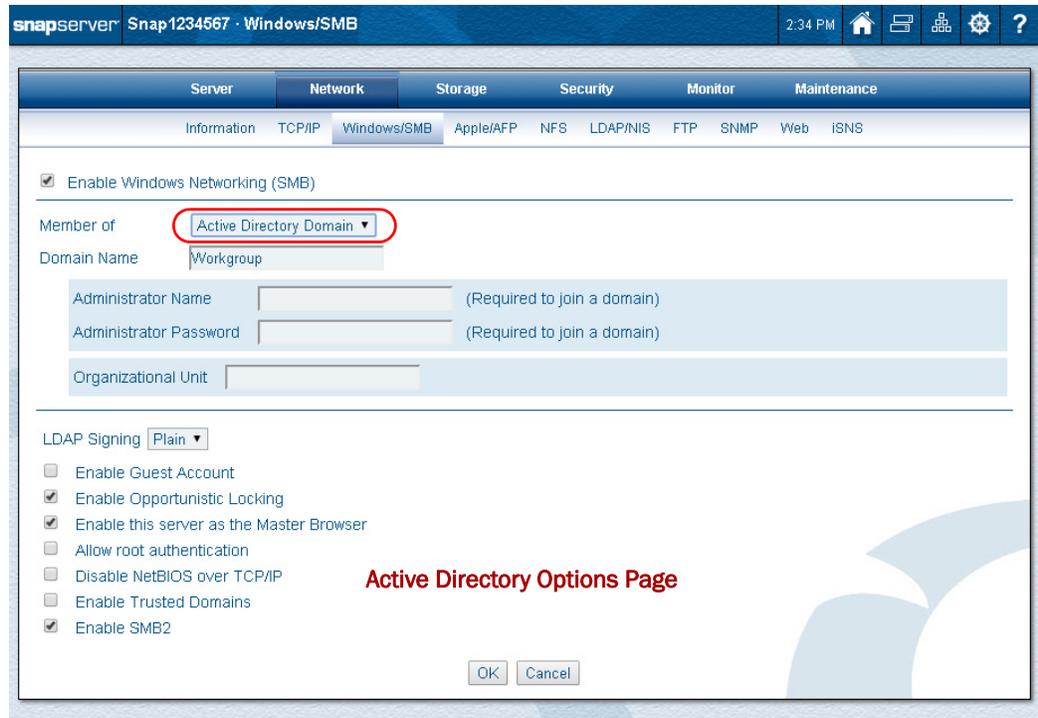
Save Changes Cancel

 **IMPORTANT:** You must reconfigure the network switch accordingly if removing Switch Trunking or Link Aggregation (802.3ad).

Windows/SMB Networking

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 in these two screens:





To use Active Directory domain mode, you cannot configure Date/Time to synchronize with an NTP server. NTP is not supported in Active Directory.

Support for Windows/SMB Networking

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 **Advanced Share Properties** to access the **Hide this share** option.

- For existing shares, select the share, click **Properties**, and click **Advanced Share Properties** to access the **Hide this 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 Service (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 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 family of servers that run in ADS mode. Any SnapServer can join Active Directory domains as a member server. References to the SnapServer 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.

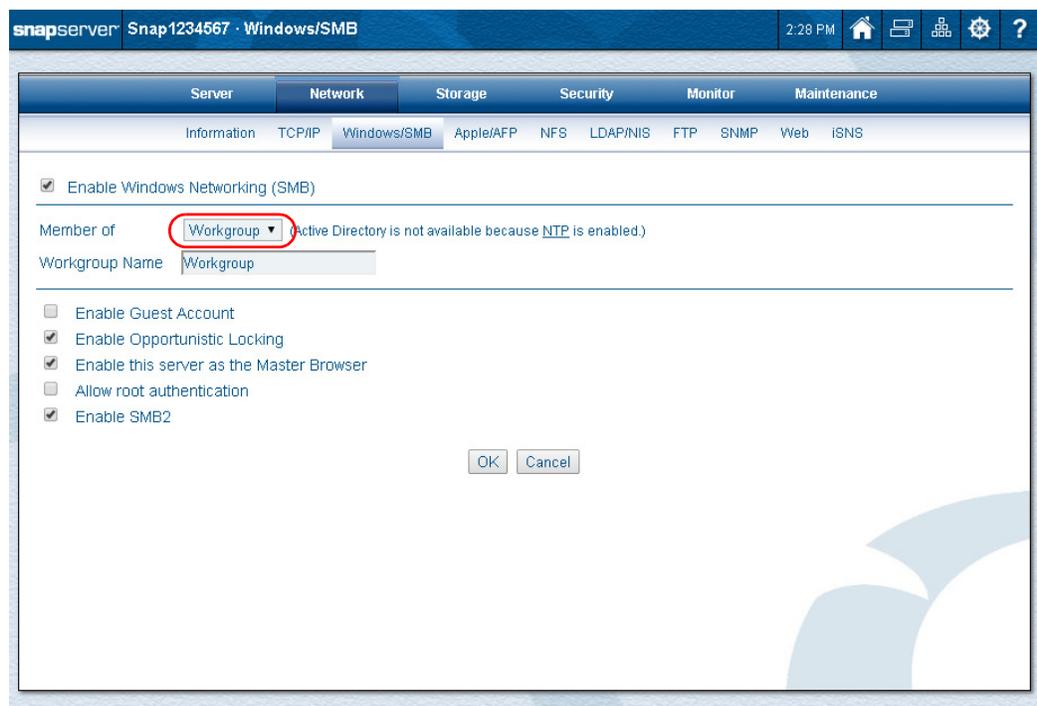
Configure Windows/SMB Networking

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/SMB Networking on page 60](#) and [Support for Windows Network Authentication on page 61](#).

To Join a Workgroup

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



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

Option	Settings
Enable Windows Networking (SMB)	Check the box to enable SMB and activate the options. Clear the box to disable.
Member Of	Verify that it is set to Workgroup . NOTE: For the Active Directory Domain option, see To Join an Active Directory Domain below.
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.

Option	Settings
Enable Guest Account	Check the box to allow unknown users (or users explicitly logging in as Guest) to access the SnapServer using the guest account. Clear the box 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 box to disable this feature.
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 box 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 box to disable this feature.
Allow Root Authentication	Check the box to allow root login to the server; clear the box to disable this feature. NOTE: The root password is synchronized with the server admin password.
Enable SMB2	Enabled by default. This more robust version of SMB reduces protocol overhead and is used by default by Windows Vista and later clients. Clear the box to disable this feature (clients that default to SMB2 will automatically connect via SMB1).

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.

The screenshot shows the SnapServer configuration interface for Windows/SMB. The 'Member of' dropdown menu is highlighted with a red circle and set to 'Active Directory Domain'. Below it are fields for 'Domain Name', 'Administrator Name', 'Administrator Password', and 'Organizational Unit'. At the bottom, there are checkboxes for various network settings, including 'Enable Guest Account', 'Enable Opportunistic Locking', 'Enable this server as the Master Browser', 'Allow root authentication', 'Disable NetBIOS over TCP/IP', 'Enable Trusted Domains', and 'Enable SMB2'. The 'OK' and 'Cancel' buttons are at the bottom right.

NOTE: You cannot select Active Directory Domain if NTP is enabled.

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

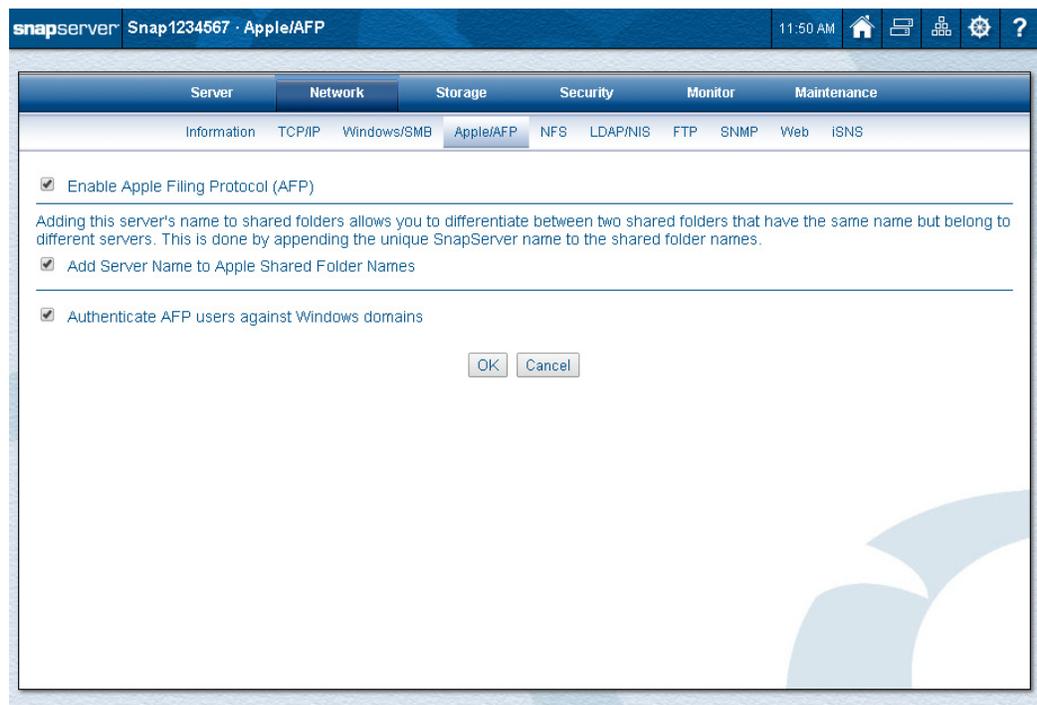
Option	Description
Enable Windows Networking (SMB)	Check the box to enable SMB and activate the options. Clear the box to disable.
Member Of	Verify it shows <i>Active Directory Domain</i> .
Domain Name	The default settings make the SnapServer available in the workgroup named <i>Workgroup</i> . Enter the domain name to which the server belongs. NOTE: Windows 2000 domain controllers must run SP2 or later.
Administrator Name / Administrator Password	If joining a domain, enter the user name and password of a user with domain join privileges (typically an administrative user).
Organizational Unit	To create a machine account at a different location than the default, enter a name in the field. By default, this field is blank, signaling the domain controller to use a default defined within the controller. NOTE: Sub-organizational units can be specified using Full Distinguished Name LDAP syntax or a simple path ([org_unit]/[sub-unit1]/[sub-unit1a])
LDAP Signing	Use the drop-down list to 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 box to allow unknown users (or users explicitly logging in as Guest) to access the SnapServer using the guest account. Clear the box 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 box to disable this feature.
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 box 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 box to disable this feature.
Allow Root Authentication	Check the box to allow root login to the server. Clear the box to disable this feature. NOTE: The root password is synchronized with the server's admin password.
Disable NetBIOS over TCP/IP	Some administrators may wish to disable NetBIOS over TCP/IP. Check the box to disable NetBIOS; clear the box 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 (such as, "actdirdomainname.companyname.com"). A short form such as "ActDirDomName" does not work.

Option	Description
Enable Trusted Domains	<p>SnapServer recognizes trust relationships established between the domain to which the SnapServer is joined and other domains in a Windows environment by default. Check the box to enable this feature; clear the box to disable this feature.</p> <p>NOTE: SnapServer remembers 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, is retained.</p>
Enable SMB2	<p>Enabled by default. This more robust version of SMB reduces protocol overhead and is used by default by Windows Vista and later clients. Clear the box to disable this feature (clients that default to SMB2 will automatically connect via SMB1).</p>

- 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).

AFP Configuration Considerations

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	A SnapServer share appears as a Mac volume that can be accessed through the Finder. 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.
Volume	A volume on a SnapServer is a logical partition of a RAID's storage space that contains a filesystem.
Right-click	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 ." NOTE: Mac users with a single-button should substitute control-click to achieve the same result.

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 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.

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 Filing Protocol (AFP) box to enable AFP; leave the box 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.

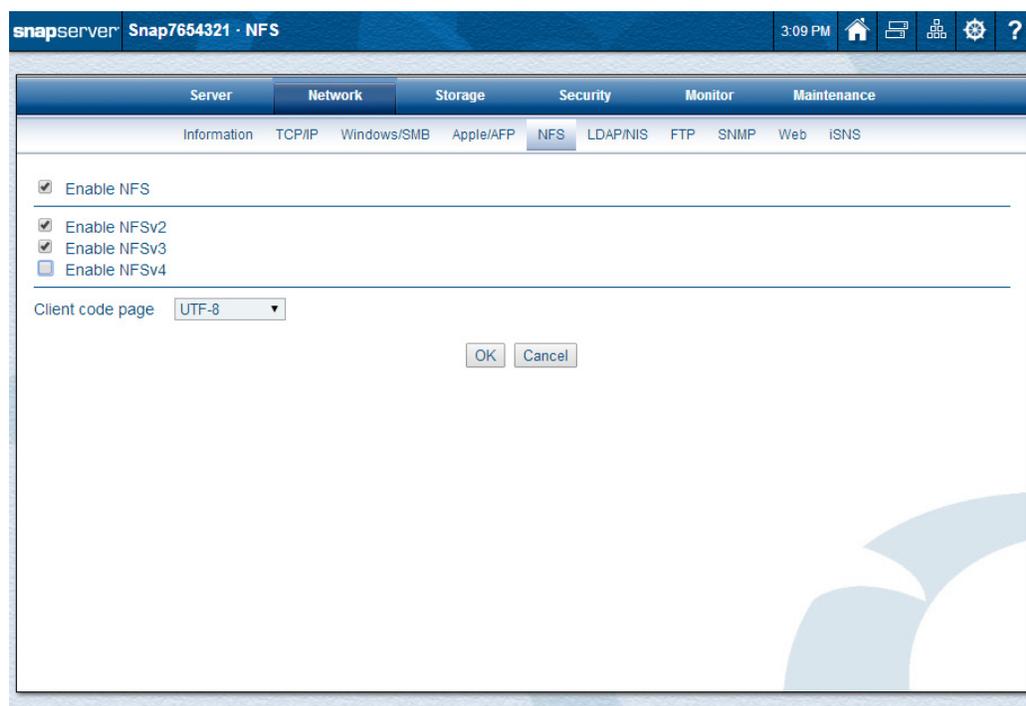
Options	Usage
Authenticate AFP Users Against Windows Domains	Select this option to automatically authenticate AFP users against a Windows domain, if configured. NOTE: By default, users are authenticated against the domain first, then against the local database, so if the same user name exists on both 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 servername. For example: <i>windowsdomain\username</i> or <i>snap12345\username</i>

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

NFS Access

NFS access to the server is enabled on the **Network > NFS** page of the Web Management Interface. By default, most 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, check the **Enable NFSv4** box 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. To ensure proper Unicode representation on the file system, set the client code page to indicate the code page used by NFS clients to represent characters in filenames (usually UTF-8 on modern Unix/Linux-based operating systems).

These versions of the NFS protocol are supported:

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.

Assigning Share Access to NFS Users

The NFSv2/3 protocol does not support user-level share access control, but rather supports host- and subnet-based access control. NFSv4 supports user-level access control via Kerberos configuration, but otherwise uses the same form of host-based access control. On a standard Unix server, share access is configured in an “exports” file. On SnapServers, the exports for each share are configured on the **NFS** page independently of user-based share access for other protocols.

Enable NFS Access to the Server

1. Go to **Network > NFS**.
2. Check the **Enable NFS** box.
3. Check the **versions** you want to enable.
Select one or more from **NFSv2**, **NFSv3**, and **NFSv4**.
4. Choose the desired **Client code page** from the drop-down list.
Select **UTF-8**, **ISO-8859-1**, **ISO-8859-15**, or **EUC-JP**.
5. Click **OK**.

Configure NFSv4 Access

1. Go to **Network > NFS**.
2. Check the **Enable NFS** and **Enable NFSv4** boxes.

A new set of security options are displayed below the **Enable NFSv4** option.

3. Use this table to select the **level of security** you want to apply:

Option	Description
Domain Name	The default domain name “localdomain” is shown in the field. If necessary, you can change it.
Security Type	<ul style="list-style-type: none"> • Standard NFS Security – Choose this option if you want to use standard NFS host- and subnet-based security. • RPSEC GSS Security (Unix Kerberos) – Choose this option and complete the fields that appear if you want to use Unix Kerberos security to authenticate NFSv4 connections. <p>NOTE: Kerberos security can only be configured for Unix-based Kerberos implementations. Windows ADS Kerberos is not supported for NFSv4 authentication.</p>

4. If you selected **RPSEC GSS Security (Unix Kerberos)** security, complete the new options displayed using the table below. Note the following:
- The service will not start unless the TCP/IP domain name 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 Networking on page 50](#).

Option	Description
KDC Host Name	Enter the host name of the Kerberos server (for example, "kerberos-2000.mit.edu").
Realm Name	Enter the Kerberos realm name (For example, "ATHENA.MIT.EDU"). 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, "zeus.keytab"). 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.

5. 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.

LDAP/NIS

LDAP and NIS databases are configured on the **Network > LDAP/NIS** page of the Web Management Interface. Use the drop-down list to choose either LDAP or NIS as the user database type to be configured.

The screenshot displays the configuration interface for LDAP/NIS. The 'User Database Type' is set to 'LDAP'. The 'Enable LDAP' checkbox is checked. The 'LDAP Server' field is empty. The 'LDAP Base DN' field is empty with a 'Search' button to its right. The 'LDAP Bind Type' is set to 'Anonymous'. 'OK' and 'Cancel' buttons are located at the bottom of the configuration area.

LDAP vs. NIS Overview

LDAP (Lightweight Directory Access Protocol) is an open, industry standard application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network. The SnapServer can be configured to query an LDAP directory for user/group names and UIDs/GIDs for configuration of quotas, ID mapping, and home directories. As such, you must use the LDAP directory to make modifications.

NOTE: A SnapServer currently can't be configured to authenticate users against an LDAP directory.

NIS (Network Information Service) is a client-server directory service protocol for distributing system configuration data such as user and host names between computers on a computer network. 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 to translate user/group names to UIDs/GIDs for configuration of quotas, ID mapping, and home directories. As such, you must use the NIS server to make modifications.

NOTE: 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.

Configuring LDAP

Use this procedure to configure LDAP on your SnapServer:

1. Go to **Network > LDAP/NIS**.

2. Verify **LDAP** is displayed in the **User Database Type** drop-down list.
3. Check **Enable LDAP**.
4. Edit the **settings** shown in the following table:

Options	Description
LDAP Server	Enter the host name or IP address for the LDAP server.

Options	Description
LDAP Base DN	Click Search to locate the Base DN on the LDAP server, or enter the Base DN in LDAP syntax such as: <code>cn=accounts,dc=mydir,dc=mydomain,dc=com.</code>
LDAP Bind Type	From the drop-down list, select the LDAP bind type: <ul style="list-style-type: none"> • Anonymous • Simple If Simple is selected, two new fields are shown: Bind DN and Bind Password .

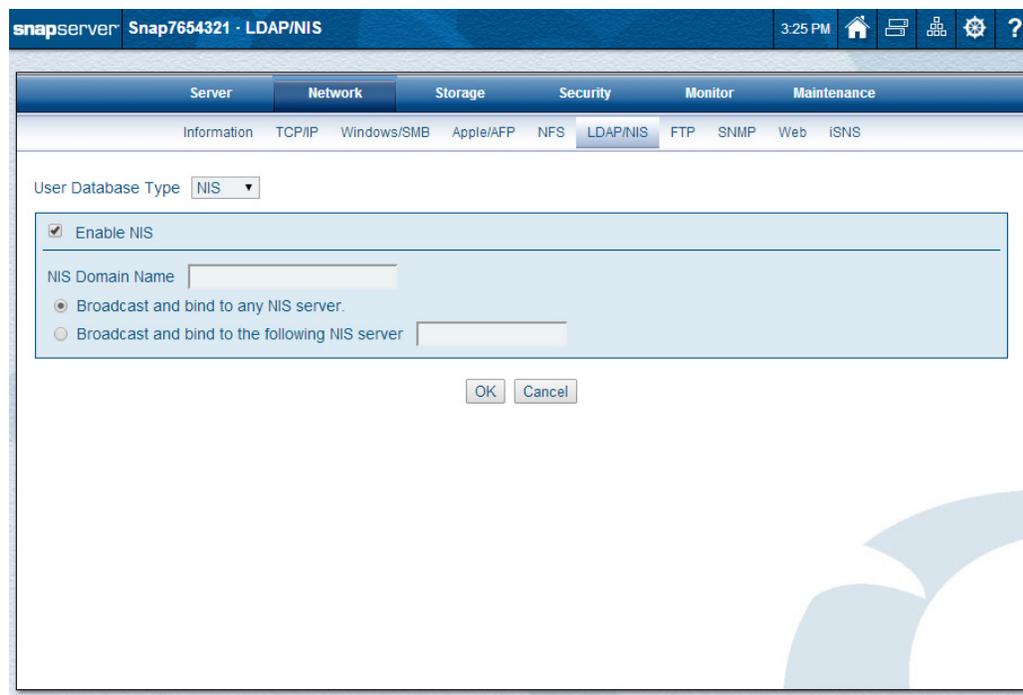
5. If you selected **Simple** as the bind type, complete the two new options (**Bind DN** and **Bind Password**).
6. Click **OK** to update the settings immediately.
If NIS is enabled, you are warned that existing quotas or ID mappings for NIS users will be applied automatically to LDAP users and groups that have the same UID or GID.
7. Click **Enable LDAP** to complete the process.

Configuring NIS

NOTE: 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 in Chapter 8](#).

NIS uniquely identifies users by UID, not user name, and although it is possible to have duplicate user names, Overland Storage does not support that configuration. To configure NIS on your SnapServer:

1. Go to **Network > LDAP/NIS**.
2. Verify **NIS** is displayed in the **User Database Type** drop-down list.



3. Check **Enable NIS**.
4. Edit the **settings** shown in the following table:

Options	Description
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 and enter the IP address for a specific NIS server in the field provided.

5. Click **OK** to update the settings immediately.
If LDAP is enabled, you are warned that existing quotas or ID mappings for LDAP users will be applied automatically to NIS users and groups that have the same UID or GID.
6. Click **Enable NIS** to complete the process.

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 *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.

For FTPS, it is recommended that your FTPS client application use 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.

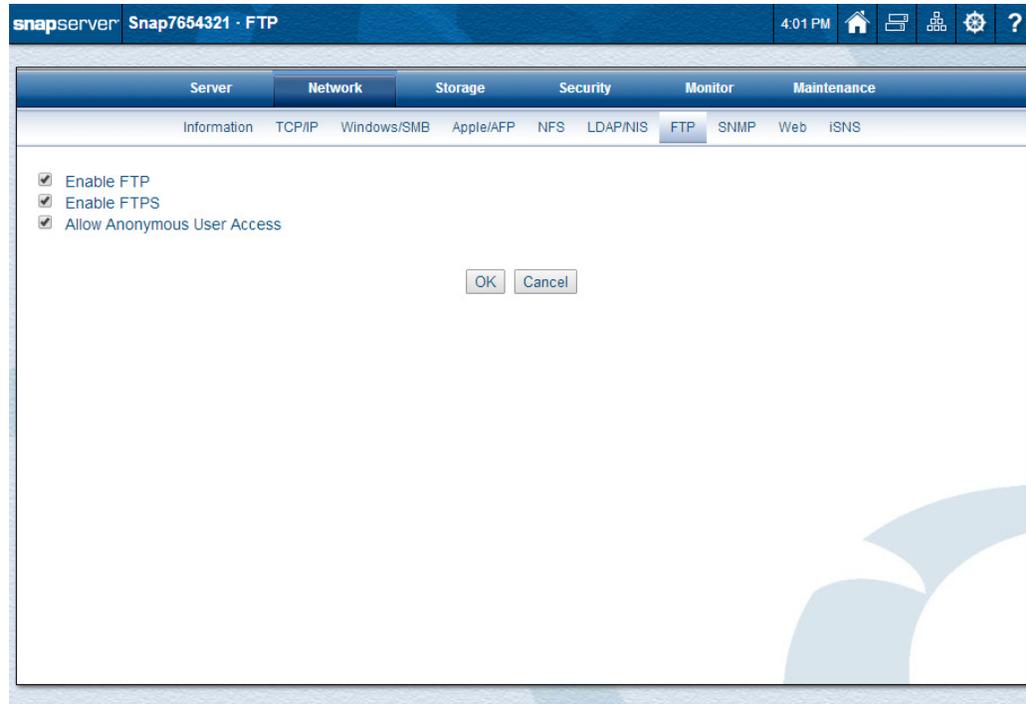
Supported FTP Clients

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 most browsers for standard FTP.

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** as shown in the following table:

Option	Settings
Enable FTP	Check the box to enable standard FTP services; leave the box blank to disable access to this server via standard FTP.
Enable FTPS	Check the box to enable FTPS services; leave the box 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.</p> <ul style="list-style-type: none"> • Check the box to allow users to connect to the server using the anonymous user account. The anonymous user is mapped to the 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. • Leave the box blank so 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 **SnapServer**:

- For **standard FTP**, enter the name of the server 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)
- For **secure FTPS**, configure your FTPS client application to use explicit FTPS (such as, FTPES or “Auth TLS”) and enter the name of the server 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.

2. 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. SnapServer supports the default traps shown in this table:

Trap	Initiating Action
coldStart	Whenever SNMP is enabled and the server boots.
linkDown	A server's Ethernet interface has gone offline.
linkUp	A server's Ethernet interface has come back online.
authenticationFailure	An attempt to query the SNMP agent using an incorrect read-only or read-write 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 OIDs:</p> <ul style="list-style-type: none"> • 1.3.6.1.4.1.6411.2000.1000.1:loglevel 0 syslog messages (<i>emergency</i>) • 1.3.6.1.4.1.6411.2000.1001.1:loglevel 1 syslog messages (<i>alert</i>) • 1.3.6.1.4.1.6411.2000.1002.1:loglevel 2 syslog messages (<i>critical</i>) • 1.3.6.1.4.1.6411.2000.1003.1:loglevel 3 syslog messages (<i>error</i>) <p>NOTE: There is no specific MIB that defines traps sent by SnapServer.</p>

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

The SNMP configuration page can be found at **Network > SNMP**:

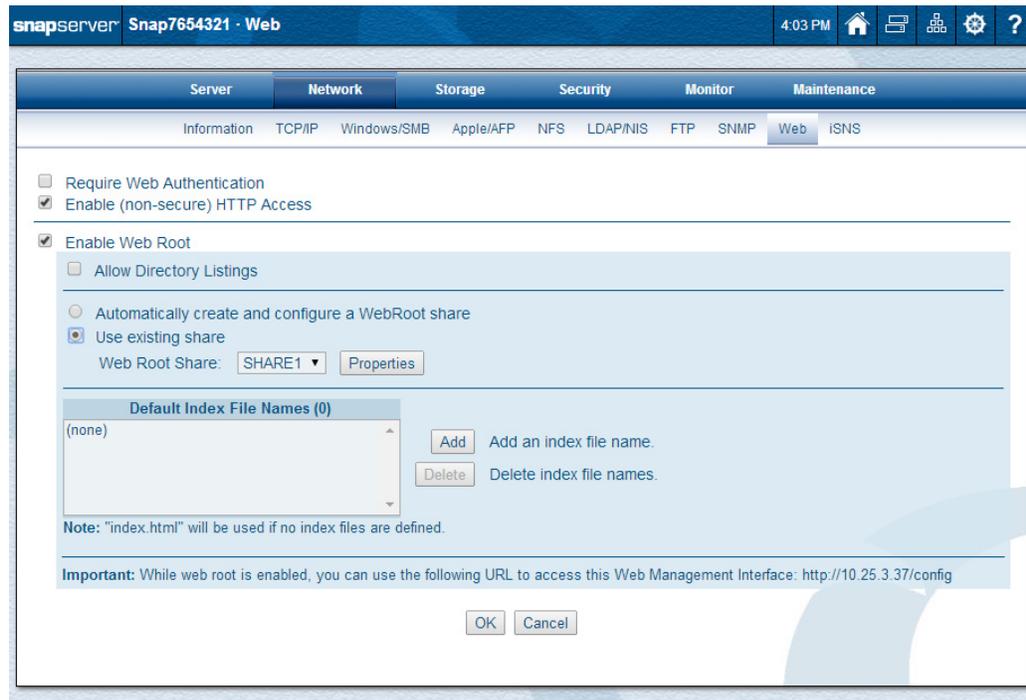
Edit settings as described in the following table and then click **OK**. Once enabled, SNMP managers can access MIB-II and Host Resources MIBs management data on the server.

Option	Description
Enable SNMP	To enable SNMP, check the Enable SNMP box. Leave the box blank to disable SNMP.
Read-Only Community	To allow SNMP managers to read data from this server, enter a read-only community string or accept the default <i>snap_public</i> . NOTE: As a precaution against unauthorized access, Overland Storage recommends that you create your own community string.
Read-Write Community	While SNMP support is read-only, the optional read-write string is used for compatibility purposes. Enter a read-write community string or accept the default <i>snap_private</i> . NOTE: As a precaution against unauthorized access, Overland Storage recommends that you create your own community string.

Option	Description
Location	Optionally 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	Optional. 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 box to enable traps. Clear the box to disable SNMP traps.
IP Address 1-4	Only available when SNMP traps are enabled. Enter the IP address of at least one SNMP manager in the first field as a trap destination. Optionally, you can enter up to three additional IP addresses in fields 2-4.
Send a Test Trap	Only available when SNMP traps are enabled. 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 this **Web** page. Additionally, you can require browser-based clients to authenticate to the server.



Configuring HTTP/HTTPS

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

Edit the options as needed and click **OK**:

Option	Description
Require Web Authentication	<p>Check the Require Web Authentication box to require clients to enter a valid user name and password in order to access the server via HTTP/HTTPS. Leave the box blank to allow all HTTP/HTTPS clients access to the server without authentication.</p> <p>NOTE: This option applies to both Web View and Web Root modes.</p>
Enable (non-secure) HTTP Access	<p>Check the Enable HTTP Access box to enable non-secure HTTP access. Leave the box blank to disable access to the server via HTTP.</p> <p>NOTE: This option applies to both Web View and Web Root modes.</p>
Enable Web Root	<p>Checking the Enable Web Root box displays the settings for web root access that can be configured.</p>

Connect via HTTPS or HTTP

1. Enter the **server name** (or server 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://Snap2302216/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 **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://[server_name] 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://[server_name]` 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:

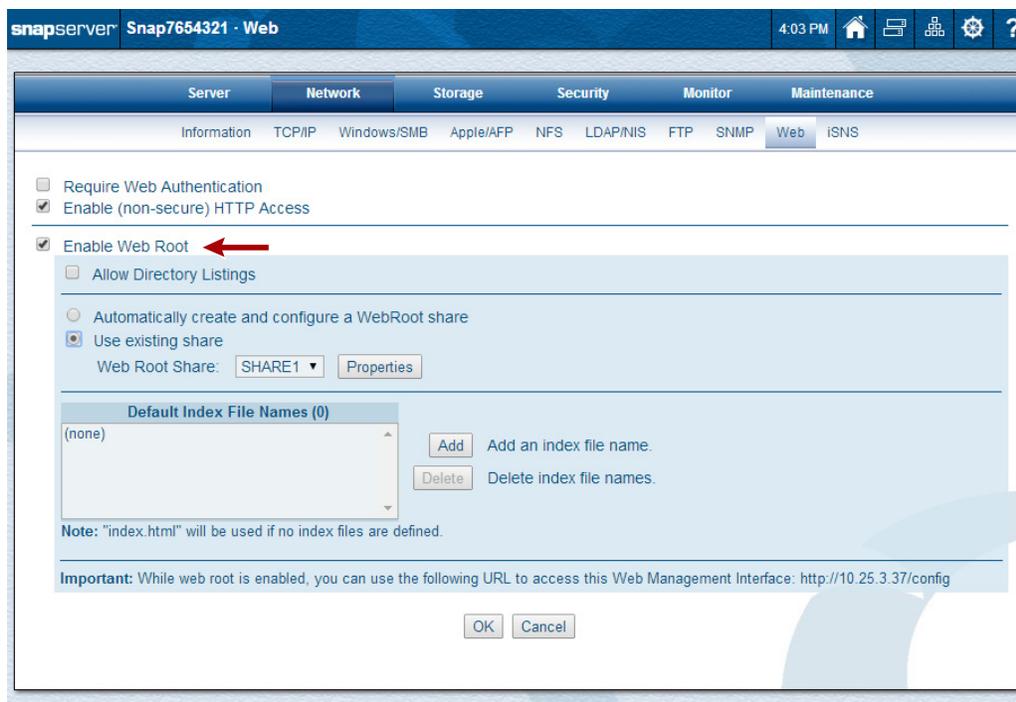
```
http://[server_name]/photos/slideshow.html
```

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

NOTE: 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 SnapServer.

Configure Web Root

Check the **Enable Web Root** box to configure the SnapServer to serve the Web Root directory as the top level web access to the SnapServer and, optionally, automatically serve an HTML file inside. When the box is checked, the options described below appear.



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

Option	Description
Allow Directory Listings	If Allow Directory Listings is checked and no user-defined index pages are configured or present, the browser opens 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" is automatically created. By default, the share is hidden from network browsing and has 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 at Security > Shares. • Use existing share: From the drop-down list of existing shares for selection, select a share and click Properties to edit the selected share's properties (see Security > Shares).

Option	Description
Default Index File Names	<p>Files found underneath the Web Root with names matching those in this list is automatically served to the web browser when present, according to their order in the list. To add a filename, click Add, enter the name of one or more index HTML files, then click OK. The file you entered is shown in the Index Files box.</p> <p>NOTE: If no files are specified, <code>index.html</code> is automatically used if found.</p> <p>To delete a name, highlight it and click Delete. At the confirmation page, click Delete again.</p>

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 of the HTML files 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 loads 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 or change a password), enter the following in the browser address field:

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

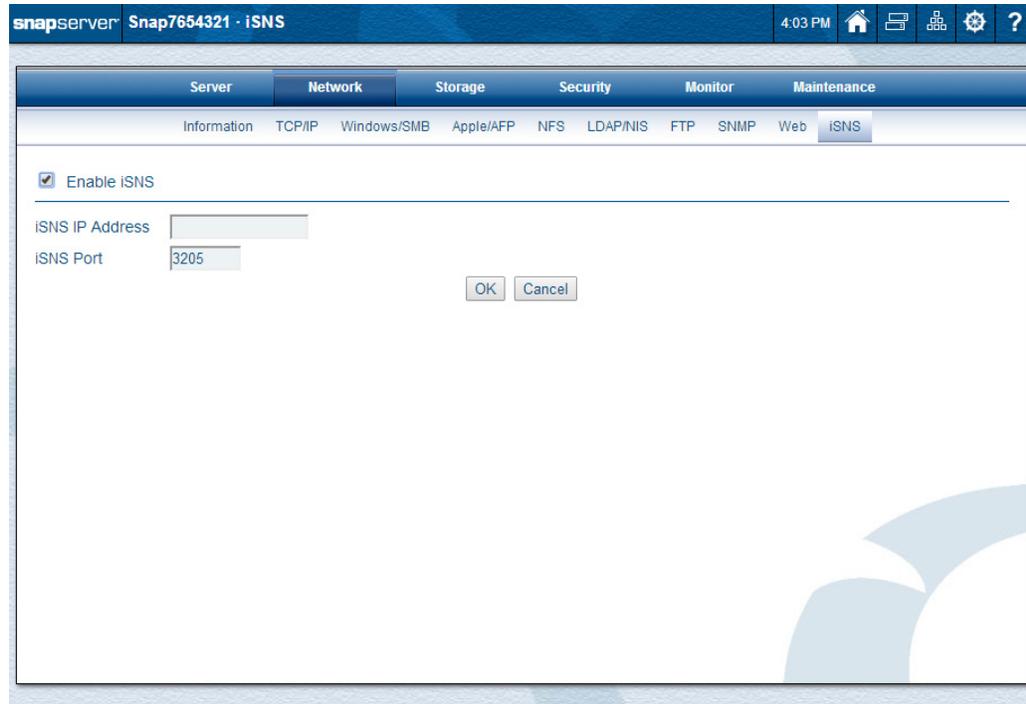
You are prompted for your User ID and password, then you are placed into the Web Management Interface.

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

iSNS Configuration

Microsoft iSNS Server can be used for the discovery of SnapServer iSCSI targets on an iSCSI network.

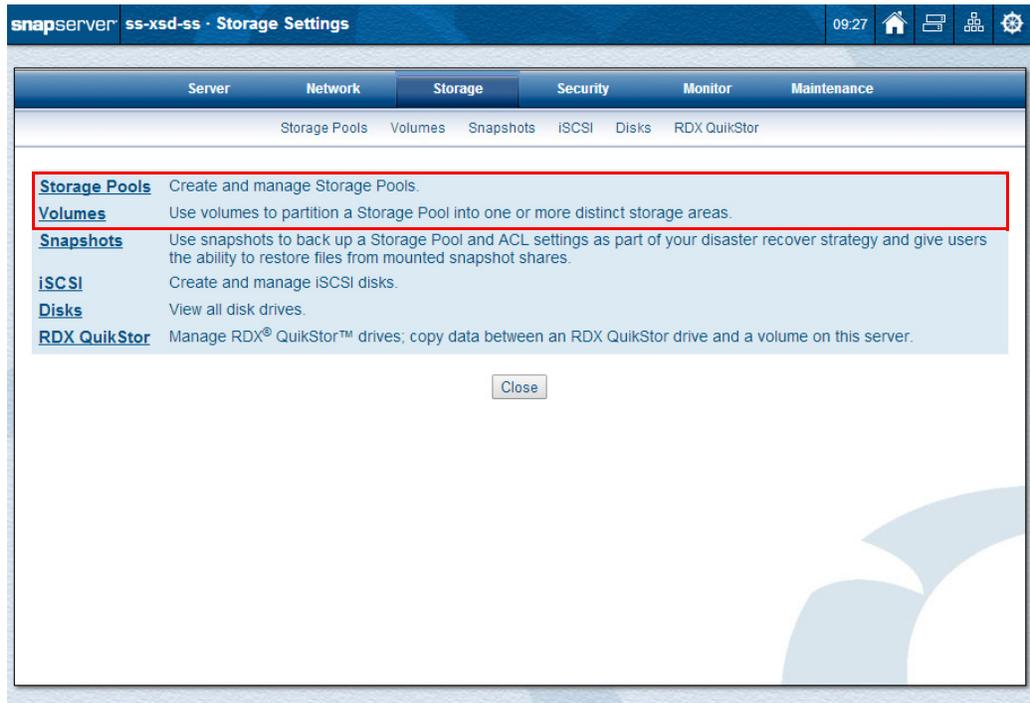


To configure the iSNS settings:

1. If not already installed, install the iSNS service on a **Windows server**.
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 the **Enable iSNS** box, enter the IP address of the iSNS server, and then click **OK**. If the iSNS server does not use the default port, the iSNS port default value of 3205 can be changed on this page as well.
3. Configure the **iSCSI initiator** to discover iSCSI targets via the iSNS server.

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 maximum flexibility.



To determine which RAID configuration is appropriate for your needs, see the [Should I use DynamicRAID or Traditional RAID?](#) section in [Appendix C](#).

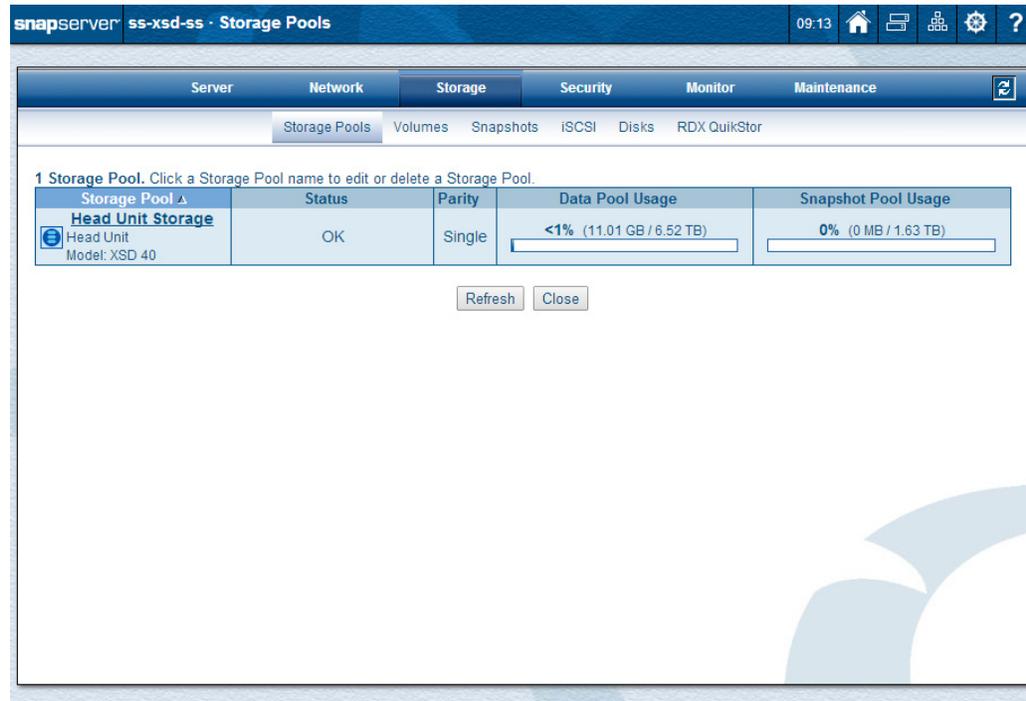
For information on the Traditional RAID configuration option, see [Traditional RAID Storage](#) in [Chapter 6](#). For other storage features, see [Other Storage Options](#) in [Chapter 7](#).

Topics in DynamicRAID Storage:

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

Storage Pools

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

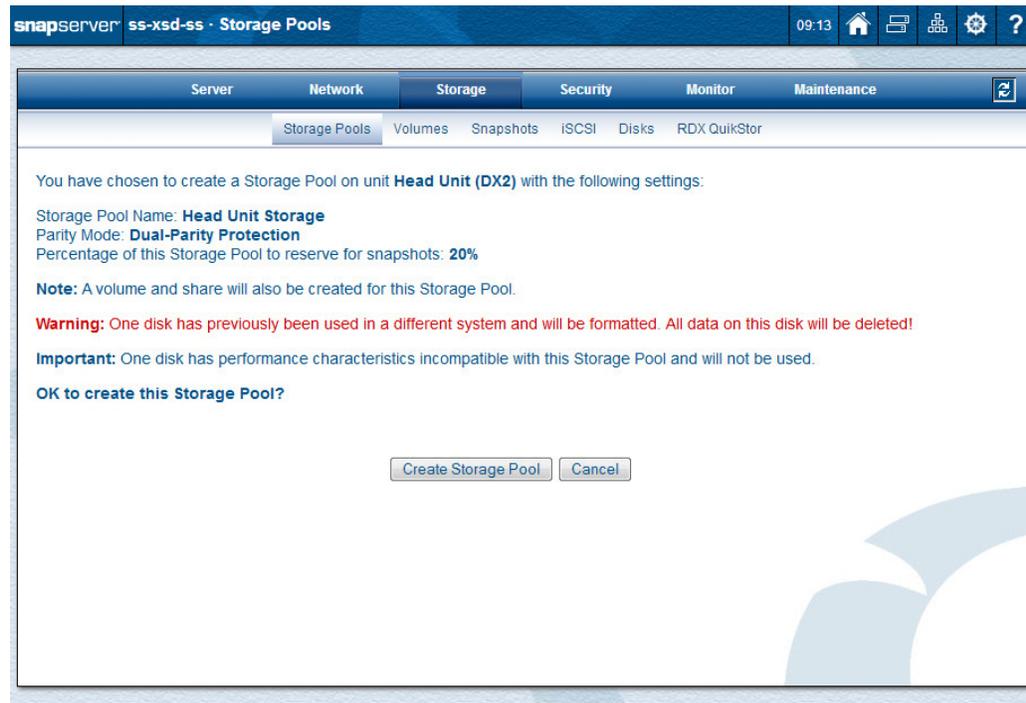


IMPORTANT: A SnapServer head unit or expansion unit supports only **one** storage pool created from its drives and contained within that enclosure. Multiple volumes can be created on that storage pool.

All the 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 (foreign drives) can be added to a head or expansion unit and are then 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. This also applies to drives that are current RAID members and may have been removed inadvertently. Upon reinsertion, they will not be automatically incorporated, regardless of whether automatic incorporation of unassigned drives is turned on.

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



Storage Pool Creation

Storage pools can be created on head and expansion units that do not yet have pools, one pool per unit, with each pool completely contained within the unit.

During the initial setup process, storage pools are created on the head and expansion units using all disk drives available in each unit. DynamicRAID always maximizes the space available based on both the parity mode type and the snapshot pool size requested.

NOTE: 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 size characteristics of the first drive. Otherwise, the drives are not used.

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.

 **CAUTION:** When a storage pool is created, any disks in the storage pool that have previously been used in a different system will be reformatted and all data on the disks will be deleted.

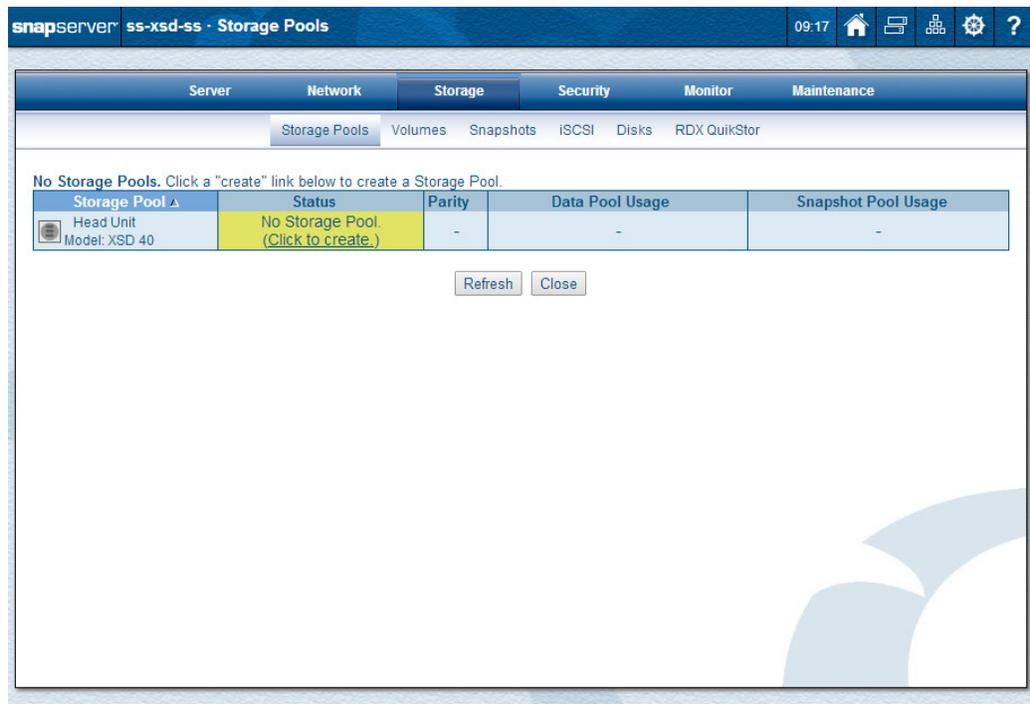
At the **Create Storage Pool** page, you can configure these options:

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.

Option	Description
Parity Mode	<p>Based on the total number of disks that are available for a storage pool, you can set the parity mode of the storage pool:</p> <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	<p>Use the drop-down list to choose a percentage of the storage pool that is reserved for snapshots.</p> <p>For more details about snapshots, refer to Snapshots in Chapter 7.</p> <p>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.</p> <p>Default: 20%</p>

Create a Storage Pool

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



1. At the **Storage Pools** page (**Storage > Storage Pools**), click **No Storage Pool** in the **Status** column to access the **Create Storage Pool** page.

ss-xsd-ss · Create Storage Pool 09:18

Server Network **Storage** Security Monitor Maintenance

Storage Pools Volumes Snapshots iSCSI Disks RDX QuikStor

Create a Storage Pool for unit **Head Unit (XSD 40)** with the following settings.

Storage Pool configuration is based on these settings:

Pool	Estimated Available Space	Percent of Storage
Data Pool	6.52 TB	80%
Snapshot Pool	1.63 TB	20%

Storage Pool Name: Disks Detected: 4 · Available Disk Slots: 4 (Unit is fully populated.)

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.

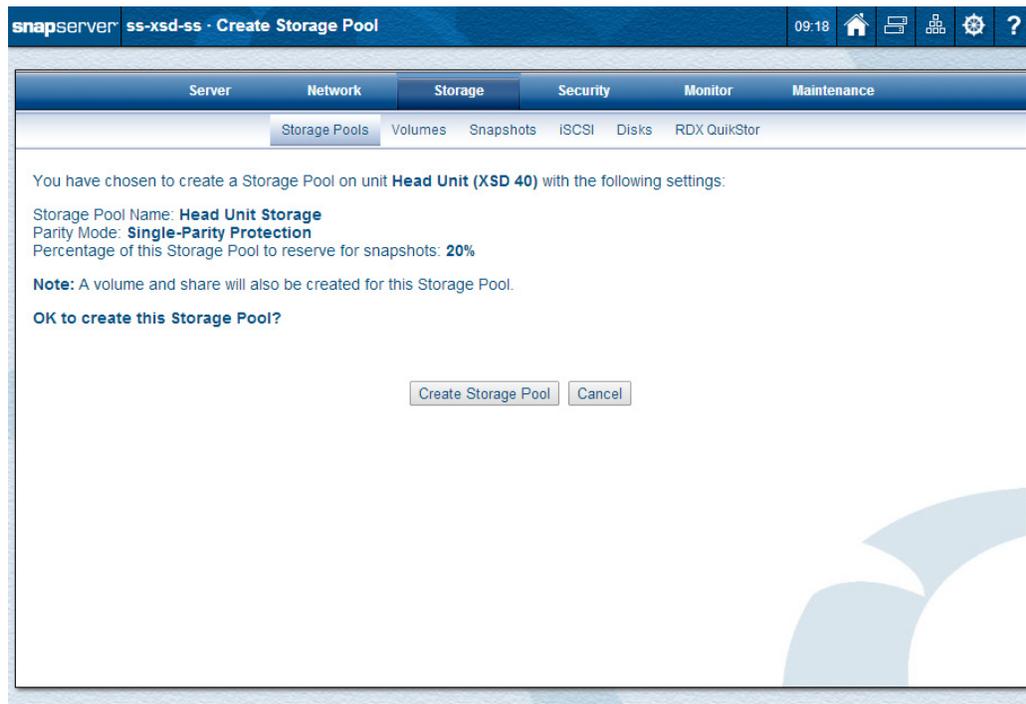
Snapshot Pool

If you plan on using snapshots, it is recommended that you reserve at least 20% of your Storage Pool for snapshots. You can adjust the snapshot pool percentage at a later time; however to increase it, you will first need to add more capacity to this Storage Pool.

Percentage of this Storage Pool to reserve for snapshots:

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** at the bottom.

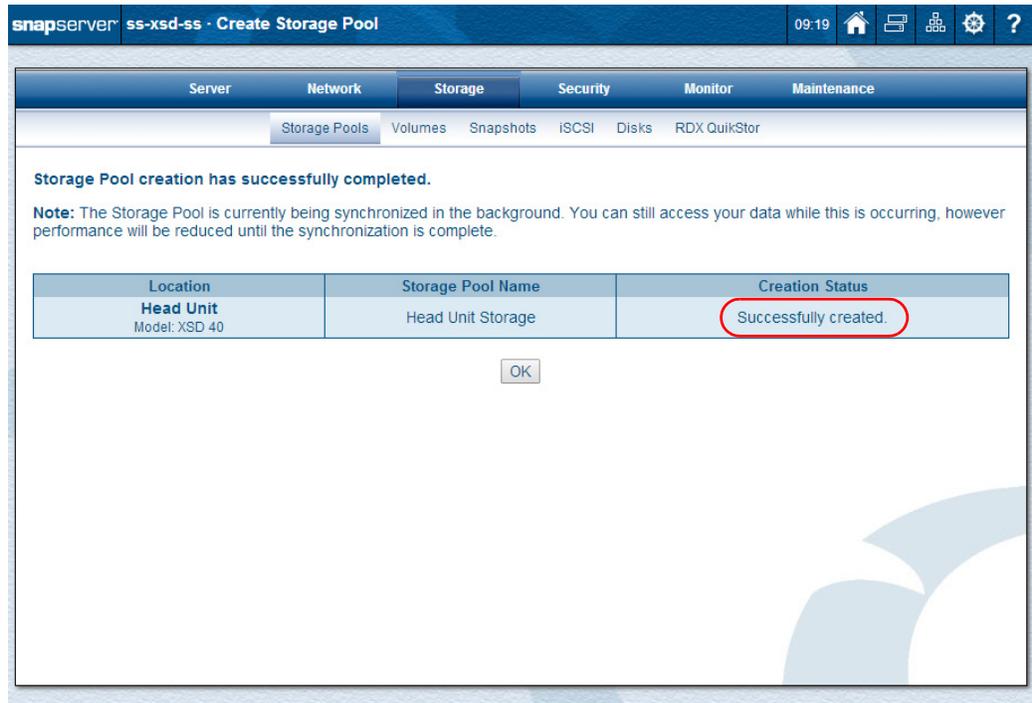
4. At the confirmation page, verify your selections and, if everything is correct, click **Create Storage Pool** again.



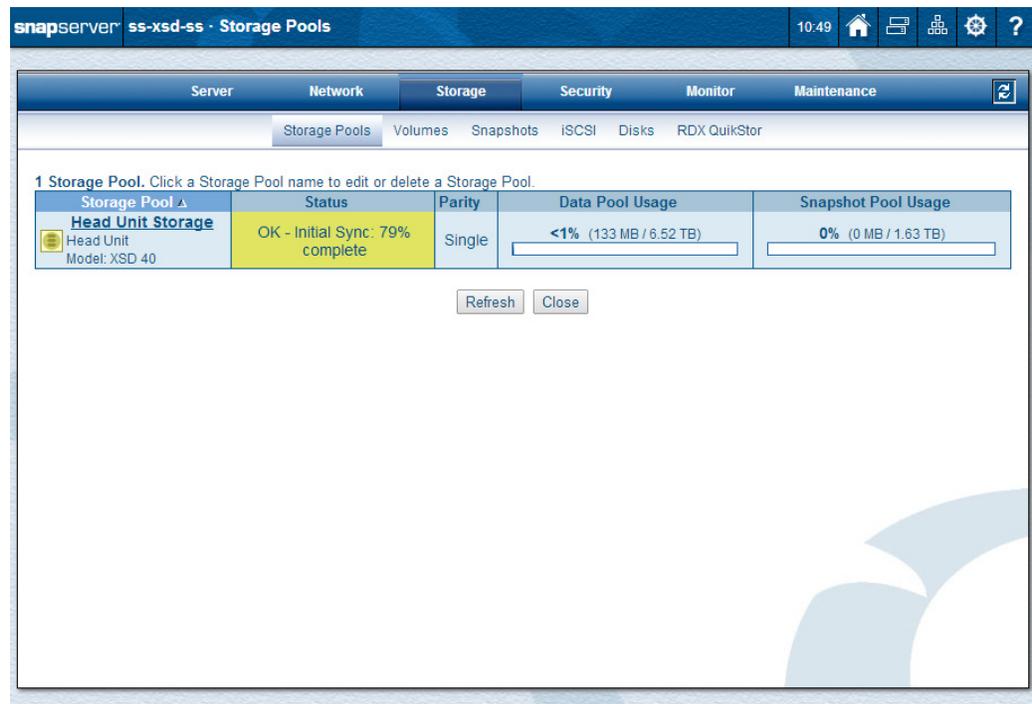
NOTE: If any disk in a storage pool has previously been used in a different system, a warning appears that it will be reformatted and all data on the disk will be deleted. If any disk has size characteristics that are incompatible with the other disks in a storage pool, a message appears that this disk 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 successfully created, click **OK** to continue.



6. You are returned to the **Storage Pools** page where the **Status** shows a resync underway.

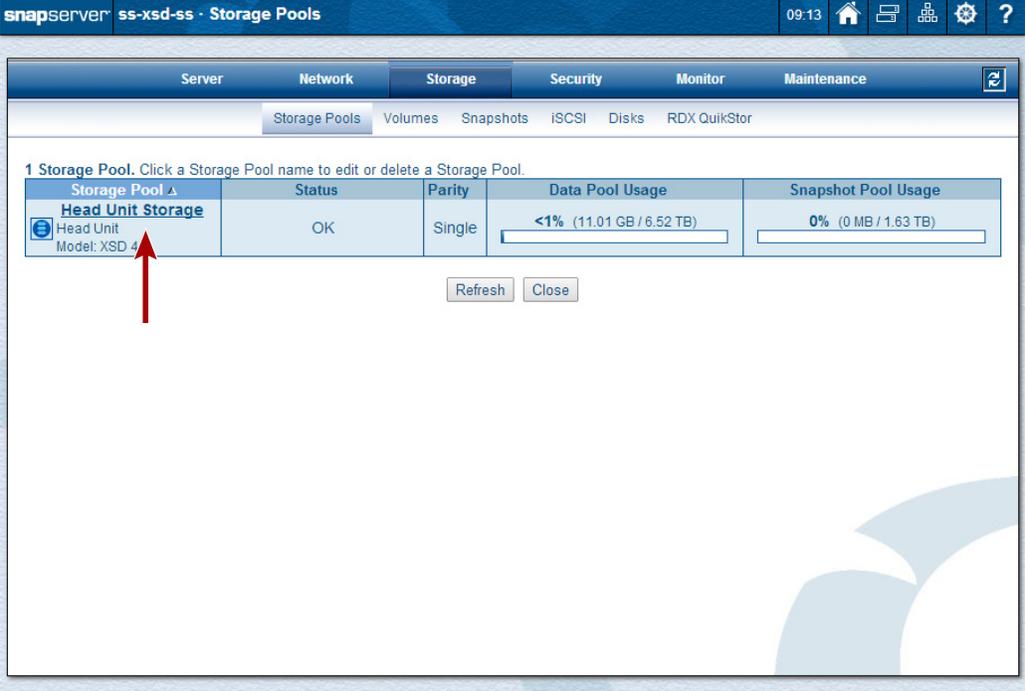


Click **Refresh** now and then to see the current **Status** and to determine when the resync is complete.

 **IMPORTANT:** 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 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.



The screenshot shows the SnapServer web management interface. The top navigation bar includes 'snapserver', 'ss-xsd-ss', and 'Storage Pools'. The main content area has tabs for 'Storage Pools', 'Volumes', 'Snapshots', 'iSCSI', 'Disks', and 'RDX QuikStor'. Below the tabs, there is a table with the following data:

Storage Pool	Status	Parity	Data Pool Usage	Snapshot Pool Usage
Head Unit Storage Head Unit Model: XSD 4	OK	Single	<1% (11.01 GB / 6.52 TB)	0% (0 MB / 1.63 TB)

Below the table, there are 'Refresh' and 'Close' buttons. A red arrow points to the 'Head Unit Storage' link in the table.

After you click the head unit storage pool name in the Storage Pool list, the properties page for the head unit is shown:

The screenshot shows the 'Storage Pool Properties' page for a 'Head Unit Storage' pool. The page is titled 'snapserver ss-xsd-ss - Storage Pool Properties' and includes a navigation bar with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. The 'Storage' tab is active, and the 'Storage Pools' sub-tab is selected. A table lists the storage pool details:

Storage Pool	Status	Parity	Data Pool Usage	Snapshot Pool Usage
Head Unit Storage Head Unit Model: XSD 40	OK	Single	<1% (11.01 GB / 6.52 TB)	0% (0 MB / 1.63 TB)

Below the table, the 'Storage Pool Name' is 'Head Unit Storage'. The 'Parity Mode' section shows 'Single-parity protection' selected, with a note that dual-parity protection is not available. The 'Snapshot Pool' section shows a note that the snapshot pool percentage cannot be increased and a dropdown menu set to 20%. At the bottom, there are buttons for 'OK', 'Refresh', 'View Disks', 'Delete Storage Pool', and 'Cancel'.

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

The screenshot shows the 'Storage Pools' page for an 'Exp Unit 1 Storage' pool. The page is titled 'snapserver ss-xsd-ss - Storage Pools' and includes a navigation bar with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. The 'Storage' tab is active, and the 'Storage Pools' sub-tab is selected. A table lists the storage pool details:

Storage Pool	Status	Parity	Data Pool Usage	Snapshot Pool Usage
Exp Unit 1 Storage Expansion Unit 1 Model: SE DX	OK	Single	<1% (133 MB / 1.44 TB)	0% (0 MB / 367.62 GB)

Below the table, the 'Storage Pool Name' is 'Exp Unit 1 Storage'. The 'Parity Mode' section shows 'Single-parity protection' selected, with a note that dual-parity protection requires adding another disk. The 'Snapshot Pool' section shows a note that increasing the snapshot pool percentage requires adding more capacity and a dropdown menu set to 20%. At the bottom, there are buttons for 'OK', 'Refresh', 'View Disks', 'Delete Storage Pool', and 'Cancel'.

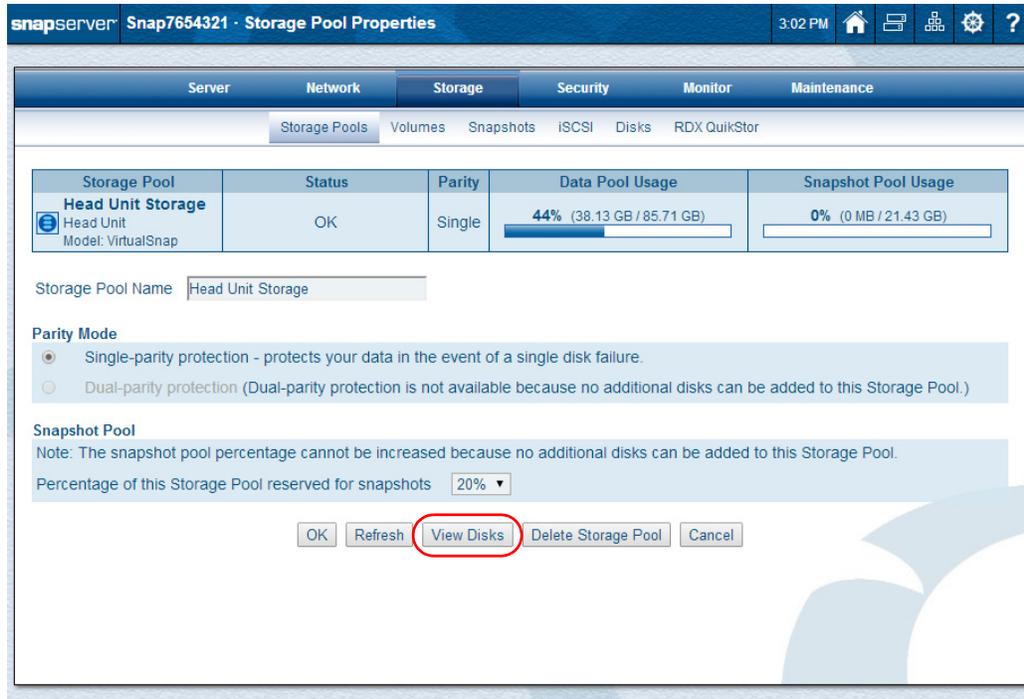
At the **Storage Pool Properties** page, you can edit these options:

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 on page 94.</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 in Chapter 7.</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).



Click **View Disks** to display the **Storage Pool Disks** page and see all the disk drives:



NOTE: Any disk that is incompatible is shown with a highlighted message.

Storage Pool Deletion



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

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**.

The screenshot shows the 'Storage Pool Properties' page for a storage pool named 'Head Unit Storage'. The page is part of the SnapServer interface, with a breadcrumb trail: 'Storage Pools > Volumes > Snapshots > iSCSI > Disks > RDX QuikStor'. The main content area displays the following information:

Storage Pool	Status	Parity	Data Pool Usage	Snapshot Pool Usage
Head Unit Storage Head Unit Model: XSD 40	OK	Single	<1% (11.01 GB / 6.52 TB)	0% (0 MB / 1.63 TB)

Storage Pool Name:

Parity Mode

- Single-parity protection - protects your data in the event of a single disk failure.
- Dual-parity protection (Dual-parity protection is not available because no additional disks can be added to this Storage Pool.)

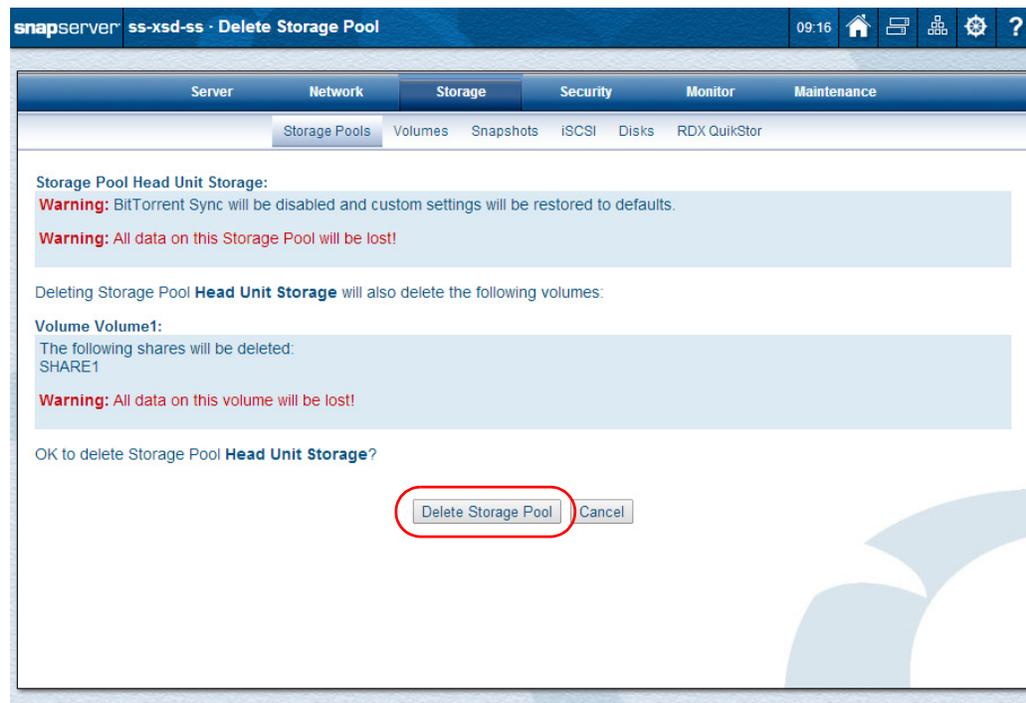
Snapshot Pool

Note: The snapshot pool percentage cannot be increased because no additional disks can be added to this Storage Pool.

Percentage of this Storage Pool reserved for snapshots:

At the bottom of the page, there are five buttons: 'OK', 'Refresh', 'View Disks', 'Delete Storage Pool', and 'Cancel'. The 'Delete Storage Pool' button is highlighted with a red circle.

- 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 [Create a Storage Pool on page 94](#).

Parity Management

Parity is used to achieve redundancy in the SnapServer. If a drive in the array fails, remaining data on the other drives can be combined with the parity data to reconstruct the missing data.

- **Single Parity** – Protects your data in the event of a single disk failure.
- **Dual Parity** – Uses more disk space than single parity, but protects your data in the event of up to two disk failures.

Parity is usually set when creating a new storage pool. It can also be changed 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](#) in [Appendix C](#).

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.

Add a Disk Drive to Upgrade Parity

To increase the parity protection of the storage pool, new disk drives are 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 Unit 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 on page 95 .

When a new disk drive is added, the Administration pages display a message banner that new drives were detected. At the **Storage > Storage Pools** page, the same message is shown with a clickable link. Clicking the link opens the **Storage Pool Properties** page, which allows you to change the parity and snapshot settings for the storage pool, taking advantage of the additional space.

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

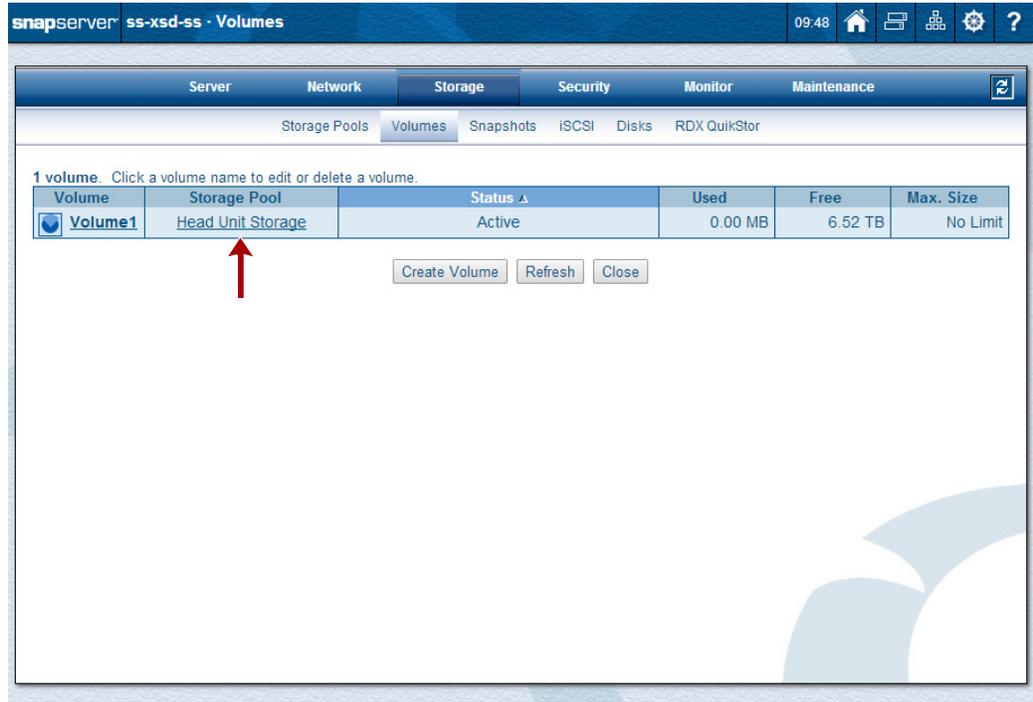
There are no separate spare or global spare disk drives when using the DynamicRAID option. With single parity, if a disk drive fails, a warning is issued and the system reverts to degraded mode with no protection, so a second drive failure will cause the system to fail. With dual parity, if two disk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

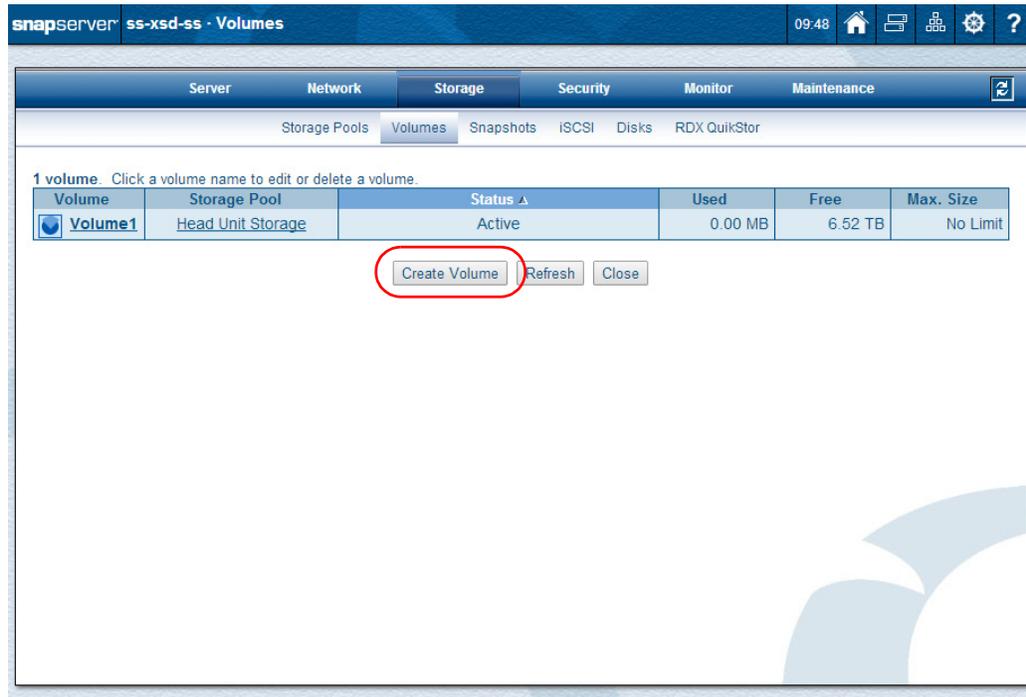
GuardianOS supports multiple volumes in a storage pool. During the initial creation of your DynamicRAID storage pool, an initial volume was also created. To view that volume (and create other volumes if needed), navigate to **Storage > Volumes**. To access the **Properties** page for a volume, click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



Clicking the storage pool name in the second column takes you to the **Storage Pool Properties** page.

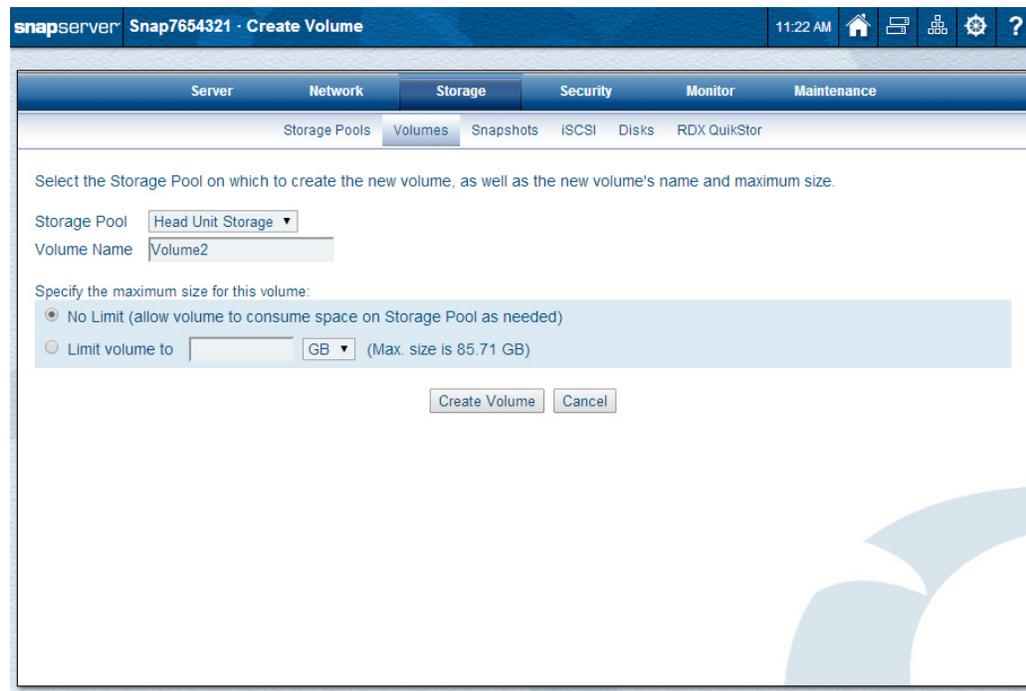
Volume Creation

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



Create a New Volume

1. Navigate to **Storage > Volumes**.
2. Click **Create Volume**.



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 on page 258.](#))

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 **Create Volume** 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 **Create Share**. See [Shares on page 179](#) for more information about creating shares.

Volume Properties

Click the volume name to show the **Volume Properties** page where the volume settings are edited.

The screenshot displays the 'Volume Properties' page for 'Volume1'. The page has a navigation bar with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. Under the 'Storage' tab, there are sub-tabs for Storage Pools, Volumes, Snapshots, iSCSI, Disks, and RDX QuikStor. A table lists the volume details:

Volume	Storage Pool	Status	Used	Free	Max. Size
Volume1	Head Unit Storage	Active	0.00 MB	6.52 TB	No Limit

Below the table, the 'Volume Name' is 'Volume1'. The 'Specify the maximum size for this volume:' section has two radio button options: 'No Limit (allow volume to consume space on Storage Pool as needed)' (selected) and 'Limit volume to' (with a text input field, a unit dropdown set to 'MB', and '(Max. size is 6.52 TB)'). At the bottom are buttons for 'OK', 'Refresh', 'Delete Volume', and 'Cancel'.

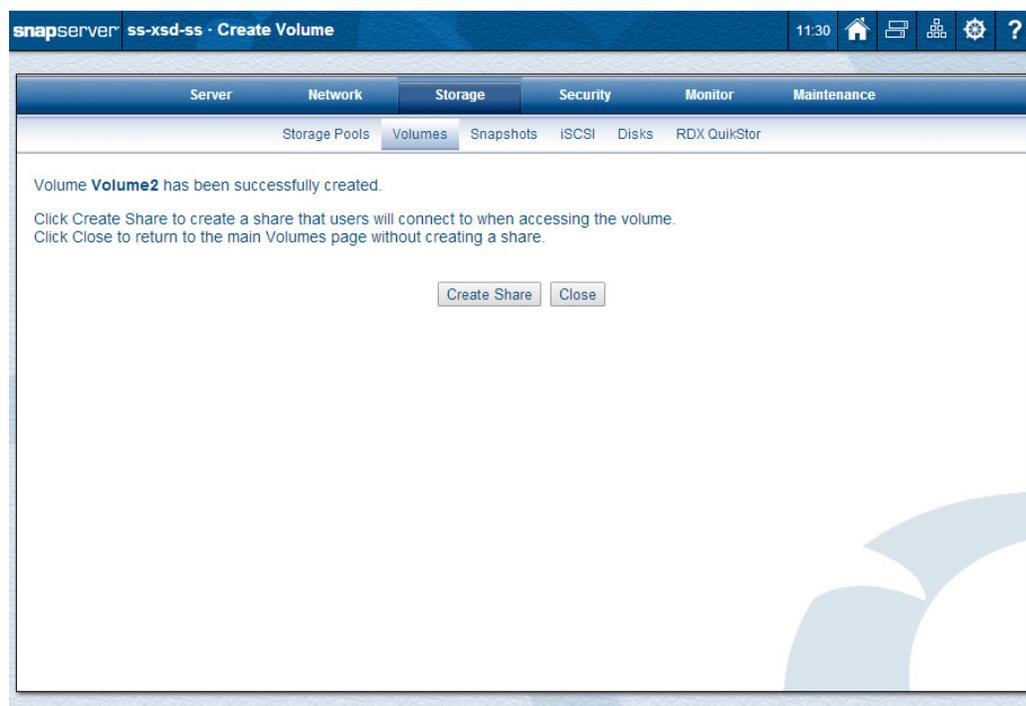
Edit Volume Properties

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 on page 258](#).)

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.

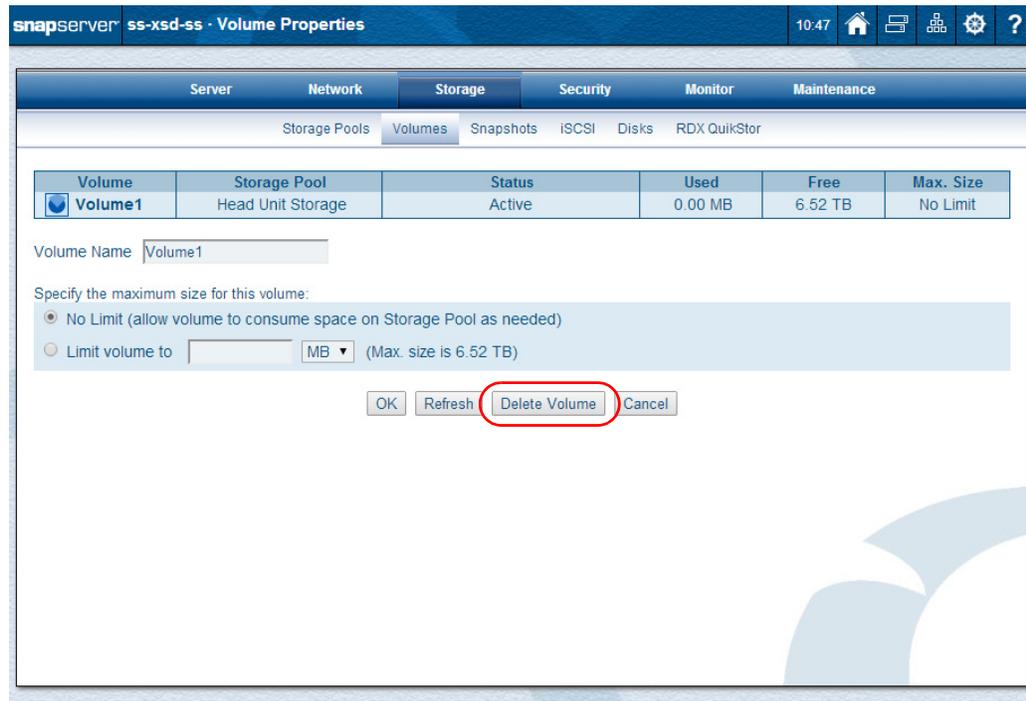
4. When you are done, click **OK**.
5. When the process is complete, a confirmation screen is shown.



You can click **Create Share** to go to the **Create Share** page under **Security** and start the creation process directly.

Volume Deletion

To delete a volume, go to the **Storage > Volumes > Volume Properties** page.



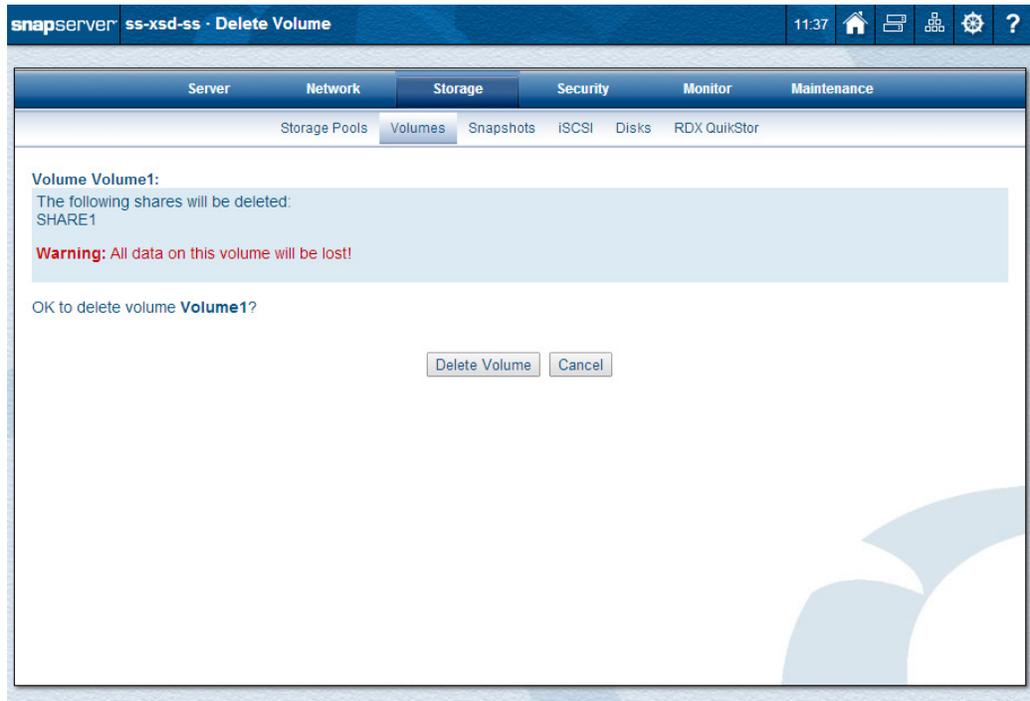
Delete a Volume

1. Navigate to **Storage > Volumes**.
2. Click the **volume name** in the table to go to **Volume Properties**.
3. 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 [DynamicRAID Storage](#) in [Chapter 5](#). For other storage features, see [Other Storage Options](#) in [Chapter 7](#).

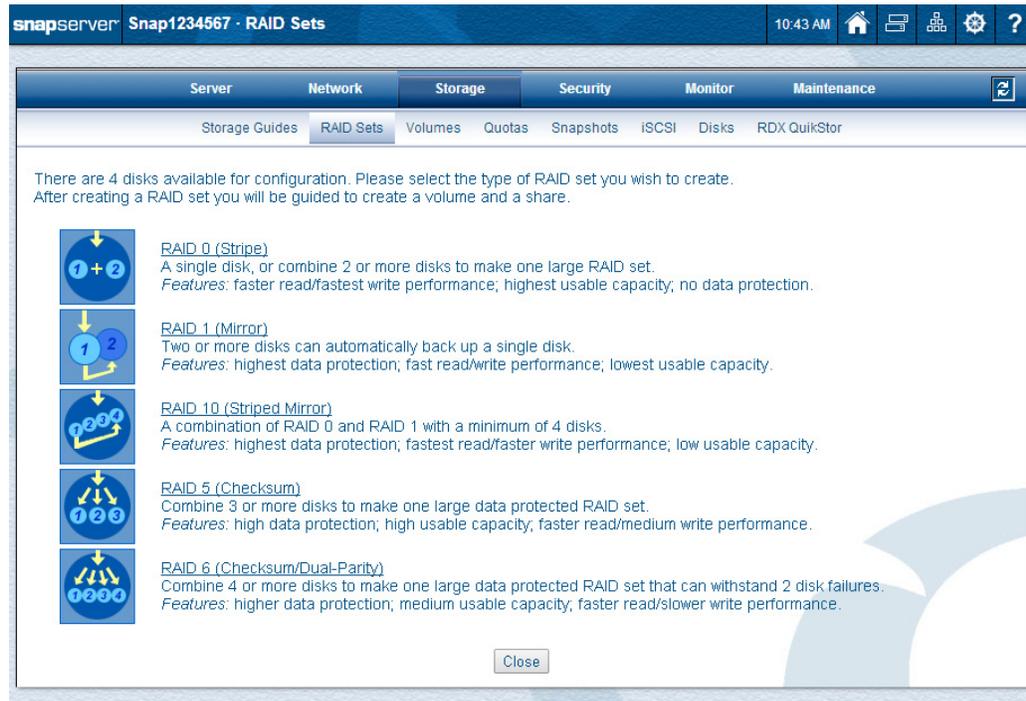
Topics in Traditional RAID Storage

- [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, the options will be grayed out and unavailable.



The basic steps for storage configuration are:

Step 1: Create a RAID set.

Step 2: Create a volume on the new RAID set.

Step 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.

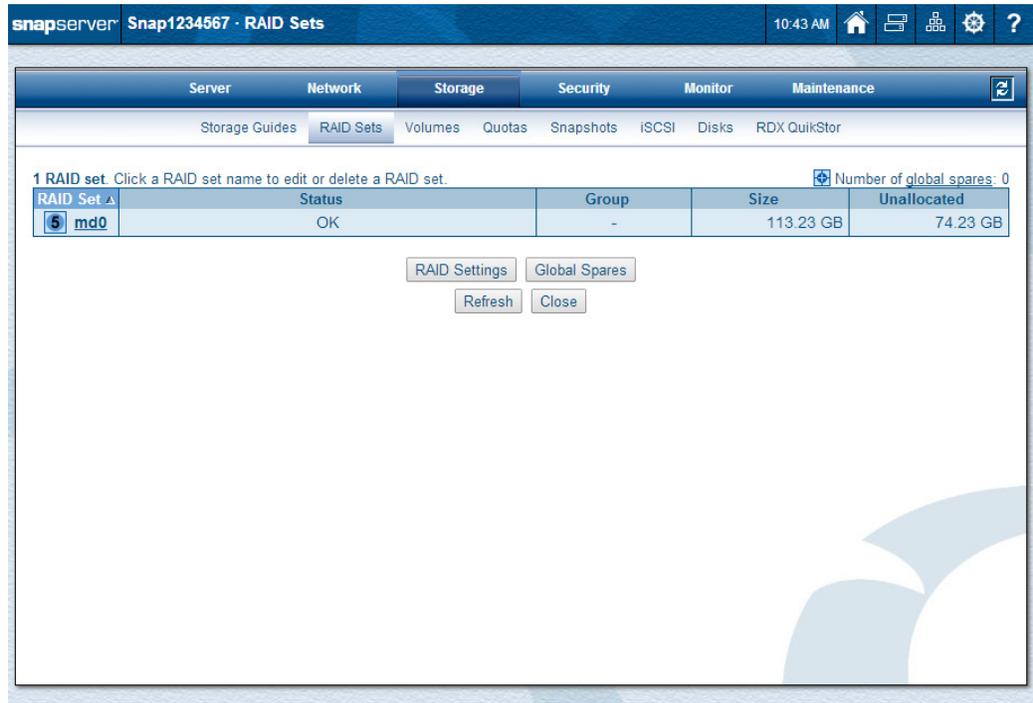
Local and Global Spares

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 style="text-align: center;">  Local Spare  Global Spare (GS) </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:

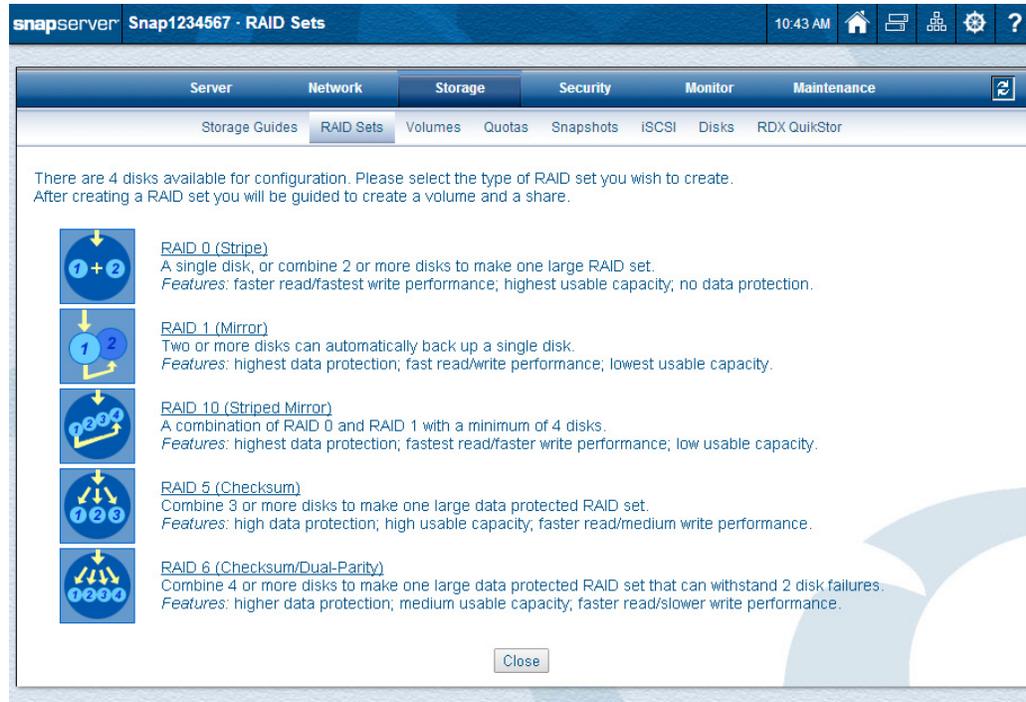
- **Create RAID Sets** – If unassigned drives exist, a new RAID set can be created by launching a wizard.
- **Group RAID Sets** – If more than one RAID set exists, they can be grouped together.
- **Change RAID Settings** – Change two RAID options:
 - Enable/disable automatic incorporation of unused disks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

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.

- 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 capacities 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 160](#) 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
Drive 1	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

- 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 on page 114](#).

5. Click Next.

The screenshot shows the SnapServer GUI for creating a RAID set. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'Storage' section is active, with sub-tabs for 'Storage Guides', 'RAID Sets', 'Volumes', 'Quotas', 'Snapshots', 'iSCSI', 'Disks', and 'RDX QuikStor'. The main content area displays the RAID configuration details:

- Create RAID set: **RAID 5 (Checksum)**
- RAID Set Name: **md0**
- Total Storage Capacity: **113.24 GB**

Below this, a message states: "You have chosen the following 4 disks for this RAID set." A table lists the selected disks:

Disk	Location	Status	Usable Space	Usage	Type
<input checked="" type="checkbox"/> 50 GB SAS	Head Unit, disk 1	OK	37.75 GB	100%	Active Member
<input checked="" type="checkbox"/> 50 GB SAS	Head Unit, disk 2	OK	37.75 GB	100%	Active Member
<input checked="" type="checkbox"/> 50 GB SAS	Head Unit, disk 3	OK	37.75 GB	100%	Active Member
<input checked="" type="checkbox"/> 50 GB SAS	Head Unit, disk 4	OK	37.75 GB	100%	Active Member

At the bottom of the configuration area, there are three buttons: 'Back', 'Next' (highlighted with a red circle), and 'Cancel'.

6. Verify your configuration, then click **Next** to create the RAID.

A message appears confirming the successful creation of the RAID set. It details how much storage space is available. The RAID will be syncing in the background.

The screenshot shows the SnapServer GUI after the RAID set has been created. The top navigation bar is the same as in the previous screenshot. The main content area displays a confirmation message:

- A new RAID Set named **md0** has been created, providing approx. **113.24 GB** of storage space.
- Note:** The RAID array is currently being synchronized in the background. You can still access your data while this is occurring, however performance will be reduced until the synchronization is complete.

Below the message, a text block states: "You must create a volume on this RAID set before any data can be placed on it. You may choose to create a volume at a later time using the Create Volume page." At the bottom, there are two buttons: 'Create Volume Now' (highlighted) and 'Create Volume Later'.

7. Before you can place any data on this RAID set, you must create a volume. You use the buttons on this page to choose whether you want to create the volume now or later:
 - Click **Create Volume Now** to create the volume now by following the procedure outlined in [Volume Creation on page 119](#).
 - Click **Create Volume Later** to be returned to the **RAID Sets** page. You will need to remember to go to the **Volumes** option at a later time to create the required volume.

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, which 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 **Group RAID** to show the **Group RAID Sets** page.

The screenshot displays the SnapServer RAID Sets management interface. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'Storage' section is active, showing 'RAID Sets', 'Volumes', 'Quotas', 'Snapshots', 'iSCSI', 'Disks', and 'RDX QuikStor'. A table lists 2 RAID sets:

RAID Set	Status	Group	Size	Unallocated
md0	OK	-	1.80 TB	1.80 TB
md1	OK	-	3.61 TB	3.61 TB

Below the table, there are buttons for 'Create RAID', 'Group RAID', 'RAID Settings', and 'Global Spares'. The 'Group RAID' button is highlighted with a red circle. There are also 'Refresh' and 'Close' buttons below the main action buttons.

1. Select (check) the **RAID sets** you want to include in the group and click **Next**.

You can group two RAID sets (of the same type) together to expand the capacity of both RAID sets. The new RAID group will be created with the name you specify below.

RAID Group Name:

Select two RAID sets of the same type and click Next.

RAID Set	Status	Type	Size	Unallocated
<input checked="" type="checkbox"/> 5 md0	OK	RAID 5 (Checksum)	1.80 TB	1.80 TB
<input checked="" type="checkbox"/> 5 md1	OK	RAID 5 (Checksum)	3.61 TB	3.61 TB

2. At the confirmation page, click **Create RAID Group** to complete the process.

The RAID group will be created with the following settings.

RAID Group Name: **group0**
 RAID Group Size (approx.): **5.41 TB**

RAID sets for new RAID group.

RAID Set	Status	Type	Size	Unallocated	Preserve Volumes
5 md0	OK	RAID 5 (Checksum)	1.80 TB	1.80 TB	Yes
5 md1	OK	RAID 5 (Checksum)	3.61 TB	3.61 TB	Yes

3. At the primary **RAID Sets** page, click the **group name** to see the details of the group.

The screenshot shows the SnapServer interface for a RAID Group named 'group0'. The page has a navigation bar with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. Under the Storage tab, there are sub-tabs for Storage Guides, RAID Sets, Volumes, Quotas, Snapshots, iSCSI, Disks, and RDX QuikStor. The main content area displays two tables:

RAID Group	Status	Type	Size	Unallocated
group0	Active	RAID 5 (Checksum)	5.41 TB	5.41 TB

Below this is a sub-table for RAID sets for the group 'group0':

RAID Set	Status	Type	Size	Unallocated
md0	OK	RAID 5 (Checksum)	1.80 TB	-
md1	OK	RAID 5 (Checksum)	3.61 TB	-

At the bottom of the RAID sets table, there are buttons for Refresh, Delete RAID Group, Add RAID, and Close.

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 in two tables:

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. See Disks on page 156 .
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.

Label	Description
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

NOTE: When a new undiscovered expansion unit is detected at start up, an alert message is shown detailing the number of units found and if they have existing storage.

In a common scenario, a SnapServer is nearing full utilization. The administrator decides to add an expansion unit. The administrator creates the same RAID type 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. By clicking the **Expand Volume** button that appears, these things are done automatically (see [Expand Volume Capacity on page 119](#)).

Expansion unit RAID sets are created in the same way as head unit RAID Sets. See [Create RAID Sets on page 105](#).

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 on each of the expansion units, then group the two of them together. The resulting RAID 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 on page 109](#).

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 on page 117](#).

Snapshot Pools are Combined

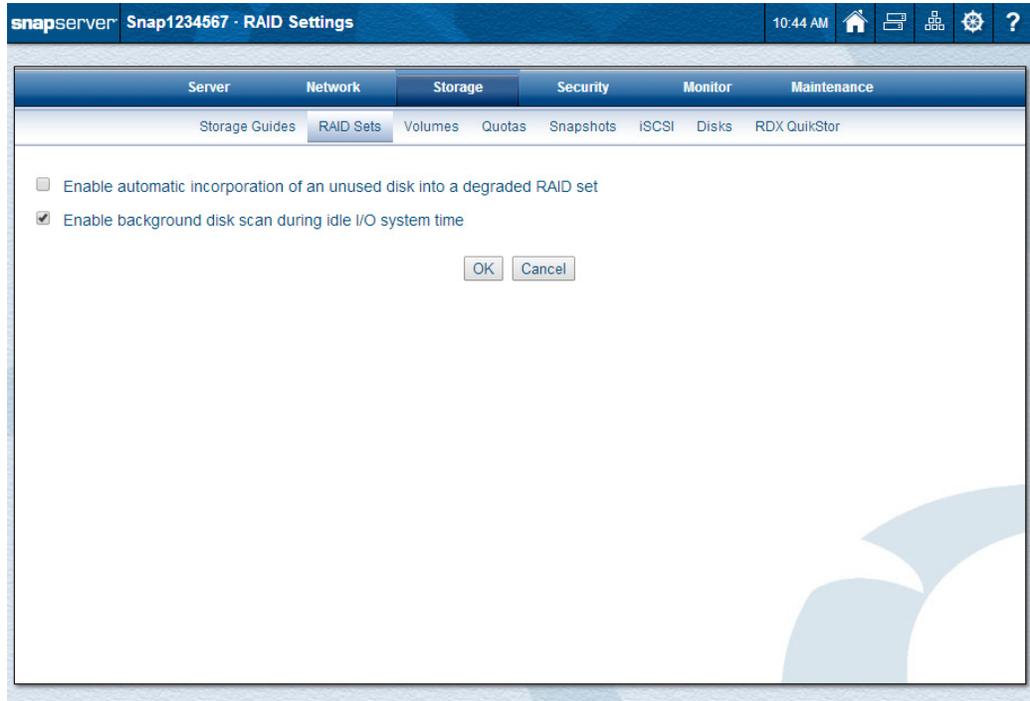
When two RAID sets are grouped, the size of the resulting snapshot pool is the sum of each RAID set's formerly separate snapshot pools.

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 **RAID Settings** 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 **RAID Settings**.

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, the scan process 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 **RAID Settings**. 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 RAID sets that are degraded, syncing, or rebuilding.

Manage Global Spares

A **spare** is a unused 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 on page 104](#)).

Click **Global Spares** to view all the disks either available for use or in use as global spares.

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 type="checkbox"/> 2.73 TB SATA	Head_Unit_1_disk 7	OK	2.72 TB
<input checked="" type="checkbox"/> 931.51 GB SATA	Exp_Unit_1_disk 1	OK	919.26 GB
<input type="checkbox"/> 931.51 GB SATA	Exp_Unit_1_disk 2	OK	919.26 GB
<input type="checkbox"/> 3.64 TB SATA	Exp_Unit_1_disk 4	OK	3.63 TB

OK Cancel

To enable a disk as a global spare, check the box 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 box next to the disk and click **OK**.

Edit 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.

The screenshot shows the RAID Set Properties page for RAID set md0. The RAID set is in an OK status, belongs to group0, and has a size of 1.80 TB. It has 3 active members and 0 spares. The member disks are three 931.51 GB SATA disks at Head Unit, disk 1, 2, and 3, all with OK status and 919.26 GB usable space. A note indicates that the RAID set is a member of a RAID Group and cannot be deleted individually. Buttons for Refresh, Add Disk, and Close are visible.

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.

Label	Description
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).

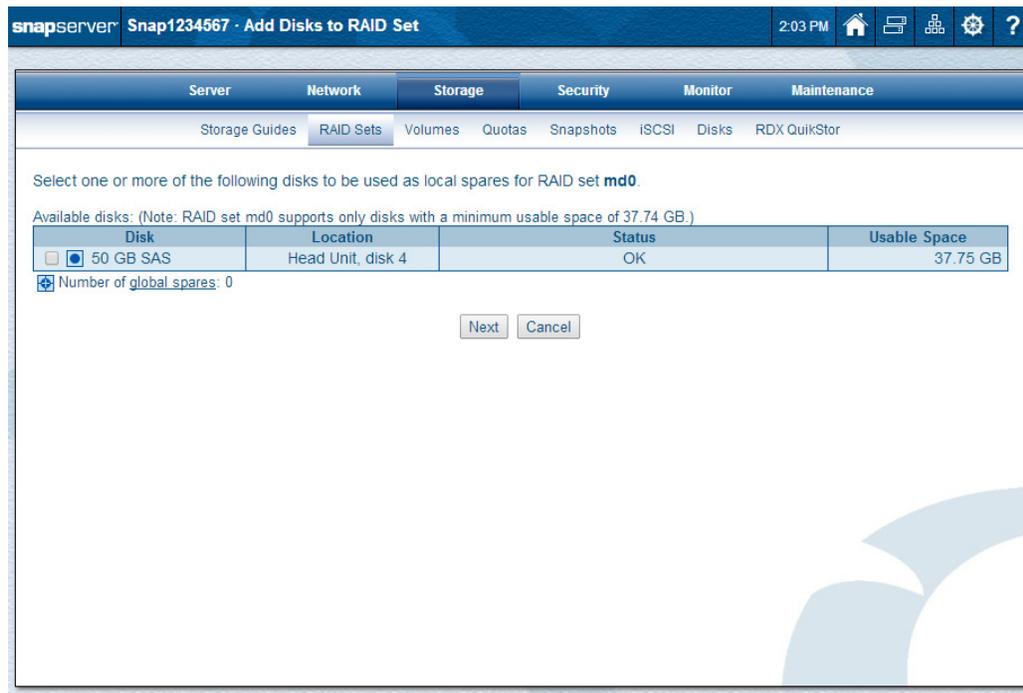
Remove a RAID Drive

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 warning of RAID running in a degraded mode (with 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 **Add Disk** at the bottom of the **RAID Set Properties** page displays a table of available disk drives.



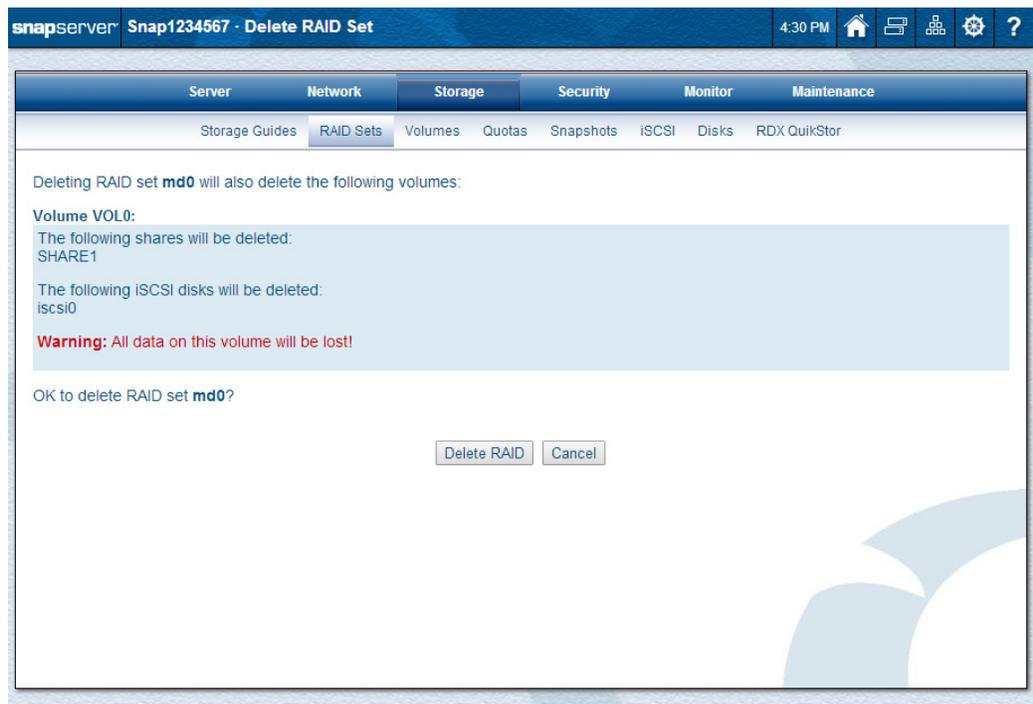
Check one or more boxes to add disks to the RAID set and click **Next** for the confirmation page. Adding disks is limited based on the type of RAID it is being made a member of:

- 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

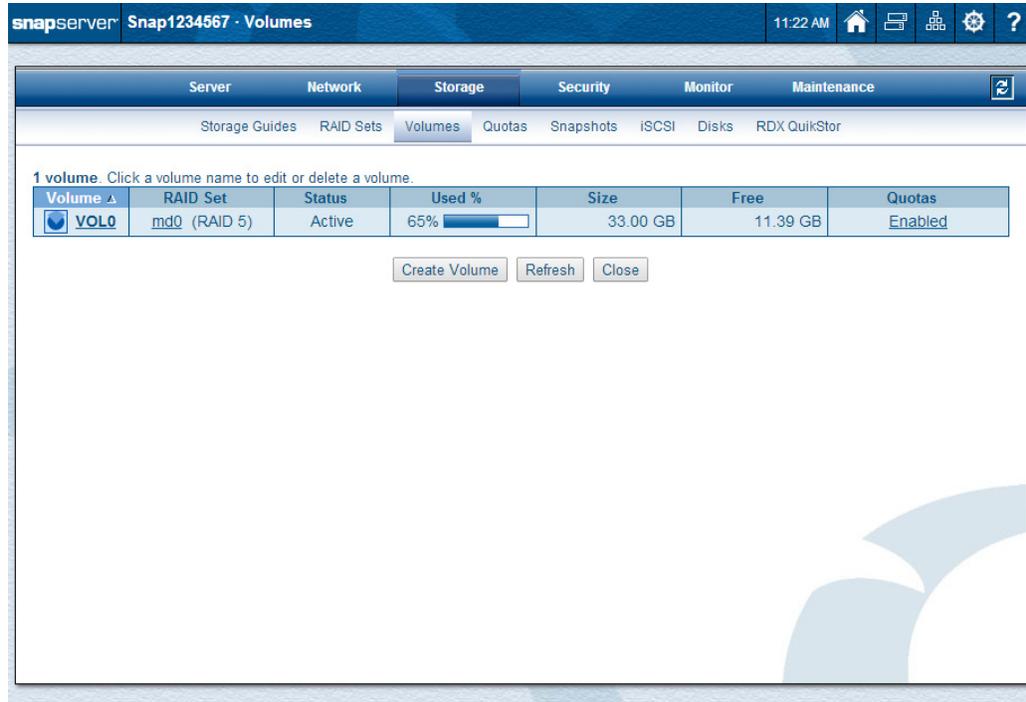
Click **Delete RAID** at the bottom to completely delete the RAID set.



This also deletes the Volume and Shares, including any data on them. Click **Delete RAID** 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 (by clicking the name).
- Enable/disable quotas on the volume (click the **Quotas** link in far right column).

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 on page 141](#).

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 153](#).

Volume Creation

To create a volume on a RAID set, click **Create Volume** on the main **Volumes** page. When manually creating a RAID set, at the end you can click **Create Volume Now** to launch the same **Create Volume** page.

Select the RAID set on which to create the new volume, as well as the new volume's name and size.

RAID Set: md0 - RAID 5 (74.23 GB available for volume)

Volume Name: VOL1

Volume Size: 74.23 GB

Enable Write Cache (not recommended without a configured, online UPS device)

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.)

Percentage of this volume's size to add to the snapshot pool for RAID set md0: 20%

Current allocation of RAID set md0

Total size:	113.23 GB
Allocated size (for volumes):	33.00 GB
Snapshot pool size:	6.00 GB

Create Volume Cancel

Create A New Volume (and Share)

1. Navigate to **Storage > Volumes** and click **Create Volume**.
2. Configure the **settings** for the new volume:

Label	Description
RAID Set	Use the drop-down menu to select the RAID to be used for the volume.
Volume Name	Enter a name for the volume or accept the default (VOL0). You can use up to 20 alphanumeric characters or hyphens (but not spaces)
Volume Size	Enter the size you want for the volume or accept the default of the full size of the RAID. Use the drop-down list to choose the appropriate unit.
Enable Write Cache	If you have a configured, online UPS device in use, check the box to turn on write caching.
Snapshot Percentage	Use the drop-down selector to choose the percentage you want to use for snapshots.

3. Click **Create Volume**.

- At the confirmation page, review the **settings** and click **Create Volume** again to start the configuration.

To prevent data loss, you are cautioned if the write caching option was enabled without the required UPS online. Click **Cancel** to return to the **Create Volume** settings page to make changes.



IMPORTANT: If you click **Cancel** to return to the **Create Volume** settings page, note that the settings revert to the defaults with the **Volume Size** reset to the maximum space.

The screenshot shows the SnapServer web interface for 'Snap1234567 - Create Volume'. The 'Storage' tab is active, and the 'Volumes' sub-tab is selected. The page displays the following information:

You have chosen to create a new volume with the following properties:

- RAID Set: **md0**
- Volume Name: **VOL1**
- Volume Size: **19.2 GB**
- Enable Write Cache: **Yes**

Current allocation of RAID set md0

Total size:	113.23 GB
Allocated size (for volumes):	33.00 GB
Snapshot pool size:	6.00 GB

Allocation of RAID set md0 after volume is created

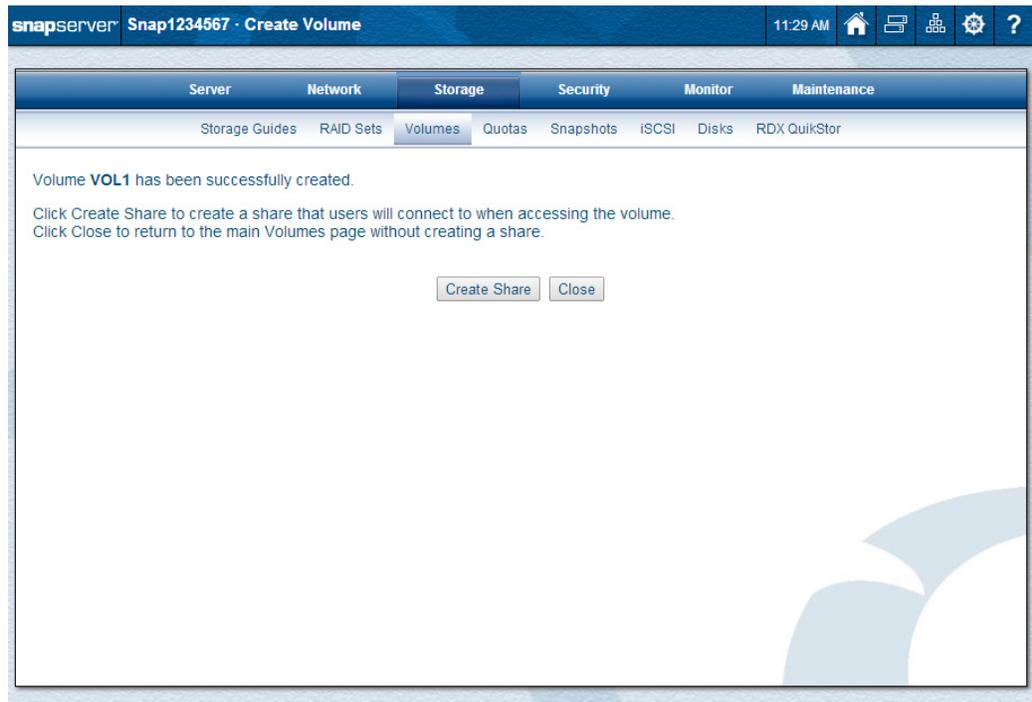
Total size:	113.23 GB
Allocated size (for volumes):	52.20 GB
Unallocated size remaining:	50.23 GB
Snapshot pool size:	10.80 GB

Warning: You have chosen to enable the write cache for this volume. This may place the volume's data at risk during a power outage because currently your server has no configured, online UPS device.

OK to create this volume?

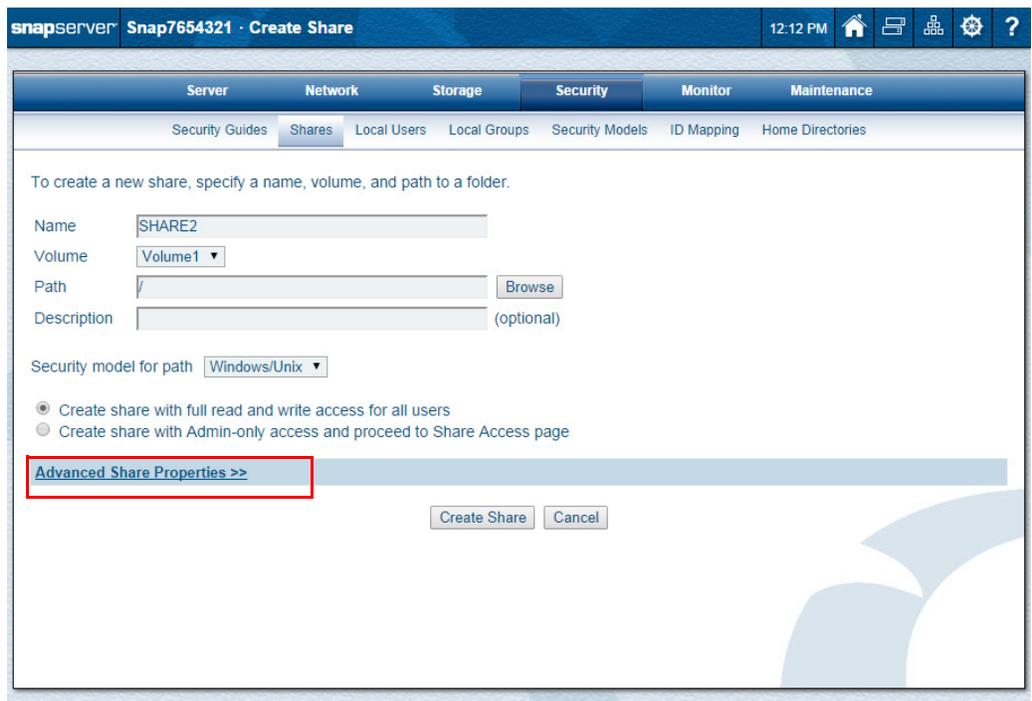
Buttons:

- At the successful volume creation page, click **Create Share** to provide access to this new volume.



This opens the **Security > Shares** option page so you can create a share pointing to this new volume.

- Enter the appropriate **data**, select the necessary **options**, and then click **Create Share**.



Click the **Advanced Share Properties** link to display additional options. See [Shares on page 179](#) for complete details.

7. Click **Create Share** again.

The share is automatically created and shown in the share table.

The screenshot shows the SnapServer GUI for the 'Shares' 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 main content area shows '2 shares' and a table with the following data:

Share	Volume	Path	Access	NFS Access	Protocols	Attributes
SHARE1	VOL0	/	Open	Default	SMB-NFS-AFP-HTTP-FTP	-
SHARE2	VOL1	/	Open	Default	SMB-NFS-AFP-HTTP-FTP	-

Attributes: H=Hidden, S=Has Snapshot Share, W=Web Root

Important Security Note: Share access for the NFS protocol is configured independently from share access for all other protocols. [View online help for more information.](#)

Buttons: Create Share, Refresh, Close

Volume Properties

By clicking a volume's name on the main **Volumes** 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 and configured).
- Delete the entire volume.

The screenshot shows the 'Volume Properties' page for volume VOL1. The page has a navigation bar with tabs for Server, Network, Storage (selected), Security, Monitor, and Maintenance. Under the Storage tab, there are sub-tabs for Storage Guides, RAID Sets, Volumes (selected), Quotas, Snapshots, iSCSI, Disks, and RDX QuickStor. A table lists the volume details:

Volume	RAID Set	Status	Used %	Size	Free	Time Created
<input checked="" type="checkbox"/> VOL1	md0 (RAID 5)	Active	<1%	19.20 GB	19.07 GB	2014-03-14 12:12 PM

Below the table, the 'Volume Name' is set to VOL1. The 'Volume Size' is 19.2 GB, with a dropdown menu set to GB and a note '(Max. size is 69.43 GB)'. The 'Enable Write Cache' checkbox is checked. At the bottom, there are buttons for OK, Refresh, Grow to Max. Size, Delete Volume, and Cancel.

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:

The screenshot shows the 'Volume Properties' page with a dialog box open. The dialog box contains the following text:

OK to grow volume **VOL1** from 19.2 GB to its maximum size of **69.43 GB**?
 (Note: This will leave RAID set md0 with no unallocated space.)

At the bottom of the dialog box, there are two buttons: 'Grow to Max. Size' and 'Cancel'.

- **Adding Unallocated Capacity** – If there is unallocated capacity remaining on the RAID, you can add this capacity to the volume:
 - Change the **Volume Size** to a size less than or equal to the maximum size of the volume and click **OK**.
 - Click **Grow to Max. Size** at the bottom and then click **Grow to Max. Size** again at the confirmation page.

NOTE: You cannot decrease the size of an existing volume. However, you can 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 a sufficient number of drives to create a new RAID set exists (or 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 Volume 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 147](#).

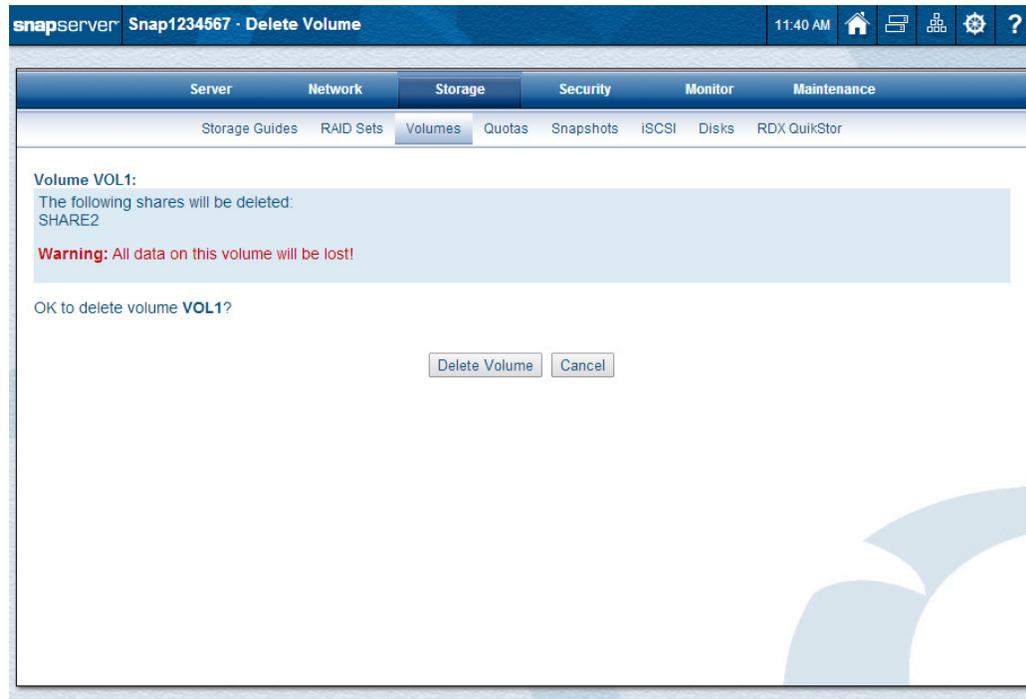
By default, write caching is enabled on all volumes. For systems that do not use a UPS device to help protect data during a power outage, or for applications that require synchronous writes to disk, write cache can be disabled on a volume-by-volume basis. When a volume's write cache is disabled, all 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.

To enable write caching, verify that a UPS device is attached and configured. Check the **Enable Write Cache** box and click **OK**. To disable, uncheck the box and click **OK**.

Delete a Volume

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

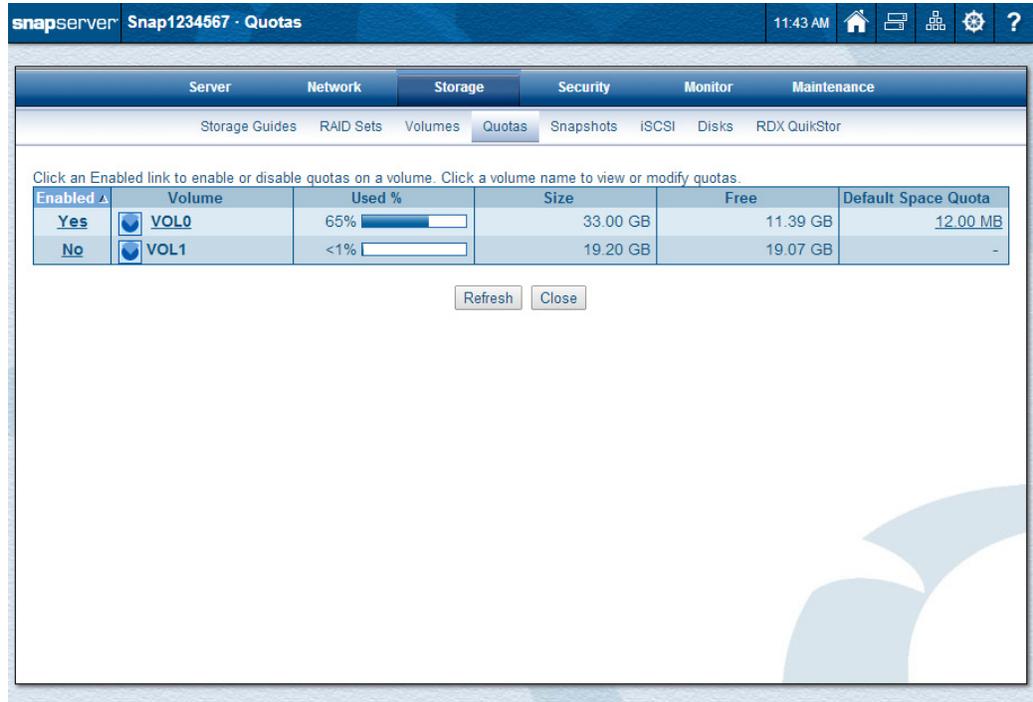
Previously installed CA Antivirus software and Snap EDR can reside on one or more volumes. If you delete a volume containing one of these applications, these components will be automatically moved to another volume, or deleted if no other volume or volumes of sufficient space are available. If deleted,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

CA Antivirus. After creating your new storage configuration, you can reenable the antivirus software by navigating to the **SnapExtensions** page and selecting **CA Antivirus**. On the next page, check the **Enable** box 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**.

Snap EDR. 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, LDAP group, or NIS group is currently consuming on the volume, allowing for precise tracking of usage patterns. You can set individual quotas for any LDAP group, NIS group, Windows domain, or local user known to the SnapServer. Group quotas are available only for LDAP and NIS groups.

For users and groups, there are no preassigned 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 the applicable user, LDAP, and NIS group quotas.

Quotas Page

The Quotas page shows if quotas are enabled on the volumes. From here you can:

- Enable or disable quotas.
- Modify the size of a quota.
- Change the way quotas are displayed.

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

When a quota is enabled on a volume, the view/modify name link in the **Volume** column is active and a status is shown in the **Default Space Quota** column.

Click an Enabled link to enable or disable quotas on a volume. Click a volume name to view or modify quotas.

Enabled	Volume	Used %	Size	Free	Default Space Quota
Yes	VOL0	65%	33.00 GB	11.39 GB	12.00 MB
No	VOL1	<1%	19.20 GB	19.07 GB	-

Refresh Close

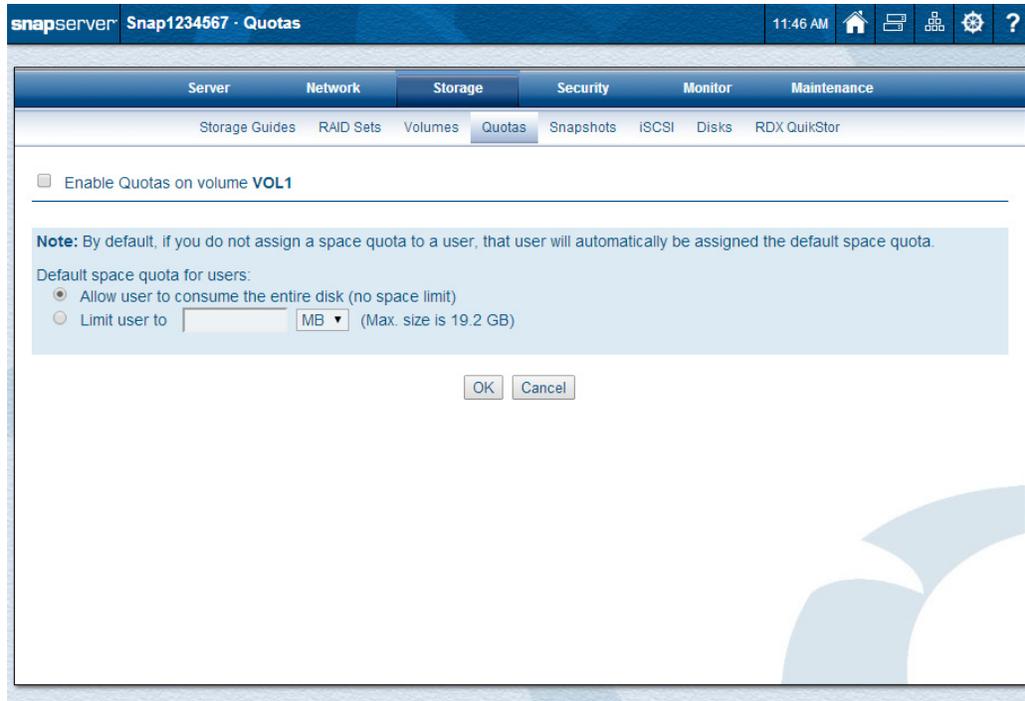
For enabled quotas, the **Default Space Quota** column shows one of the following:

- **An 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 Space Quota** amount doubles as a link to access the quota enable/disable page.

Enable/Disable Quotas

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. When you click the 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.

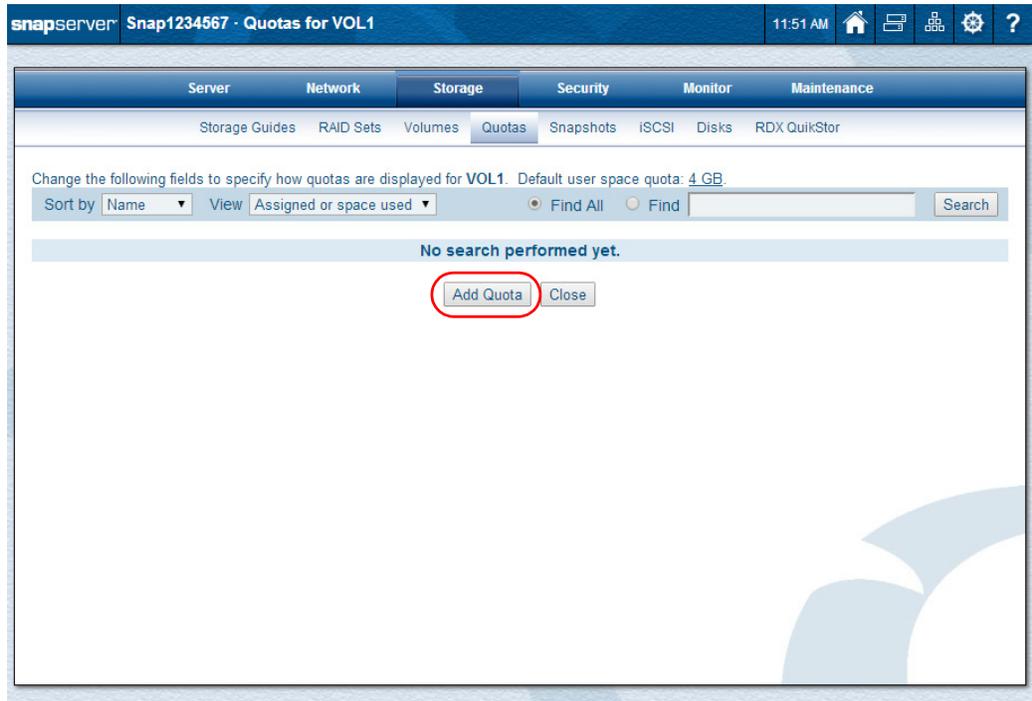
2. Select one of the two **default quota options** to set the quota applied to users who do not have individual quotas assigned to them.
3. Click **OK** (or **Cancel** to return without changes).

NOTE: The server may require a restart. If so, a warning message is displayed. Click **OK** to restart the server and continue.

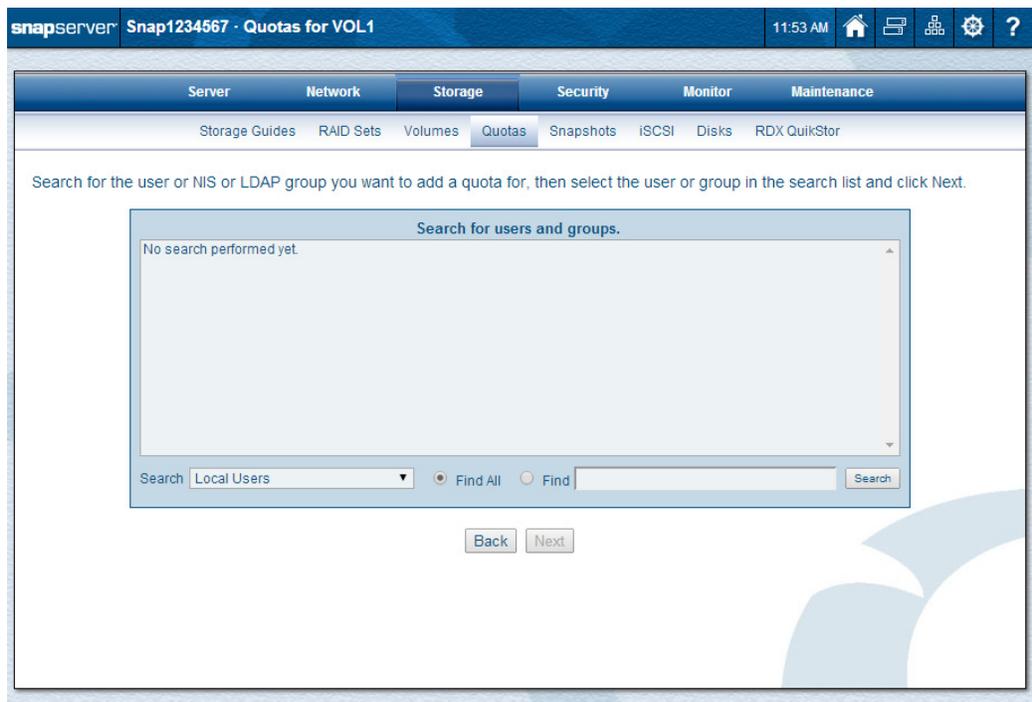
You are returned to the default **Quotas** page.

Add Quotas Wizard

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



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



3. To search for a user, LDAP group, or NIS group:
 - a. Select the **domain** from the **Search** drop-down list.

- b. Enter the **search string** (or select **Find All**).
- c. Click **Search** at the lower right.

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

The screenshot shows a search interface with a dropdown menu containing 'Cardinals User (A)'. Below the dropdown are two input fields: 'User Name' and 'Password'. A red arrow points to the 'User Name' field.

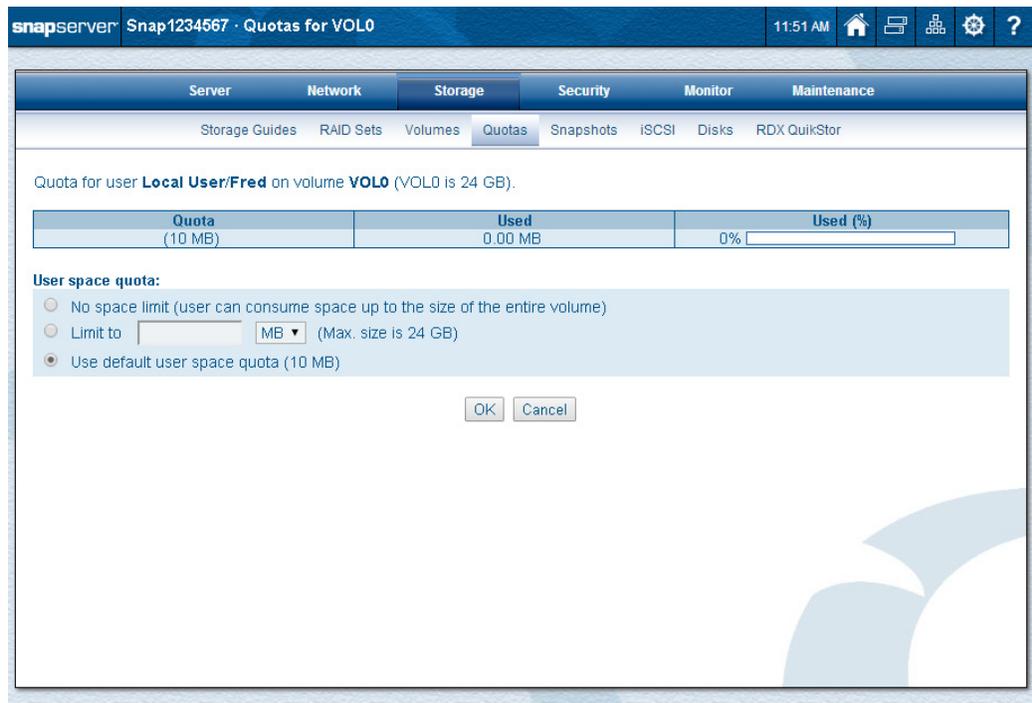
- Returned results will include all users, LDAP groups, and NIS groups whose name **begins** with the string entered in the **Search** field.
 - The search results returned may be limited. Fine tune your search by using a more specific string to return the names desired.
 - On the rare occasion you need 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**, an authentication prompt may be presented for the remote domain.
4. From the search results, select the **name** of the appropriate user, LDAP group, or NIS group, and click **Next**.

The screenshot shows the SnapServer Quotas for VOL0 interface. The search results are displayed in a table with the following data:

Local Users (2 found.)	
Fred	
Vicky	

Below the table, the search dropdown is set to 'Local Users' and the 'Next' button is visible.

- At the user quota properties page, select or enter the **quota** desired and click **OK**.



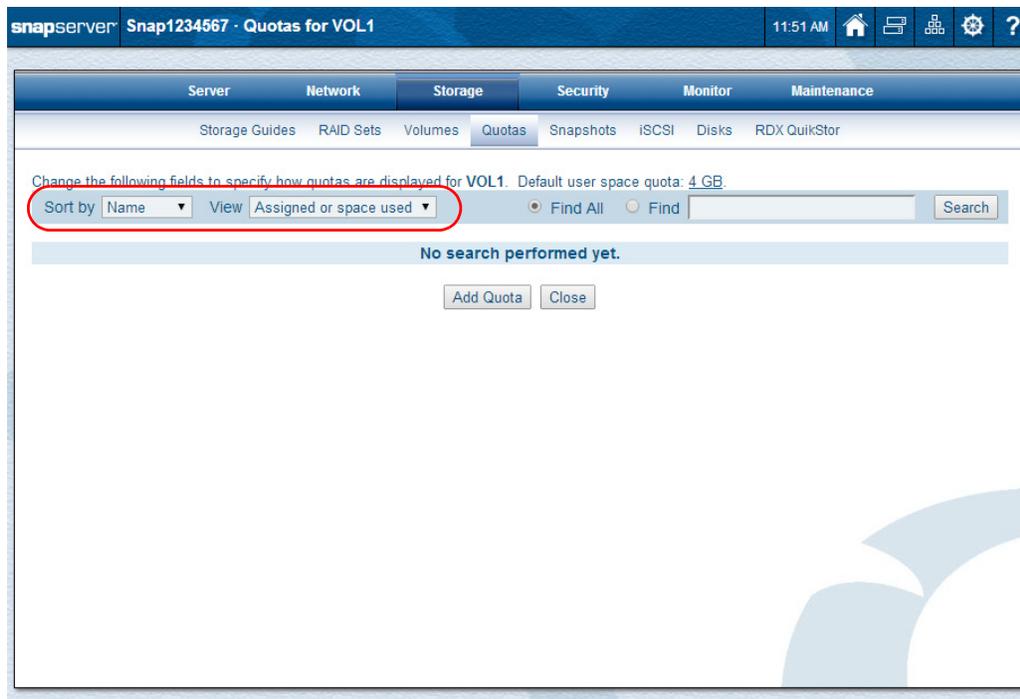
NOTE: LDAP and NIS groups do not display the third option for a default user quota.

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.

This link is only active for volumes that have quotas enabled.



2. Select the **Sort by** and **View** parameters to use.
 - From the **Sort by** 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, LDAP, 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, LDAP groups, or NIS groups that match the parameters.

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

Users are displayed with these special considerations:

- Asterisk (*) after the name indicates a Windows domain user that has been ID-mapped to a local, NIS, or LDAP user.
- Space consumption by unknown UIDs is represented as “**UID 12345**” with the numbers being the UID in question. These quotas cannot be edited.
- Space consumption by unknown UIDs that were formerly known are represented as “**UID 12345 (user_name)**.” These quotas also cannot be edited.

5. Parentheses 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 parentheses 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 LDAP and NIS groups. They don't use a volume default quota, so **no limit** (without parentheses) is shown.

6. From the search results, select the **name** of the appropriate user, LDAP group, or NIS group from the left column to open the quotas properties page.
7. Select or enter the **quota** desired and click **OK**.

LDAP and NIS groups do not display the third option for a default user quota.

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

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

The screenshot shows the SnapServer web interface for managing quotas on volume VOLUME0. The page title is "SnapServer Snap1234567 · Quotas for VOLUME0". The navigation menu includes Server, Network, Storage, Security, Monitor, and Maintenance. Under the Storage tab, there are sub-tabs for Storage Guides, RAID Sets, Volumes, Quotas, Snapshots, iSCSI, Disks, and RDX QuickStor. The main content area has a heading "Change the following fields to specify how quotas are displayed for VOLUME0. The default user quota is (50 MB)." Below this is a search filter section with "Sort by" set to "Name" and "View" set to "Assigned or space used". There are radio buttons for "Find All" and "Find", and a search input field. The results section shows "Quotas: 1 found." and a table with the following data:

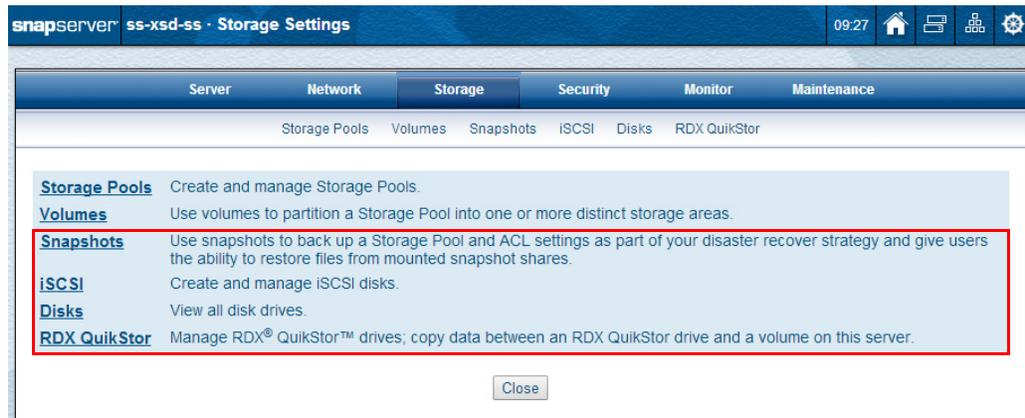
User or Group (click to edit)	Domain	Limit	Used	Used (%)
hobbiesue	Local Users	100.00 MB	0.00 MB	0% <input type="text"/>

Below the table are three buttons: "Add Quota", "Refresh", and "Close".

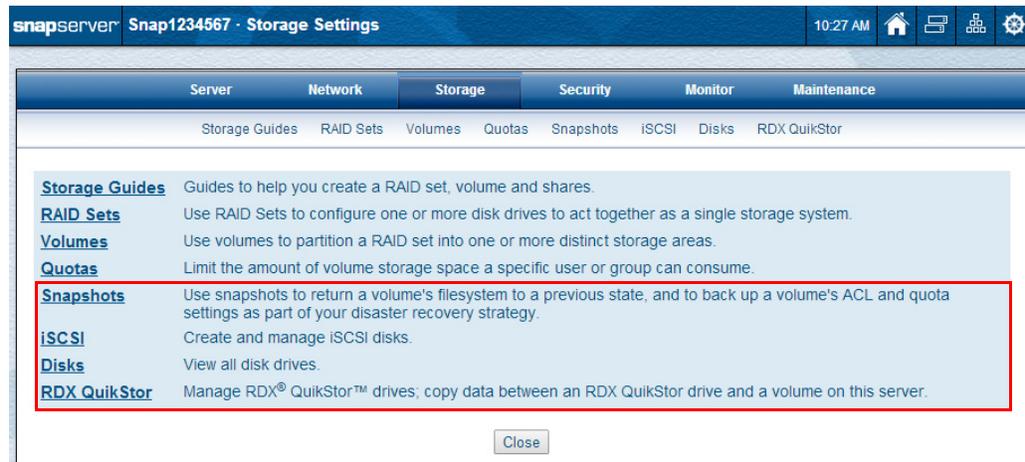
Other Storage Options

Once the RAID sets (either DynamicRAID or Traditional RAID) have been configured, you can configure the remaining four storage options for your SnapServer.

DynamicRAID Configuration



Traditional RAID Configuration



For information on the DynamicRAID configuration options, see [DynamicRAID Storage](#) in [Chapter 5](#). For information on the Traditional RAID configuration options, see [Traditional RAID Storage](#) in [Chapter 6](#).

Topics in Other Storage Options:

- [Snapshots](#)
- [iSCSI Disks](#)
- [Disks](#)
- [RDX QuikStor](#)

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. 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 154.

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

DynamicRAID Configuration

1 snapshot. Click a snapshot name to modify or delete a snapshot.

Snapshot	Status	Storage Pool	Creation Time	Expiration Time
SNAP0	OK	Head Unit Storage	2014-07-21 11:14 AM	2014-07-23 11:14 AM

Create Snapshot Snapshot Schedules Refresh Close

Traditional RAID Configuration

1 snapshot. Click a snapshot name to modify or delete a snapshot, or to rollback the volume.

Snapshot	Status	Volume	Creation Time	Expiration Time
SNAP0	OK	VOL0	2014-07-21 11:46 AM	2014-07-23 11:46 AM

Create Snapshot Snapshot Schedules Snapshot Space Refresh Close

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 total. Under normal circumstances, between nine and ten snapshots are sufficient to safely back up any system.
Edit a Snapshot Schedule	Click Snapshot Schedules and then click the snapshot name. You can then modify all snapshot parameters.
Adjust Snapshot Space	NOTE: Traditional RAID only. Click Snapshot Space and 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

NOTE: It is recommended that snapshots be taken when the system is idle or under low data traffic to minimize conflicts.

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

NOTE: The presence of one or more snapshots on a volume (Traditional RAID) or storage pool (DynamicRAID) usually has minimal performance impact, but may impact write performance when frequently overwriting file data. 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 140](#) 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. See [Schedule Snapshots on page 140](#) for more information.

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 154](#).

Create a Snapshot

1. Navigate to **Storage > Snapshots** and click **Create Snapshot**.

2. Configure the desired **settings**.
 - Enter or accept the **Snapshot Name** (20 character maximum).
 - Select the **Source Volume/Storage Pool** from the drop-down list.

- Specify **when** to create the snapshot.

Click either **Create Snapshot Now** to run the snapshot immediately or **Create Snapshot Later** to schedule the Snapshot for a later time. When you select **Create Snapshot Later**, a new input section appears below the option. Complete the following:

- Schedule a **Start Date** for the snapshot.
- Choose a **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.

- Specify the **Duration** of the snapshot.

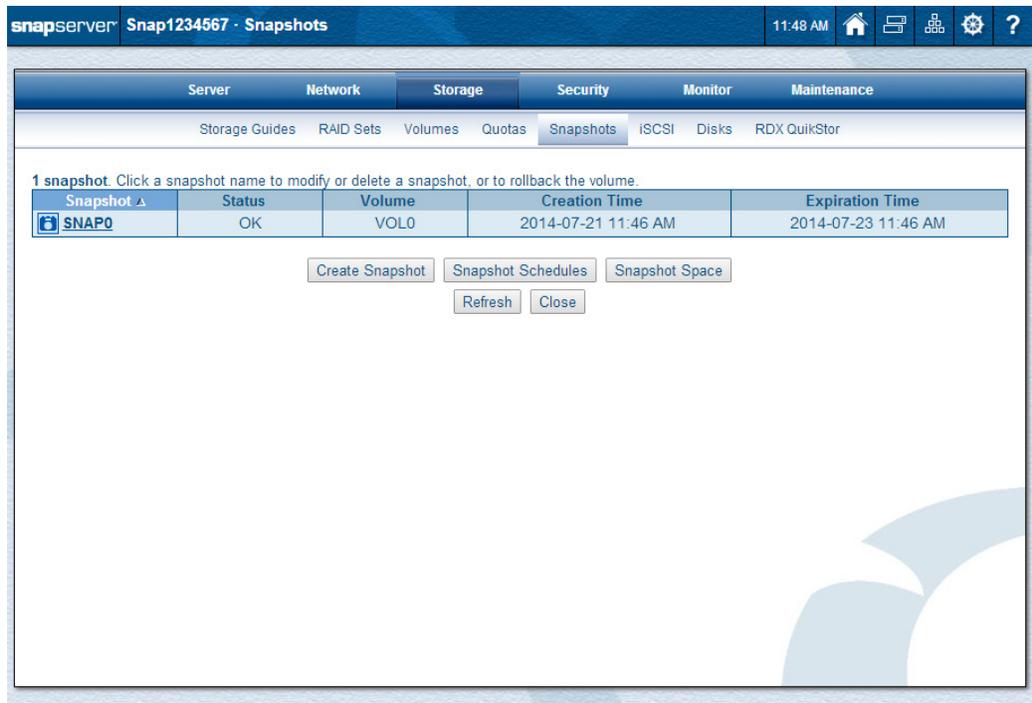
In the **Duration** field, specify how long the snapshot is to be active in hours, days, weeks, or months. The SnapServer automatically deletes the snapshot after this period expires, as long as no older unexpired snapshots exist that depend on it. If any such snapshot exists, its termination date is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

- Specify whether to create a recovery file.

If you plan to create a backup from the snapshot and want to save filesystem security configuration and quota consumption and in the backup, check the **Create Recovery File** box. See [Schedule Snapshots on page 140](#) for more information on coordinating snapshots and backup operations.

3. Click **Create Snapshot**.

If you elected to run the snapshot immediately, it appears in the current snapshots table on the **Snapshot** page.

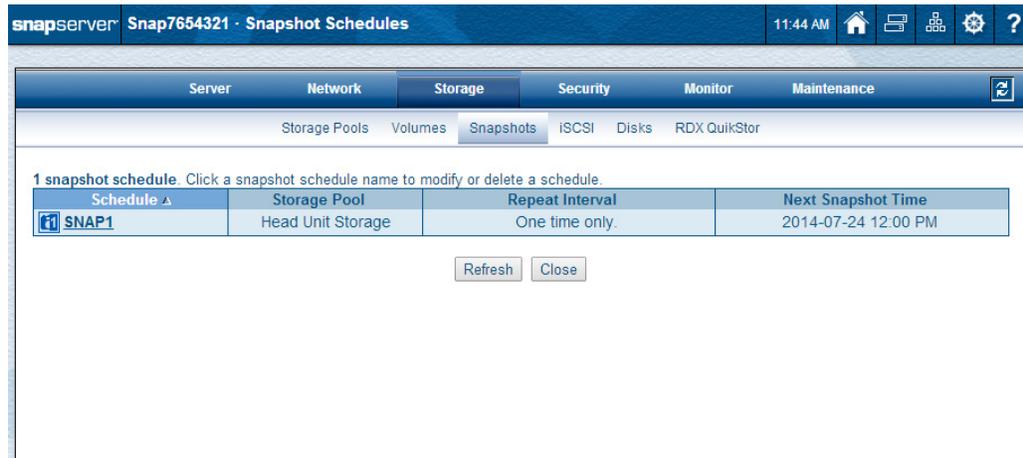


The screenshot shows the SnapServer web interface. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under 'Storage', there are sub-links for 'Storage Guides', 'RAID Sets', 'Volumes', 'Quotas', 'Snapshots', 'iSCSI', 'Disks', and 'RDX QuikStor'. The 'Snapshots' section is active, showing a table with one snapshot entry:

Snapshot	Status	Volume	Creation Time	Expiration Time
SNAP0	OK	VOLO	2014-07-21 11:46 AM	2014-07-23 11:46 AM

Below the table are several buttons: 'Create Snapshot', 'Snapshot Schedules', 'Snapshot Space', 'Refresh', and 'Close'. The page also includes a header with 'snapserver Snap1234567 · Snapshots' and a top right corner with the time '11:48 AM' and various utility icons.

If you scheduled the snapshot to run at a later time, it appears in the scheduled snapshots table under **Snapshot Schedules**.



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 on page 179](#) and [Accessing Snapshots Within the Snapshot Share on page 288](#) 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** box 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
2014-06-25.123000
2014-06-01.000100
2014-05-17.020200
```

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, `2014-06-25.123000`, or June 25, 2014, at 12:30 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.

Depending on their ability to cross bind mounts, locally-installed backup agents can access the snapshot share in one of two ways:

- via `/shares` (for example, `/shares/SHARE1_SNAP/latest`)
- via `/links` (for example, `/links/SHARE1_SNAP/latest`)

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.

The amount of space used on a Traditional RAID can be seen by navigating to **Storage > Snapshots > Snapshot Space**.

1 RAID set. Click a RAID name to modify the size of the RAID's snapshot space.

RAID Set	RAID Space Allocation	Snapshot Space Usage
5 md0	VOL0: 33.00 GB <Snapshot Space>: 6.00 GB <Unallocated Space>: 74.23 GB <Total Space>: 113.23 GB	0% (4 MB / 6 GB)

Refresh Close

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 **Snapshot Space**, 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 on page 141](#).
- **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** main 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.

The screenshot displays the 'Snapshot Properties' page in the SnapServer interface. The page title is 'SnapServer Snap1234567 - Snapshot Properties'. The top navigation bar includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under 'Storage', there are sub-tabs for 'Storage Guides', 'RAID Sets', 'Volumes', 'Quotas', 'Snapshots', 'iSCSI', 'Disks', and 'RDX QuikStor'. The 'Snapshots' tab is active, showing a table with the following data:

Snapshot	Status	Volume	Creation Time	Expiration Time
SNAP0	OK	VOL0	2014-07-21 11:46 AM	2014-07-23 11:46 AM

Below the table, the 'Snapshot Name' is 'SNAP0' and the 'Duration' is '2 days'. At the bottom of the page, there are four buttons: 'OK', 'Delete Snapshot', 'Rollback', and 'Cancel'.

Edit a Snapshot

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

Delete a Snapshot

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

Rollback to a Previous State

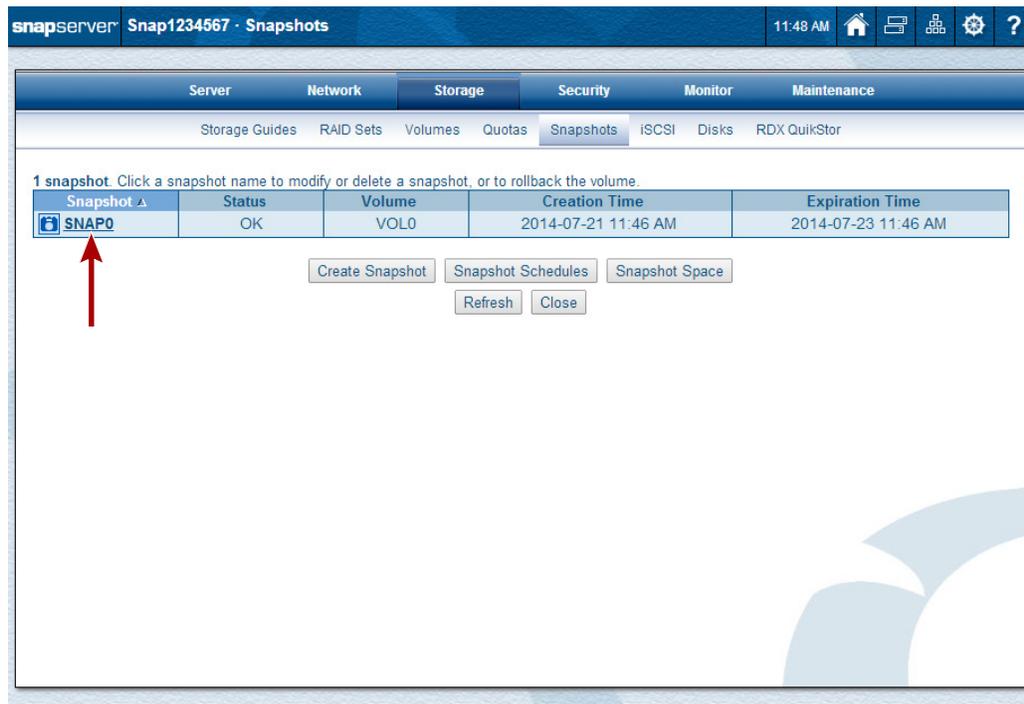
NOTE: This is only available on a Traditional RAID configuration.

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 **Rollback**. During the rollback operation, data on the volume will be inaccessible and changes blocked.



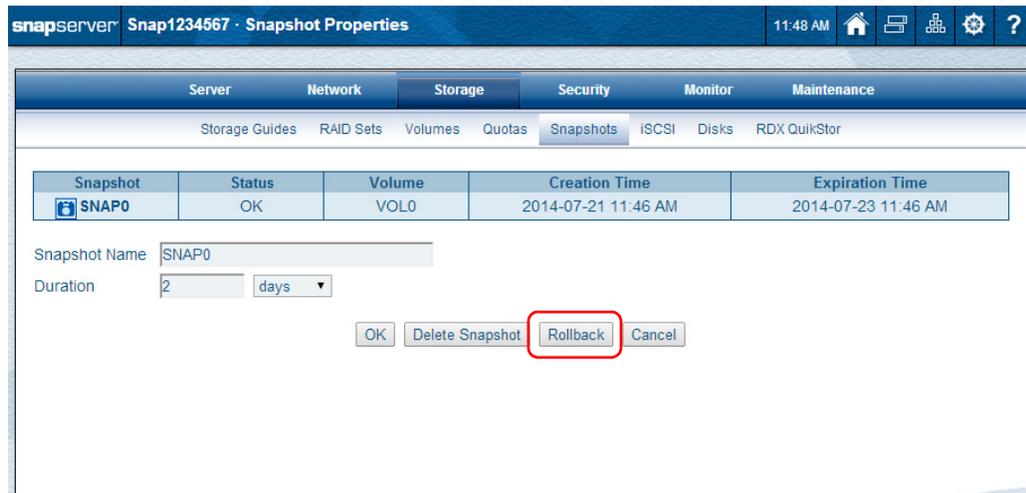
CAUTION: Rolling back a volume cannot be undone and should only be used as a last resort after attempts to restore selected directories or files have failed. Performing a rollback on a volume may disable the antivirus software. If you are using the antivirus software, take the necessary precautions.

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



2. In the left-most column, click the **name** of the snapshot you want to use.

- At the displayed **Snapshot Properties** page, click **Rollback**.

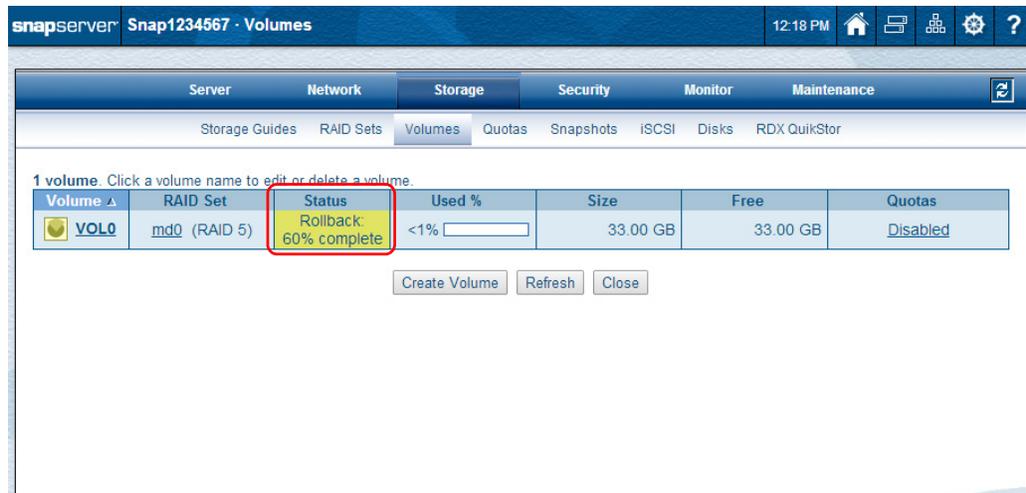


NOTE: If BitTorrent Sync is enabled, it will be disabled and reset to default settings.

- At the confirmation page, click **Rollback** 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 size of any iSCSI disk is limited to the size of a single chassis filesystem.
- 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.

- **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 the SnapServer. Refer to the vendor's documentation to properly install and configure you initiator to connect to the SnapServer iSCSI disks.

iSCSI Configuration on the SnapServer

Any iSCSI disks are created on the **Storage > iSCSI** page of the Web Management Interface. The second column in the table shows the Storage Pools for DynamicRAID systems and the Volumes for Traditional RAID systems.

2 iSCSI disks. Click an iSCSI disk name to edit or delete the iSCSI disk. (Note: Mouseover a disk name to view its IQN.)

iSCSI Disk	Storage Pool	Status ▲	Active Clients	Authentication	Size
iscsi0	Head Unit Storage	OK	0	None	21.00 GB
iscsi1	Head Unit Storage	OK	0	None	17.00 GB

Before setting up iSCSI disks on your SnapServer, carefully review the information in the sections below.

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 native filesystem.
- iSCSI disks are managed from client software rather than the SnapServer 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 154](#) for more information).

iSCSI Multi-Initiator Support

Check the **Support Multiple Initiators** box to allow two or more initiators to simultaneously access a single iSCSI target. Multiple 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.

When the box for **Support Multiple Initiators** is checked, a warning message appears:

```
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.
```

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 on page 122](#).

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 147](#) before you begin creating iSCSI disks.

The screenshot shows the 'Create iSCSI Disk' configuration page in the SnapServer web interface. The page is titled 'SnapServer Snap7654321 - Create iSCSI Disk' and includes a navigation menu with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. The 'Storage' tab is active, and the 'iSCSI' sub-tab is selected. The configuration form includes the following fields and options:

- iSCSI Disk Name:** iscsi0
- Storage Pool:** Head Unit Storage (85.58 GB available)
- Size:** 77 GB
- 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
- User Name:** [text input]
- Target Secret:** [text input] (Minimum 12 characters)
- Confirm Target Secret:** [text input]

At the bottom of the form, there are two buttons: 'Create iSCSI Disk' and 'Cancel'.

NOTE: You cannot delete or edit an iSCSI disk until all its clients have been disconnected from that specific 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.
Enable Write Cache	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 on page 148 . 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.
Support Multiple Initiators	Check this box if you want your iSCSI disk to allow multiple initiator connections. NOTE: Data corruption is possible if this option is checked. See iSCSI Multi-Initiator Support on page 148 for more information.

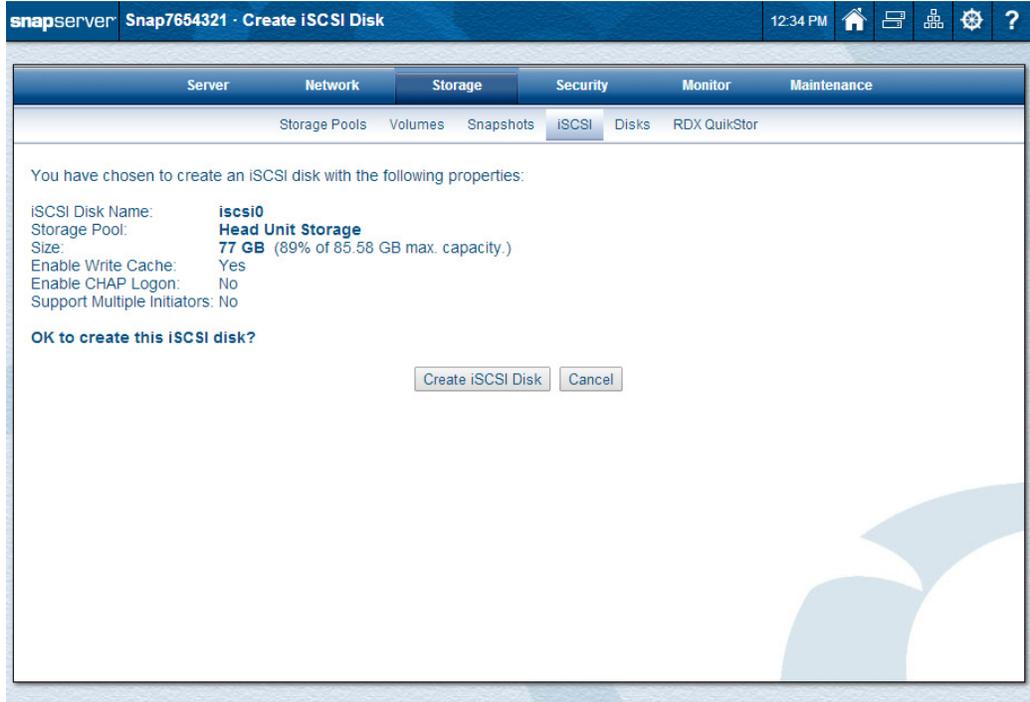
Step 2: If desired, enable CHAP authentication.

In the bottom half of the page, 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 **Create iSCSI Disk**. At the confirmation page, verify the settings and click **Create iSCSI Disk** again.



You are returned to the primary iSCSI page and the new iSCSI disk is displayed in the table there with the following information:



The iSCSI information is shown in a table with these columns:

Label	Description
iSCSI Disk	The name of the iSCSI disk.
Storage Pool/Volume	The pool or volume on which the iSCSI disk was created.
Status	Current condition of the iSCSI disk: <ul style="list-style-type: none"> • OK – The iSCSI disk is online and accessible. • Not Mounted – The iSCSI disk is offline.
Active Clients	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 **iSCSI** main 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 7.6 provides VSS and VDS hardware providers for support of 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 on page 135](#). 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.

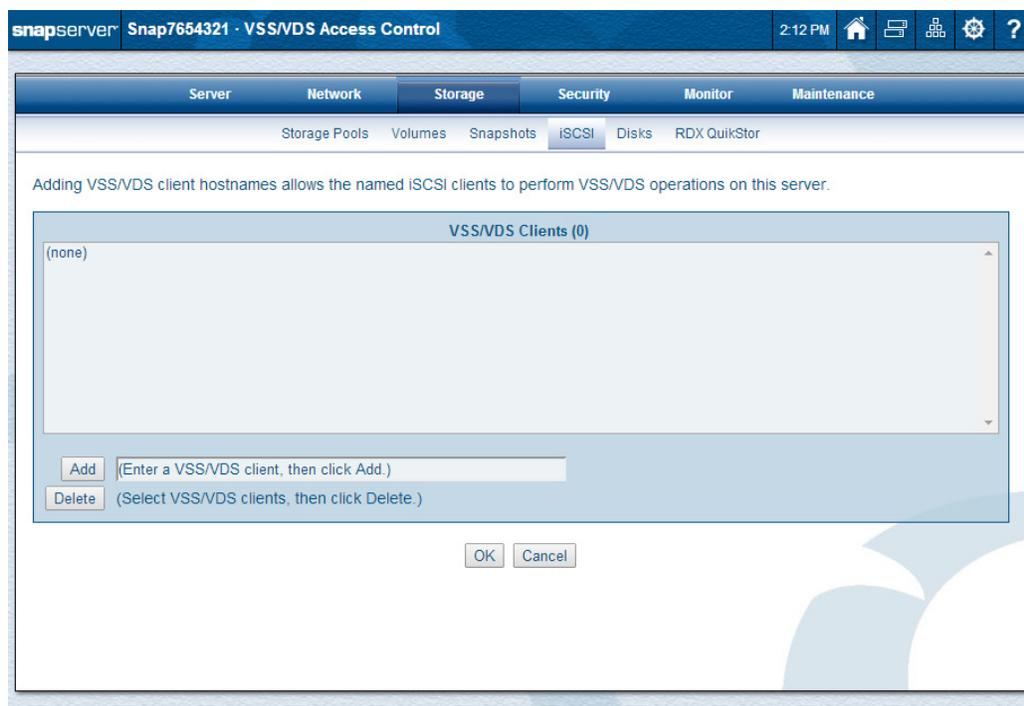
NOTE: To use Symantec Backup Exec as your VSS-compatible backup application, you must first modify the registry of the Backup Exec server and agents.

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 SnapServer 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 **VSS/VDS Access**.
 - b. Enter a VSS/VDS client and then 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 **VSS/VDS Access**.
 - 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:

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

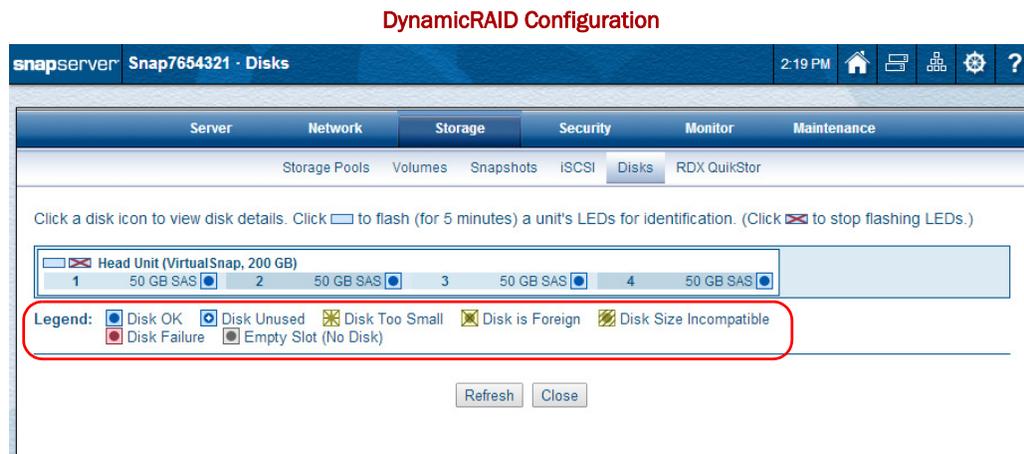
RAID set types listed in *Storage Manager for SANs* when creating an iSCSI disk reflect the types of RAID sets already configured on the SnapServer. Once a RAID set type is selected, the SnapServer automatically chooses a SnapServer RAID set of the selected type and volume to create the iSCSI disk on.

Deleting VSS/VDS Client Access

1. From the **Storage > iSCSI** page, click **VSS/VDS Access**.
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 disk icon.



Traditional RAID Configuration

The screenshot shows the 'Disks' configuration page in the SnapServer GuardianOS 7.6 interface. The page title is 'SnapServer Snap1234567 - Disks'. The navigation tabs include Server, Network, Storage, Security, Monitor, and Maintenance. Under the Storage tab, there are sub-tabs for Storage Guides, RAID Sets, Volumes, Quotas, Snapshots, iSCSI, Disks, and RDX QuikStor. The main content area displays a RAID set named 'Head Unit (VirtualSnap, 200 GB)' with four disks: 1 md0 (50 GB SAS), 2 md0 (50 GB SAS), 3 md0 (50 GB SAS), and 4 md0 (50 GB SAS). A legend below the RAID set defines the status icons: Disk OK (blue square), Disk is Foreign (yellow square), Disk Failure (red square), Empty Slot (No Disk) (grey square), Local Spare (blue square with 'L'), Global Spare (GS) (blue square with 'GS'), RAID level (yellow circle with RAID level number), Global Spare (yellow square with 'GS'), and Unassigned (grey square with '--'). There are 'Refresh' and 'Close' buttons at the bottom of the RAID set display.

- Click a disk icon (such as ) 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).

If expansion arrays are attached to your server, they will also be displayed on this page.

Click a disk icon to view disk details. Click to flash (for 5 minutes) a unit's LEDs for identification. (Click to stop flashing LEDs.)

Head Unit (DX2, 17.58 TB)			
1	0.98 TB SATA	2	0.98 TB SATA
5	0.98 TB SATA	6	0.98 TB SATA
9	0.98 TB SATA	10	0.98 TB SATA
3	1.95 TB SATA	7	1.95 TB SATA
11	1.95 TB SATA	4	1.95 TB SATA
12	1.95 TB SATA	8	1.95 TB SSD

Expansion Unit 1 (SnapExpansion DX, 10.26 TB)			
1	1.47 TB SAS	2	1.47 TB SAS
5	1.47 TB SSD	6	1.47 TB SAS
9	(No Disk)	10	(No Disk)
3	1.47 TB SAS	7	1.47 TB SAS
11	(No Disk)	4	1.47 TB SAS
12	(No Disk)	8	(No Disk)

Expansion Unit 2 (SnapExpansion DX, 9.77 TB)			
1	500.6 GB SAS	2	500.6 GB SAS
5	1.47 TB SAS	6	1.47 TB SAS
9	500.6 GB SAS	10	500.6 GB SAS
3	500.6 GB SAS	7	1.47 TB SAS
11	500.6 GB SAS	4	500.6 GB SAS
12	500.6 GB SSD	8	1.47 TB SSD

Expansion Unit 3 (SnapExpansion DX, 1.95 TB)			
1	1.95 TB SATA	2	(No Disk)
5	(No Disk)	6	(No Disk)
9	(No Disk)	10	(No Disk)
3	(No Disk)	7	(No Disk)
11	(No Disk)	4	(No Disk)
12	(No Disk)	8	(No Disk)

Expansion Unit 4 (SnapExpansion DX, 3.52 TB)			
1	1.76 TB SAS	2	1.76 TB SAS
5	(No Disk)	6	(No Disk)
9	(No Disk)	10	(No Disk)
3	(No Disk)	7	(No Disk)
11	(No Disk)	4	(No Disk)
12	(No Disk)	8	(No Disk)

Expansion Unit 5 (SnapExpansion DX, 5.86 TB)			
1	1.95 TB SAS	2	1.95 TB SAS
5	(No Disk)	6	(No Disk)
9	(No Disk)	10	(No Disk)
3	1.95 TB SAS	7	(No Disk)
11	(No Disk)	4	(No Disk)
12	(No Disk)	8	(No Disk)

Legend: Disk OK Disk Unused Disk Too Small Disk is Foreign Disk Size Incompatible
 Disk Failure Empty Slot (No Disk)

[Click here to stop flashing LEDs on all units.](#)

Refresh Close

NOTE: If GuardianOS detects an expansion unit that is not integrated with the SnapServer, a message displays across the top of the administration pages 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, 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 in DynamicRAID mode, 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 in Traditional RAID mode, 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.

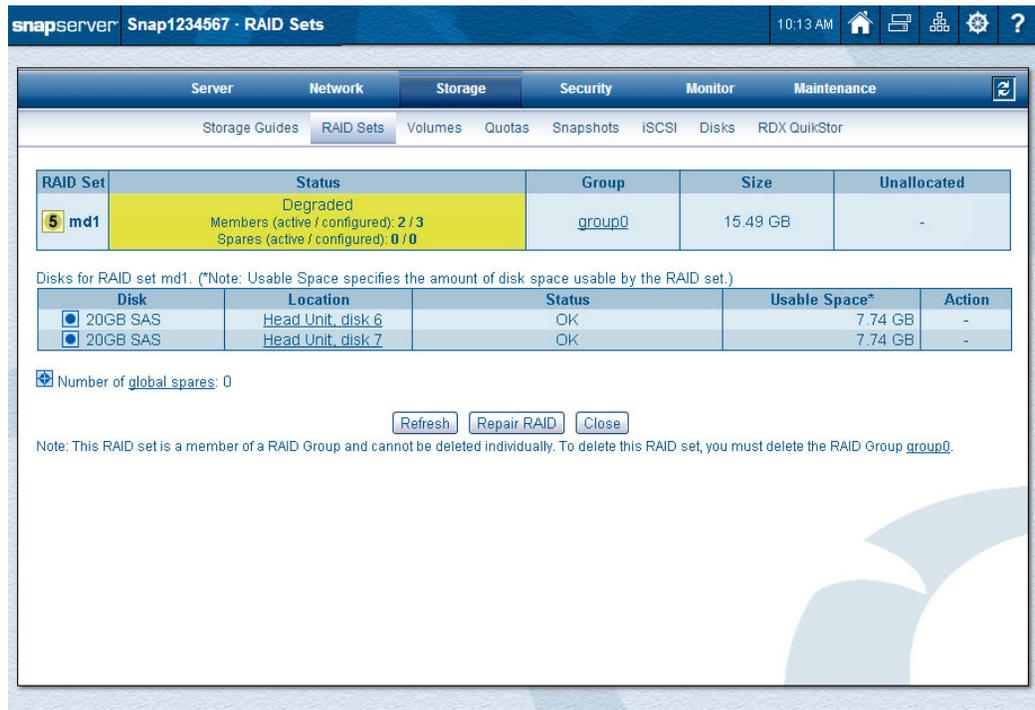
If a disk drive fails, the Traditional RAID Administration page changes to show the Disk storage as Degraded and provides a link to the **RAID Sets** page. Both the **RAID Sets** and **RAID Set Properties** pages show the degraded status.

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 **name** of the RAID set with the new drive.
2. Click **Repair RAID**.

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.



NOTE: The **Repair RAID** button only appears when a drive has failed or been removed, and the RAID is in degraded mode.

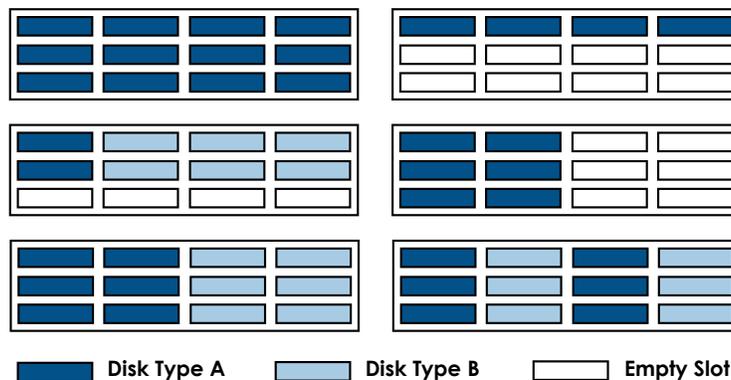
The RAID set status changes to **Resyncing** while the new drive is incorporated into the RAID set. It reads **OK** once the incorporation is complete.

Adding Disk Drives

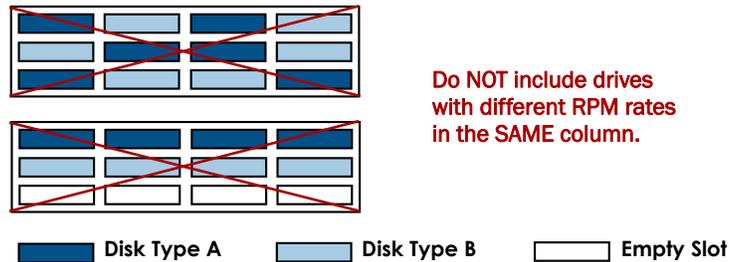
If empty slots are available, you can add an Overland-approved disk drive to expand the storage pool/volume on your SnapServer or SnapExpansion unit.

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. 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.

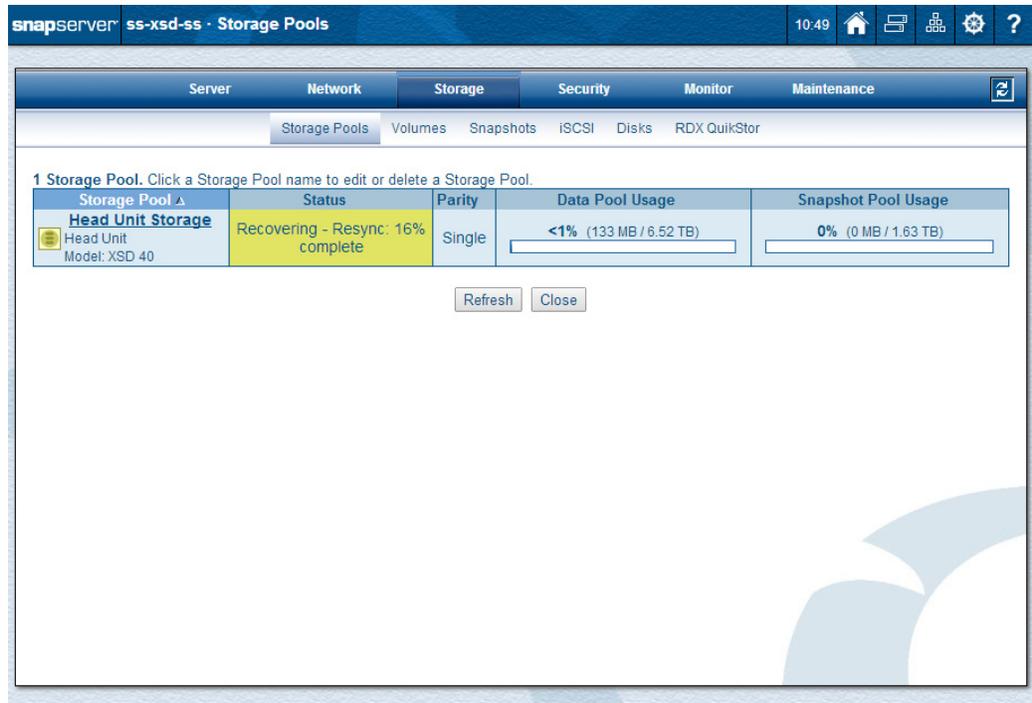
2. Click the **link** in the banner (or navigate to the **Storage > Storage Pools** page).
3. In the **Status** column, click the link to open the **Storage Pool Properties** page.

- At the lower portion of the properties page, verify or change the **Parity Mode** and **Snapshot Pool** settings.

The screenshot displays the 'Storage Pool Properties' interface for 'Head Unit Storage'. The 'Parity Mode' section is highlighted with a red box, showing 'Dual-parity protection' selected. The 'Snapshot Pool' section shows '20%' reserved for snapshots. At the bottom, there are buttons for 'OK', 'Refresh', 'View Disks', 'Delete Storage Pool', and 'Cancel', along with an important note to click OK to confirm 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.



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 on page 113](#)), 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 new drive as 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 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.

To remove the drive as a Global Spare, use the same process but uncheck the box in the **Available Disks** list.

Reintegrate Orphaned Disk Drives

An orphaned disk drive can occur in either of the following circumstances:

- A working drive from a RAID set is accidentally removed from the server
- 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 on the **Storage > RAID Sets** page, and then click the **Repair** link next to the drive in question.

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, see [Disks on page 157](#).

Click a disk icon to view disk details. Click to flash (for 5 minutes) a unit's LEDs for identification. (Click to stop flashing LEDs.)

Head Unit (DX2, 17.58 TB)																							
1	0.98 TB SATA	2	0.98 TB SATA	3	1.95 TB SATA	4	1.95 TB SATA	5	0.98 TB SATA	6	0.98 TB SATA	7	1.95 TB SATA	8	1.95 TB SSD	9	0.98 TB SATA	10	0.98 TB SATA	11	1.95 TB SATA	12	1.95 TB SATA

Expansion Unit 1 (SnapExpansion DX, 10.26 TB)																							
1	1.47 TB SAS	2	1.47 TB SAS	3	1.47 TB SAS	4	1.47 TB SAS	5	1.47 TB SSD	6	1.47 TB SAS	7	1.47 TB SAS	8	(No Disk)	9	(No Disk)	10	(No Disk)	11	(No Disk)	12	(No Disk)

Expansion Unit 2 (SnapExpansion DX, 9.77 TB)																							
1	500.6 GB SAS	2	500.6 GB SAS	3	500.6 GB SAS	4	500.6 GB SAS	5	1.47 TB SAS	6	1.47 TB SAS	7	1.47 TB SAS	8	1.47 TB SSD	9	500.6 GB SAS	10	500.6 GB SAS	11	500.6 GB SAS	12	500.6 GB SSD

Expansion Unit 3 (SnapExpansion DX, 1.95 TB)																							
1	1.95 TB 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)

Expansion Unit 4 (SnapExpansion DX, 3.52 TB)																							
1	1.76 TB SAS	2	1.76 TB 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 (SnapExpansion DX, 5.86 TB)																							
1	1.95 TB SAS	2	1.95 TB SAS	3	1.95 TB 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)

Legend: Disk OK Disk Unused Disk Too Small Disk is Foreign Disk Size Incompatible Disk Failure Empty Slot (No Disk)

[Click here to stop flashing LEDs on all units.](#)

Refresh Close

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.

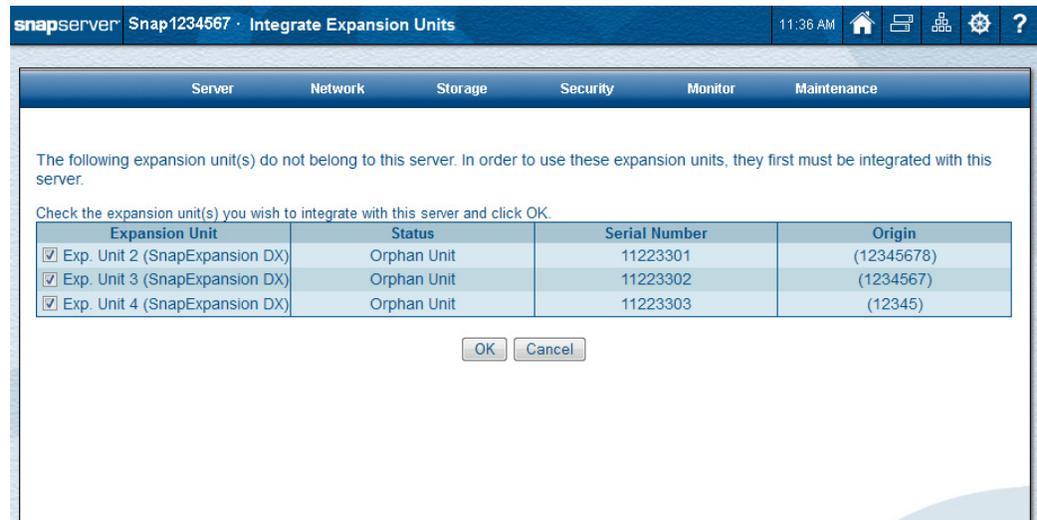
Traditional RAID

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

Newly discovered expansion units that have previously been configured by a different head unit but have not been integrated with this SnapServer are listed in the Orphan Expansion Units Table and can be accessed by the message link in the banner at the top of the administration pages.



The information shown is covered in the following table:

Property	Description
Expansion Unit	A description of the unit
Status	The status of the unit (for example, Orphan Unit)
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, check the box 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, the expansion is unconfigured or configured for the same RAID mode (Traditional RAID or DynamicRAID) as the SnapServer itself.

RDX QuikStor

RDX QuikStor, a media-based removable storage system, is now directly accessible through the Web Management Interface. It can be used to manually copy files for quick off-site data redundancy and to transfer files between SnapServers and other servers or clients without using the network.

Whenever an RDX QuikStor appliance is plugged into a USB port of the server, it automatically shows up in the Web Management Interface at **Storage > RDX QuikStor**.

NOTE: GOS is only able to mount RDX volumes that are formatted with NTFS (default) or XFS.

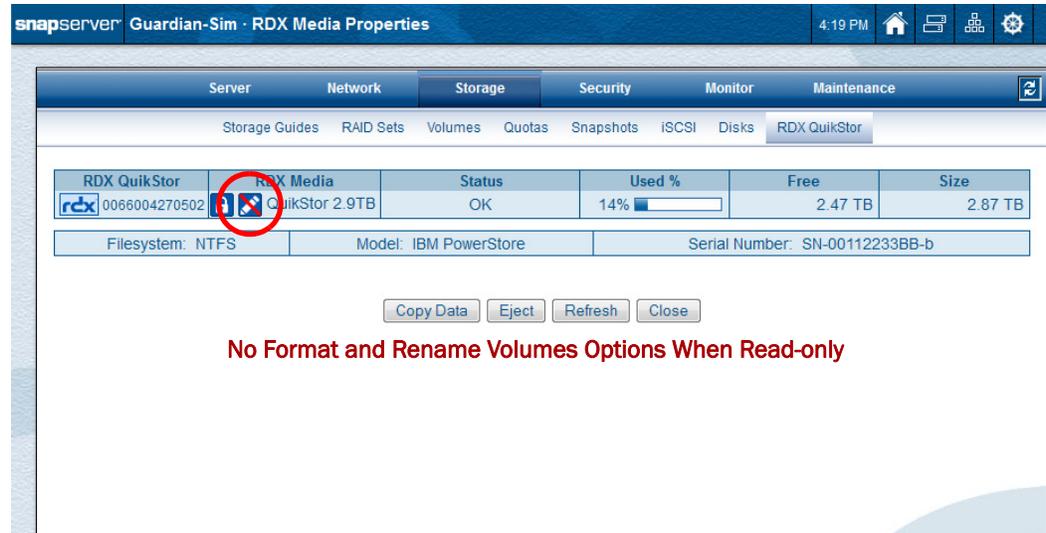
The bold text above the table provides a quick summary of the number of drives connected and media inserted and available.

RDX QuikStor	RDX Media	Status	Used %	Free	Size
0066004270502	QuikStor 2.9TB	OK	14%	2.45 TB	2.87 TB
0066004270501	QuikStor 3.2TB	OK	91%	162.06 GB	1.91 TB
0066004270503	(No media.)	(No media.)	-	-	-

Label	Description
RDX QuikStor	The identity of the RDX QuikStor. This is not editable.
RDX Media	Shows media details: <ul style="list-style-type: none"> • Eject button status icon – Shows if the eject button on the drive is either locked (🔒) or unlocked (🔓). Refer to Eject RDX Media for details. • Read/Write Access status icon – Shows if the media in the drive is either write-protected (🔒) or in read/write mode (🔓). • Media name – User-configurable RDX media volume name.
Status	Current condition of the RDX QuikStor. Some conditions are: <ul style="list-style-type: none"> • OK – The RDX QuikStor is online and accessible. • No Media – The RDX QuikStor is online but no media has been inserted. • Formatting – RDX media is currently formatting. • Ejecting – The RDX media is currently ejecting. • Media Not Formatted – RDX media is inserted but is not formatted. You must format the media before you can store data on it. • Media Failure – The RDX media has failed.
Used %	The amount of space being used. Both numeric percentage and bar graph shown.
Free	Remaining available space on the RDX media.
Size	The size of the total usable space on the RDX media.

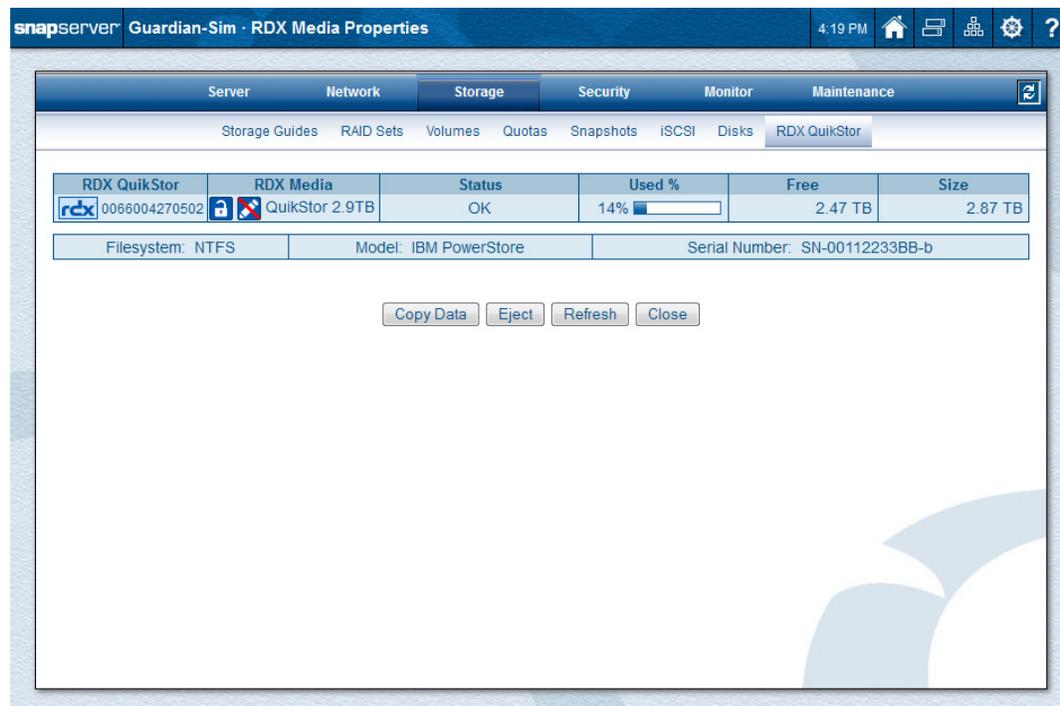
Whenever media is inserted into an RDX system, a pencil icon appears indicating whether the media is read-write (normal icon) or read-only (red line through the icon). If the RDX media is read-only, the **Format** and **Rename Volume** options are not available.

For example, if you click the RDX media name that is read-only, you would see:



RDX Media Properties

To display the **RDX Media Properties** page and access the RDX QuikStor options, click the RDX media name in the second column.

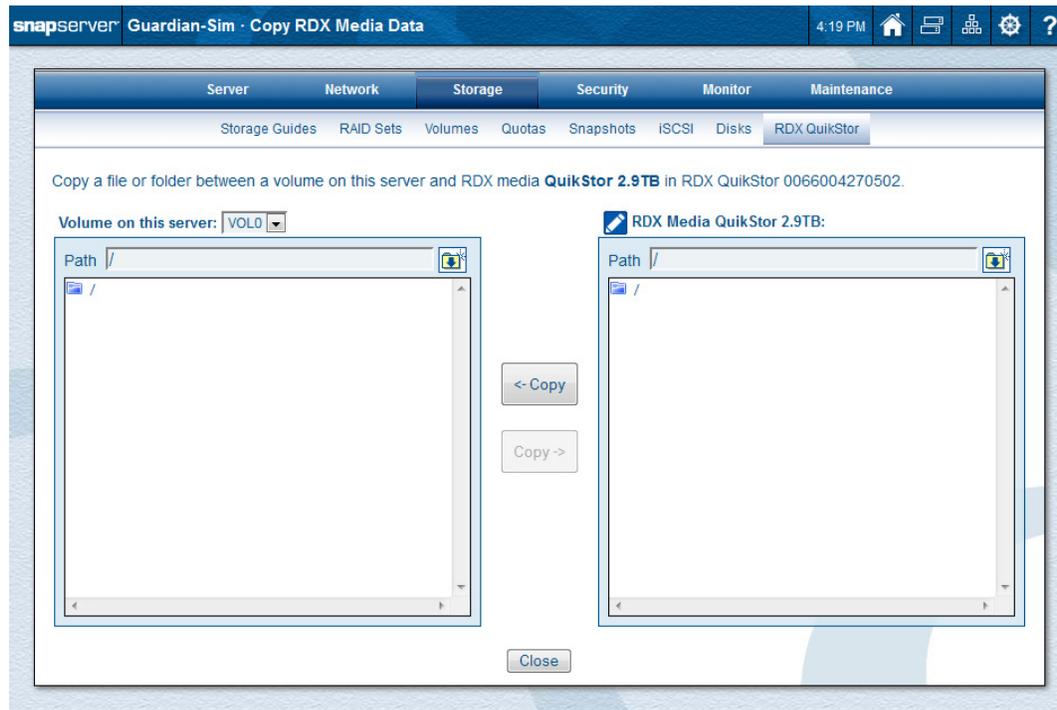


NOTE: If the RDX media is configured with multiple partitions by some other computer, GOS only sees one partition.

Although RDX Media changes are reflected automatically in the Web Management Interface, you can click **Refresh** at any time to see any changes.

Copy Data To/From RDX Media

To copy data to and from the RDX media, click **Copy Data** to go to the **Copy RDX Media Data** page.



Copying Data To and From RDX QuikStor

These steps are used to copy data to or from RDX media. If the destination folder does not yet exist, create it as you normally would on the destination volume.

1. At **Storage > RDX QuikStor**, click the RDX media name to access the **RDX Media Properties** page.
2. Click **Copy Data**.
3. At the **Copy RDX Media Data** page:
 - a. From the left drop-down **Volume** list, choose the **volume** on the SnapServer that will be used.
 - b. Select the **file or folder** involved from the volume file/folder list.

NOTE: Optionally, you can type a path in the Path field and click the **browse** button to the right of the field to verify the path. If it does not exist, you are prompted to create it.

- c. Select the **file or folder** involved from the RDX media file/folder list.

NOTE: Optionally, you can type a path in the Path field and click the **browse** button to the right of the field to verify the path. If it does not exist, you are prompted to create it.

- d. Click the appropriate **Copy** direction button to complete the process.

NOTE: If the copy operation successfully completes in only a few seconds, then the target file/folder list is refreshed and you are able to immediately either perform another copy operation or click **Close** to return to the **RDX Media Properties** page. However, if the copy operation fails (or takes longer than a few seconds to complete), you are redirected to the **RDX Media Properties** page to view the real-time status of the copy operation.

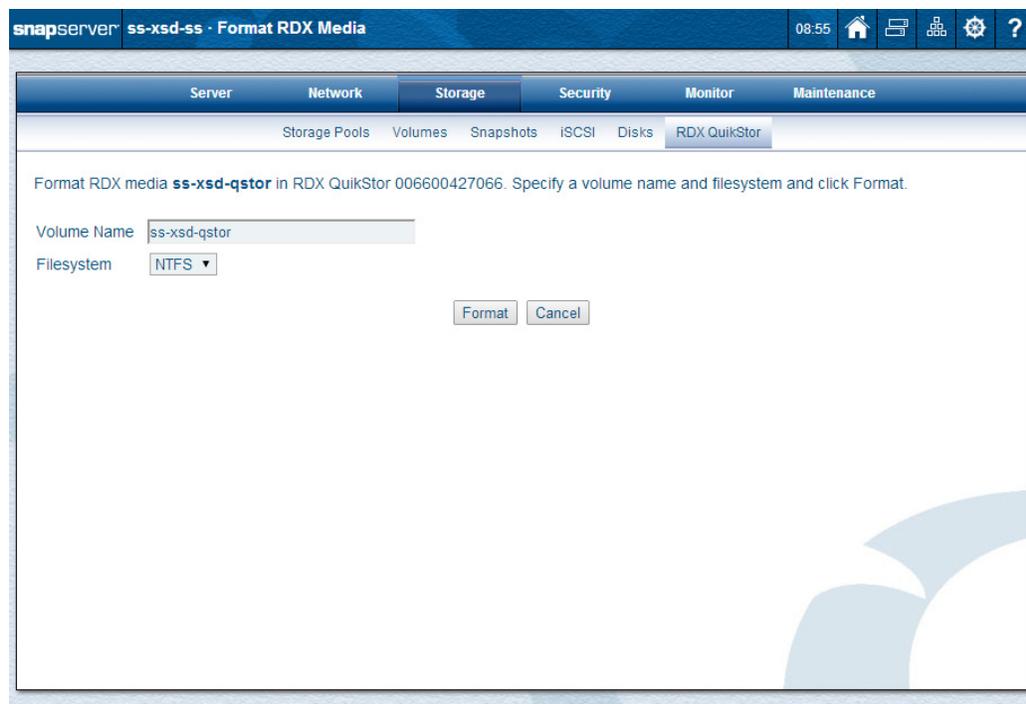
4. When done, click **Close**.

Format RDX Media



CAUTION: Reformatting RDX media erases all data currently on the media.

To format the media in the RDX QuikStor, go to the **Format RDX Media** page.



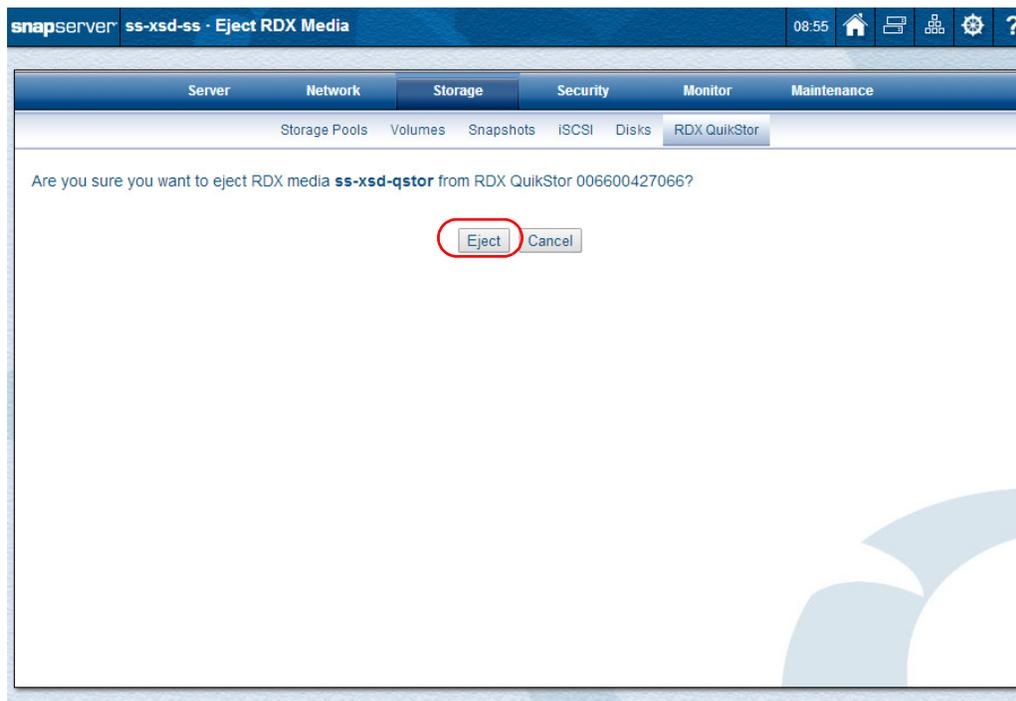
Format RDX Media

1. At **Storage > RDX QuikStor**, click the RDX media name to access the **RDX Media Properties** page.
2. Click **Format**.
3. At the **Format RDX Media** page:
 - a. Choose a **Volume Name** for the media.
 - b. Select the **Filesystem** (NTFS to share files with Windows machines or XFS to share files with Linux machines or SnapServers) to be used.
4. Click **Format** to complete the process.

Eject RDX Media

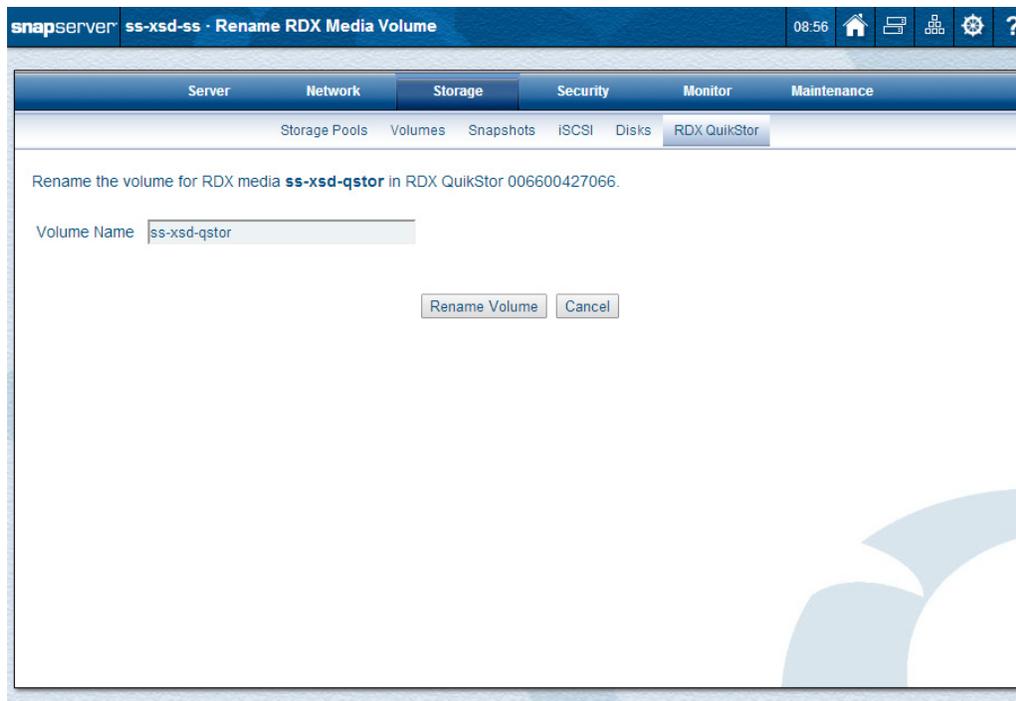
To eject the RDX media, go to **Storage > RDX QuikStor > RDX Media Properties** and click **Eject**. At the confirmation page, click **Eject** again.

NOTE: The eject lock option in RDX QuikStor screen only locks the physical eject button on the drive. This screen is not impacted by that setting.

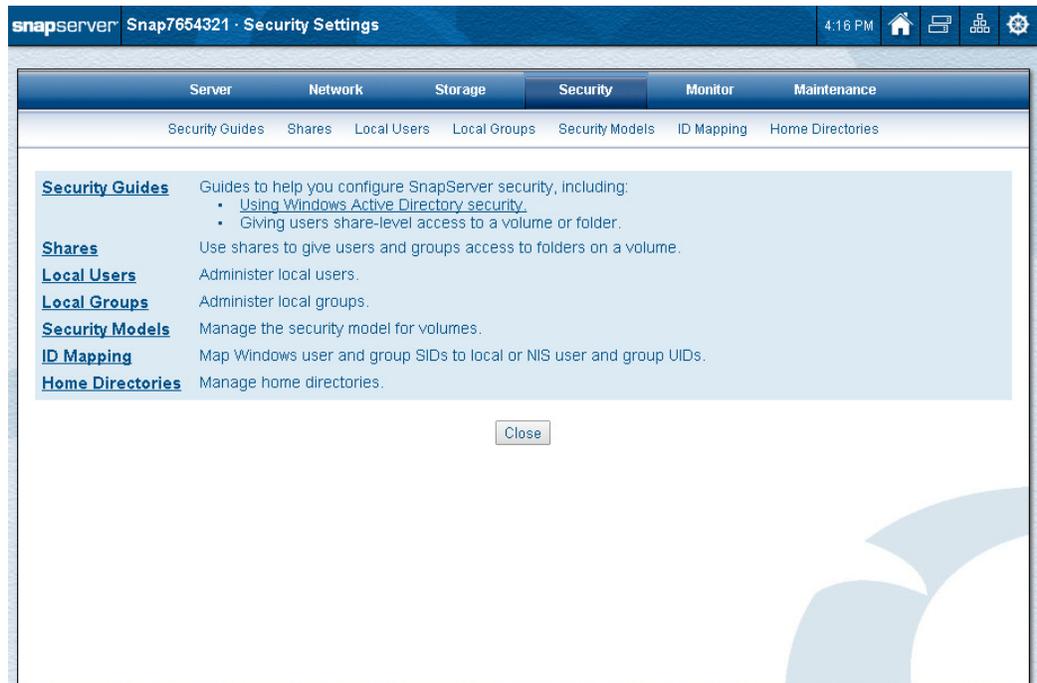


Rename RDX Media Volume

To rename the RDX media volume, go to **Storage > RDX QuikStor > RDX Media Properties** and click **Rename Volume**. At the **Rename RDX Media Volume** page, enter the new name and click **Rename Volume**.



The Security options control the access to the SnapServer and its data.



SnapServer authentication validates a user’s identity by requiring the user to provide a registered login name (User ID) and corresponding password. The server ships 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 LDAP or NIS domain, and the SnapServer can look up user IDs (UIDs) and group IDs (GIDs) maintained by the domain for configuration of quotas and ID mapping. 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 on page 175](#).

Topics in Security Options

- [Security Considerations](#)
- [Security Guides](#)
- [Shares](#)
- [Local Users](#)
- [Local Groups](#)

- [Security Models](#)
- [ID Mapping](#)
- [Home Directories](#)

Security Considerations

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 but enabled for HTTP, AFP, and FTP.

Network clients can initially access the server using the guest account (where enabled), 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 **Security > Local Users** or **Security > Local Groups** in the Web Management Interface. Local users are also used for administrative access to the server through the Web Management Interface, SSM, or the CLI through SSH.

A local user or group is one that is defined locally on a 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 and Mac clients, duplicate their local client logon credentials on the SnapServer by creating local accounts on the server that match those used to log on to client workstations. This strategy allows users to bypass the login procedure when accessing the server.



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).

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 server 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 following tools are available for creating, modifying, and editing local user and group accounts:

Function	Navigation Path
Local User Management	Navigate to the 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 Local Groups page, from which you can create, view, edit, and delete local groups.

User and Group ID Assignments

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 server, including Windows Active Directory, LDAP, NIS, and local users.

If you join the server to a Windows, LDAP, or NIS domain, IDs are assigned using available IDs only. Consider the following when creating users and groups:

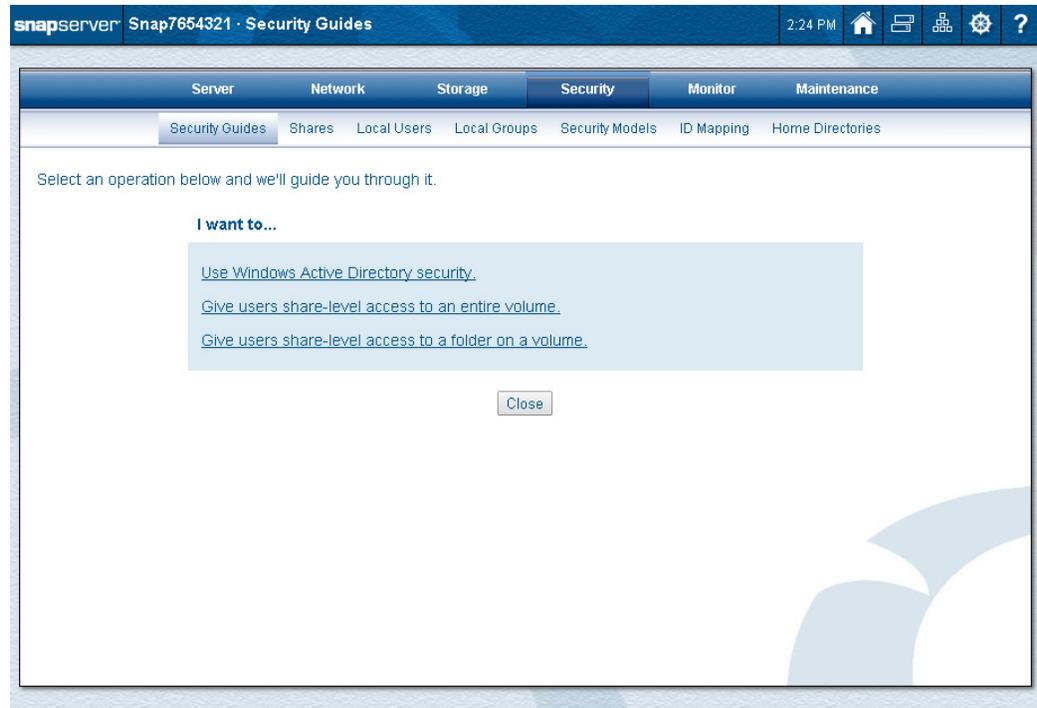
- UIDs and GIDs from 0 to 100 are unavailable for use. If you try to assign a UID or GID that is less than 101 (or in use by Windows, LDAP, or NIS domain), you will get an error message.
- When the server automatically generates UIDs or GIDs for imported Windows domain users or groups, UIDs or GIDs that are already in use by LDAP, NIS, or local users are skipped.
- When LDAP or NIS domain users and groups are imported, the server discards 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:

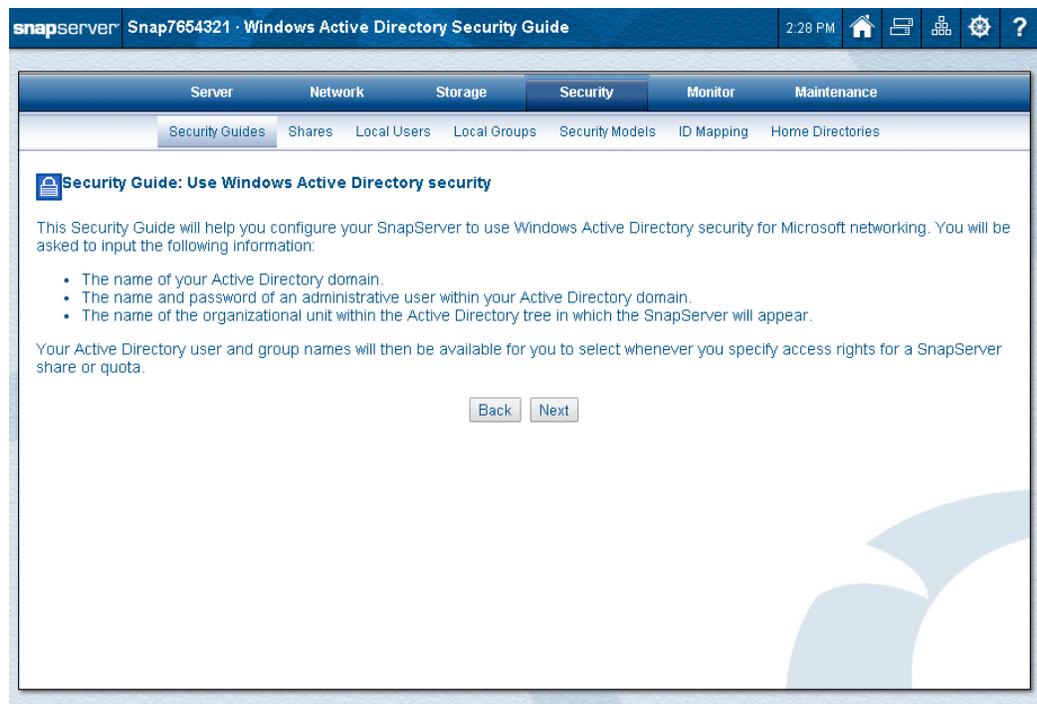
- Setting up Windows Active Directory security.
- Giving users or groups share-level access to an **entire volume**.
- Giving users or groups share-level access to a **folder on a volume**.



Security Guide for Windows Active Directory

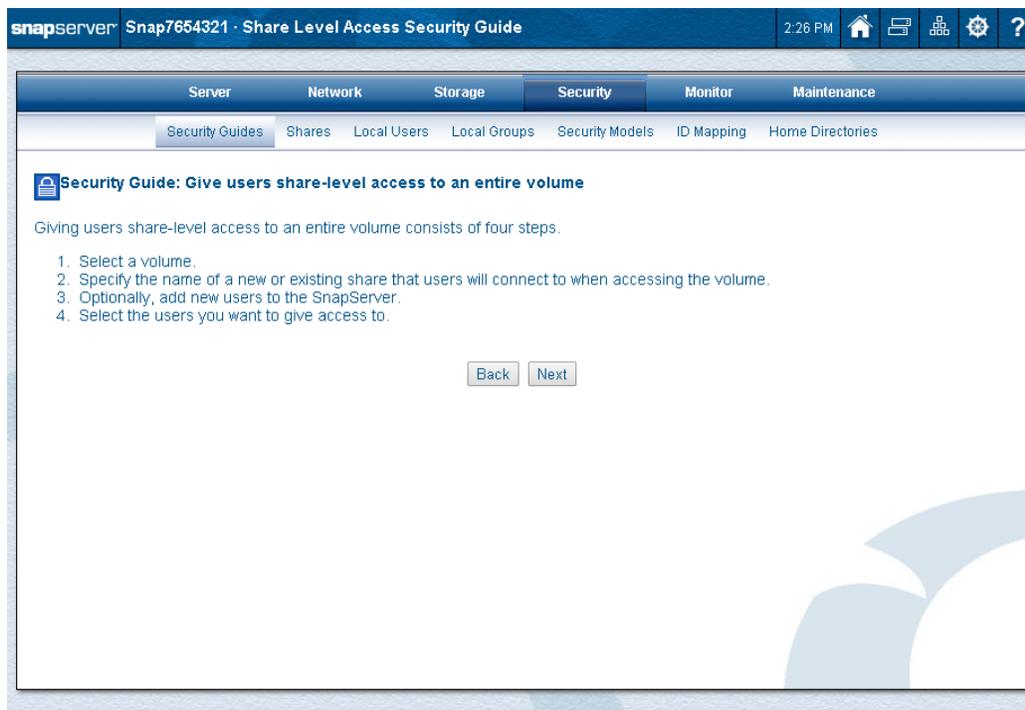
The **Windows Active Directory Security Guide** wizard guides you through the setup of Windows Active Directory on your server.

NOTE: You cannot join an Active Directory domain if NTP is enabled. If you see such a message, click the NTP link to change your settings.



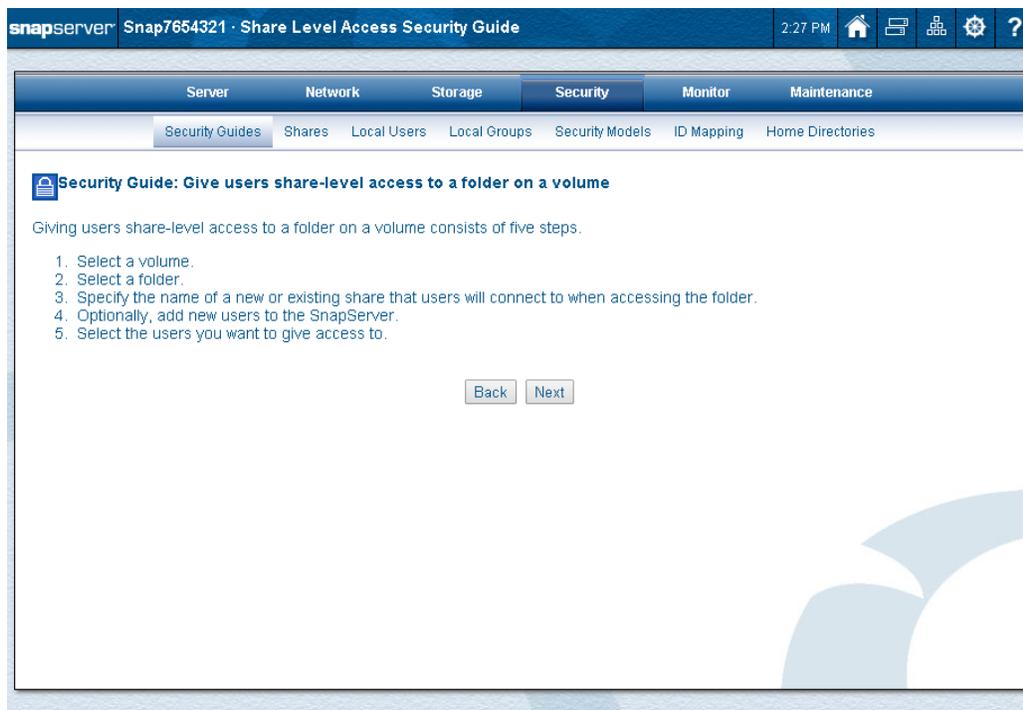
Security Guide for Entire Volume Access

This **Share Level Access Security Guide** wizard guides you through the four steps it takes to give share-level access to an **entire volume**.



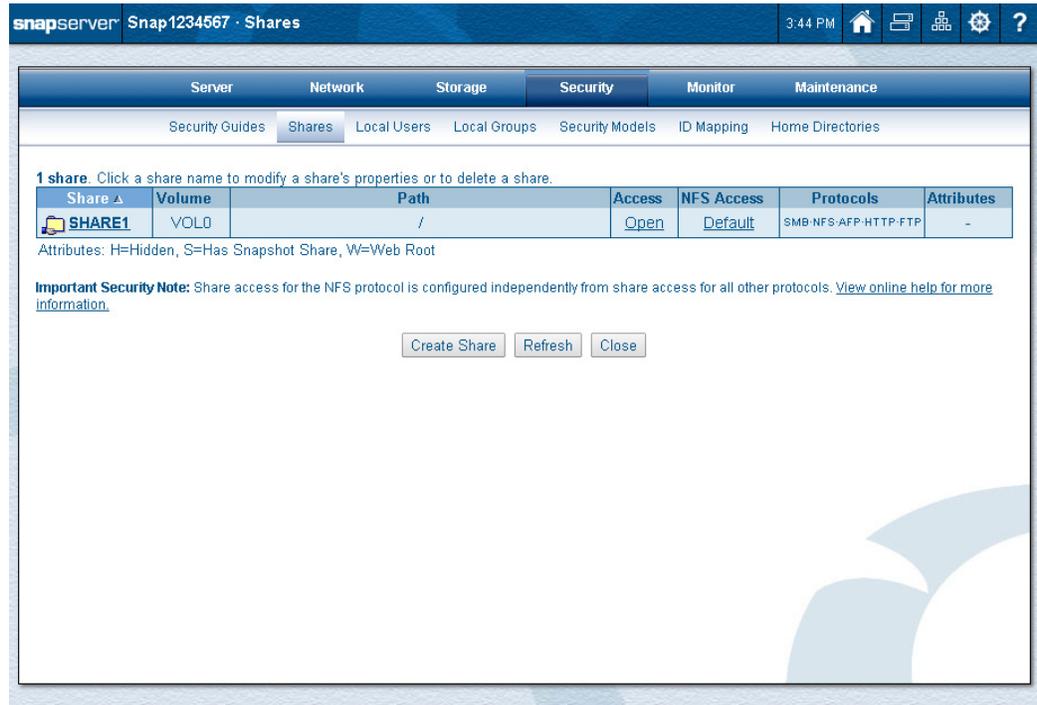
Security Guide for Folder Access on Volume

This **Share Level Access Security Guide** wizard guides you through the five steps it takes to give share-level access to a **folder on a volume**.



Shares

SnapServer provides full integration with existing Windows Active Directory domain or Unix LDAP or NIS user and group databases. At the share level, administrators can assign read-write or read-only share access to individual local and Windows domain users and groups for Windows/SMB, AFP, FTP, and HTTP. Administrators can also edit the NFS exports file to control how shares are exported to NFS client machines.



SnapServer supports file access in Windows, Apple, and Unix networks, as well as access via HTTP and FTP. New shares are created by default with full read-write access to all users, subject to the filesystem permissions on the share target directory. The first step to securing a server 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.

Create Shares

To create a new share, at a minimum you need to specify the share name, volume, and folder path. Click **Create Share** on the default **Shares** page to start the process.

The screenshot shows the SnapServer web interface for creating a share. The page title is "SnapServer Snap7654321 · Create Share". The top navigation bar includes "Server", "Network", "Storage", "Security", "Monitor", and "Maintenance". The "Security" tab is active, and the "Shares" sub-tab is selected. Below the navigation, there are links for "Security Guides", "Local Users", "Local Groups", "Security Models", "ID Mapping", and "Home Directories".

The main content area contains the following form fields and options:

- Instruction: "To create a new share, specify a name, volume, and path to a folder."
- Name: Text input field containing "SHARE2".
- Volume: Dropdown menu showing "Volume1".
- Path: Text input field containing "/", with a "Browse" button to its right.
- Description: Text input field, with "(optional)" to its right.
- Security model for path: Dropdown menu showing "Windows/Unix".
- Radio buttons for permissions:
 - Create share with full read and write access for all users
 - Create share with Admin-only access and proceed to Share Access page
- Link: [Advanced Share Properties >>](#)
- Buttons: "Create Share" and "Cancel".

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.

NOTE: The snapshot information at the bottom is only shown if snapshot space has been reserved.

Creating a Share

Creating a share includes 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).
2. Choose the **volume** you need from the drop-down menu.
3. Select from the following **path 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 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**.

- If desired, enter a **description** to clarify the purpose of the share.
- Choose a **security model for path** by selecting **Windows/Unix**, **Windows**, 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 205](#).

- Choose the user-based **share access** option desired for Windows/SMB, AFP, FTP, and HTTP users:

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 user share access. For more information, see [Share Access Behaviors on page 186](#).

NOTE: If selecting the share with Admin-only access option and 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:

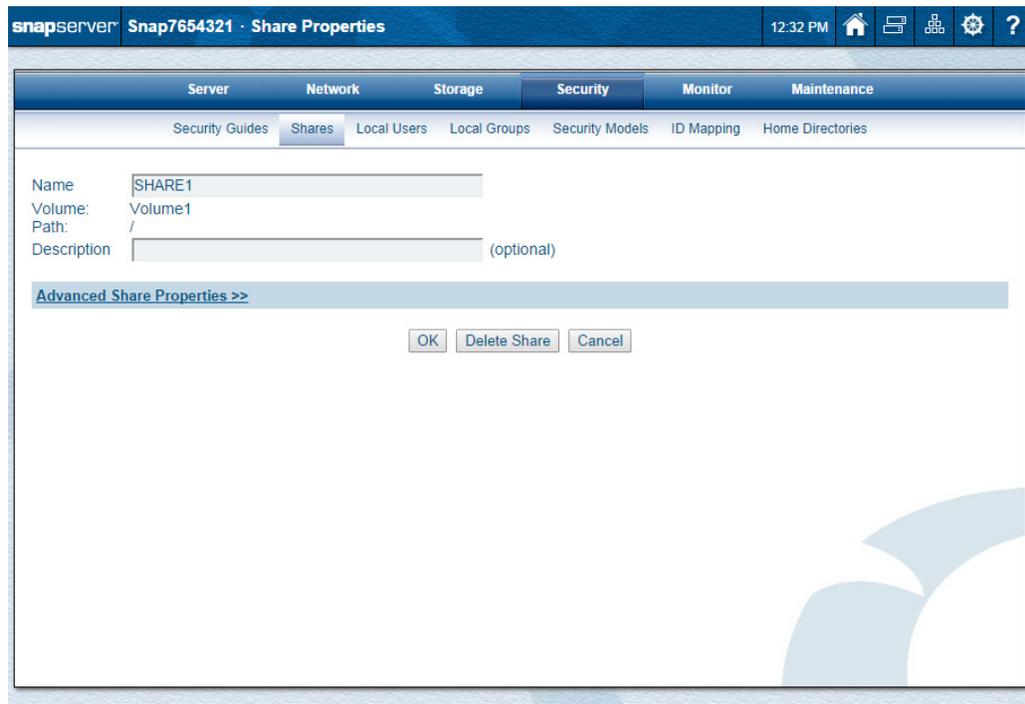
Option	Description
Hide this Share	Select this option if you want the share to be hidden from network browsing using SMB, HTTP/HTTPS, AFP, and FTP protocols (but not NFS).
Protocols	Select the access protocols for the share: Windows (SMB), Linux/Unix (NFS), Apple (AFP), Web (HTTP/HTTPS), and FTP/FTPS. Check all that apply.
Snapshot Share	To create a snapshot share, check the Create Snapshot Share box. Optionally, do either of the following: <ul style="list-style-type: none"> To hide the snapshot share from the SMB, HTTP, AFP, and FTP protocols, check the Hide Snapshot Share box. If you do not want to accept the default name provided, enter a unique name for the Snapshot Share Name field. Use up to 27 alphanumeric characters (including hyphens and spaces).

- Click **Create Share** to complete the process.

Edit Share Properties

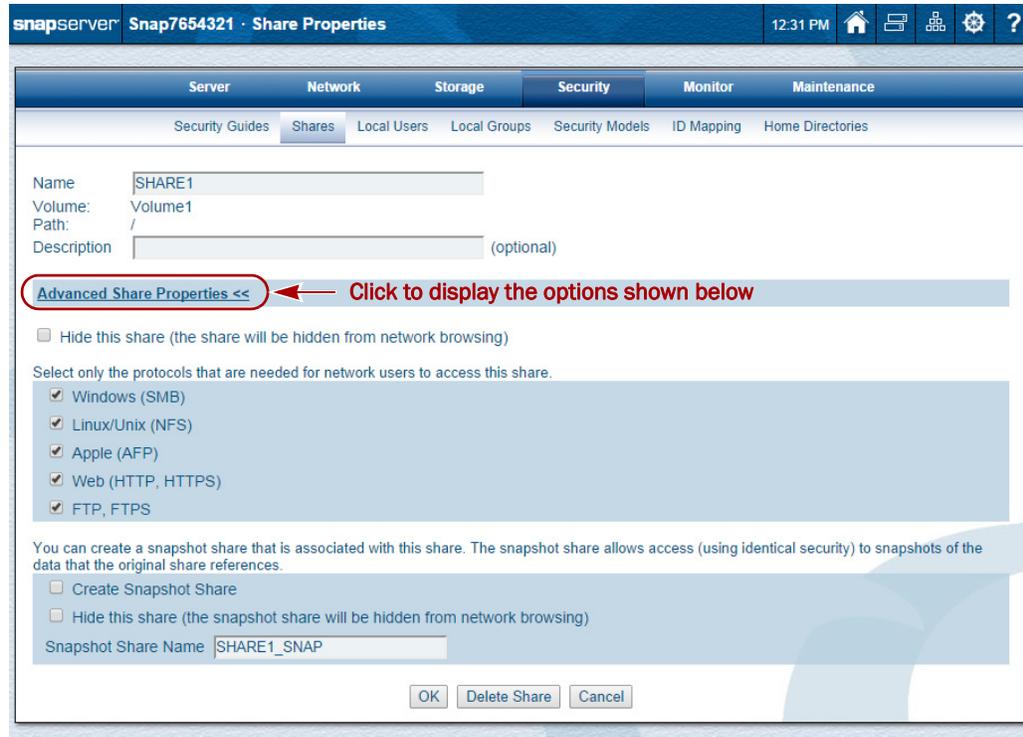
NOTE: You cannot change the volume (or path) of a share once it is created. If you need to change the share's volume, you must delete the share and create a new share on the other volume.

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 in the table).



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.

NOTE: The snapshot information at the bottom is only shown if snapshot space has been reserved.

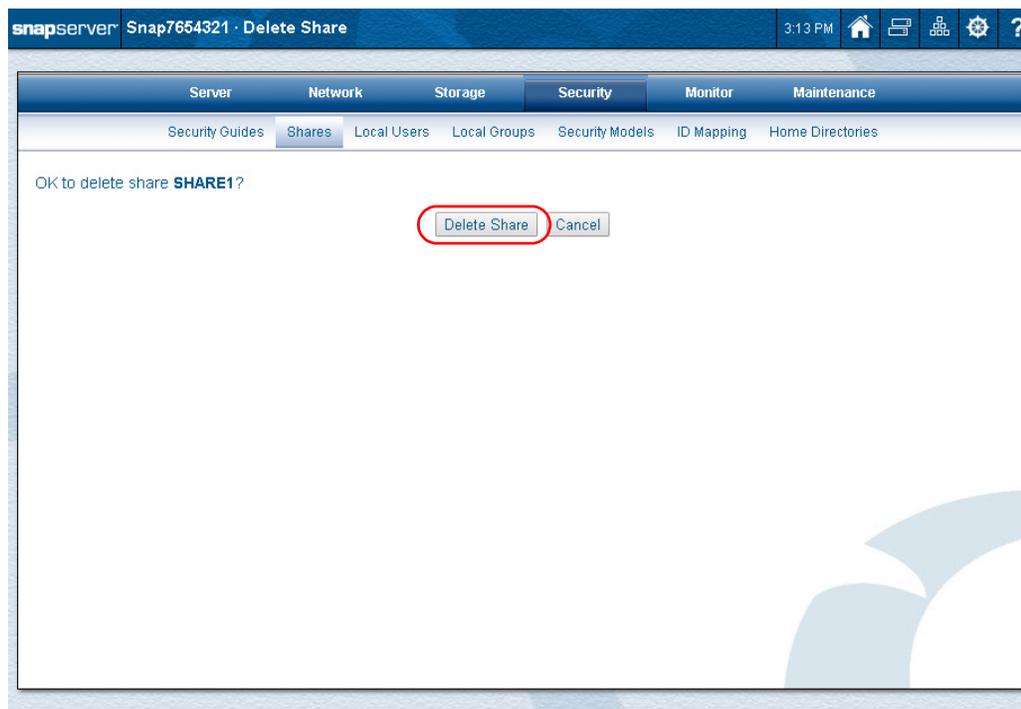


Option	Description
Name	Accept the default share name or enter a new one. If you change the default, observe the following guidelines: <ul style="list-style-type: none"> • Make sure the share name is unique on this server. • To ensure compatibility with all protocols, share names are limited to 27 alphanumeric 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 using SMB, HTTP/HTTPS, AFP, and FTP/FTPS (but not NFS) protocols.
Protocols	Select the access protocols for the share: Windows (SMB), Linux/Unix (NFS), Apple (AFP), Web (HTTP/HTTPS), and FTP/FTPS. Check all that apply.
Snapshot Share	The option that displays depends on whether a snapshot share currently exists. To create a snapshot share, check the Create Snapshot Share box. Optionally, do either of the following: <ul style="list-style-type: none"> • To hide the snapshot share from the SMB, HTTP, AFP, and FTP protocols (but not NFS), check the Hide Snapshot Share box. • If you do not want to accept the default name provided, enter a unique name for the Snapshot Share Name field. Use up to 27 alphanumeric characters (including hyphens and spaces). To remove an existing snapshot share, check the Remove Snapshot Share box.

Delete Shares

To delete a share, go to **Security > Shares > Share Properties** (displayed by clicking the share name in the table).

1. At the **Delete Share** page, click **Delete Share**.
2. At the confirmation page, click **Delete Share** 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 is displayed. 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 193 for information on creating new users.

The top screenshot shows the SnapServer GUI for 'Snap1234567 · Shares'. It features a navigation bar with 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under 'Security', there are sub-tabs for 'Security Guides', 'Shares', 'Local Users', 'Local Groups', 'Security Models', 'ID Mapping', and 'Home Directories'. A table lists shares with columns for 'Share', 'Volume', 'Path', 'Access', 'NFS Access', 'Protocols', and 'Attributes'. The 'SHARE1' row shows 'VOLD' as the volume and '/' as the path. The 'Access' column contains an 'Open' button, which is circled in red. Below the table, there are buttons for 'Create Share', 'Refresh', and 'Close'. An 'Important Security Note' is also present.

The bottom screenshot shows the 'Snap7654321 · Share Access' page. It has the same navigation structure. The main content area is titled 'Access to share SHARE1. Use Ctrl-click to select multiple users/groups. Use Command-click for Mac.' It is divided into two panes: 'Users and groups with specific access to share. (1)' and 'Search for users and groups.' The first pane lists 'AllUsers' and has a 'Remove' button. The second pane has a search box and a dropdown menu set to 'Full Access'. At the bottom, there are 'OK' and 'Cancel' buttons.

Share Access Behaviors

Administrators tasked with devising security policies for 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, HTTP, AFP, and 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 both 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. Consider the following examples:

Example A: More restrictive file-level access is given precedence over more permissive share-level access.

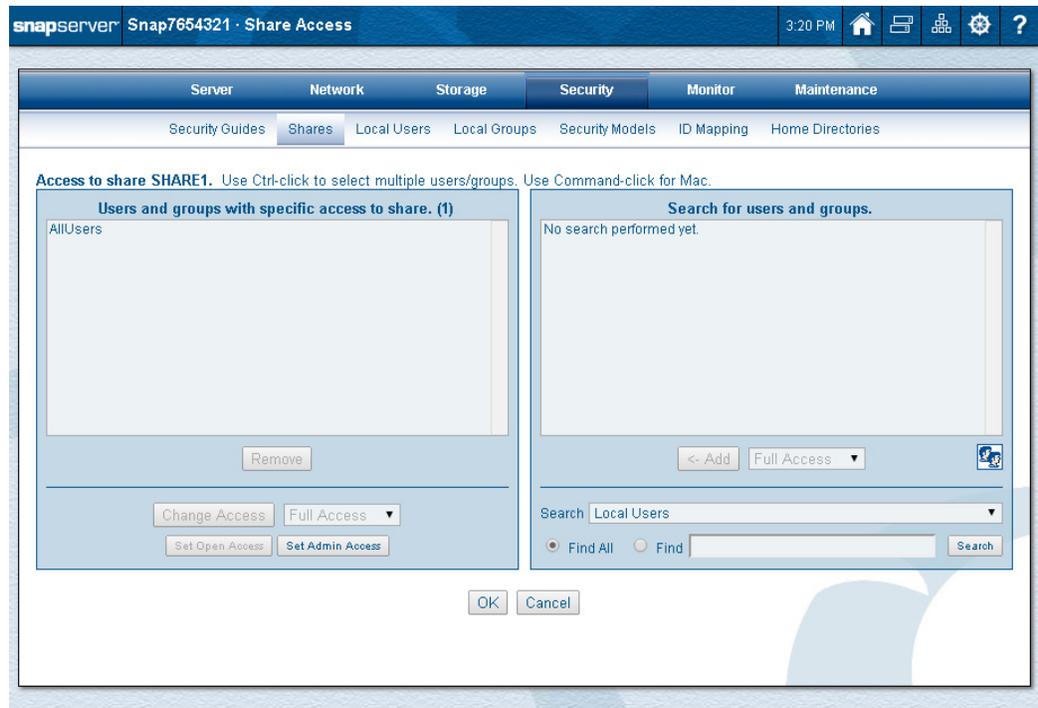
Share Level	File Level	Result
Full control	Read-only to File A	Full control over all directories and files in SHARE1 <i>except</i> where a more restrictive file-level permission applies. The user has read-only access to File A.

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, HTTP, and FTP 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 190 for more information.



User-based share access permissions apply to users connecting over SMB, AFP, HTTP, or 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	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 and 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 is 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`).

c. Click **Search** to display any matches.

After you click **Search**, another authentication prompt may be presented to authenticate with the remote domain.

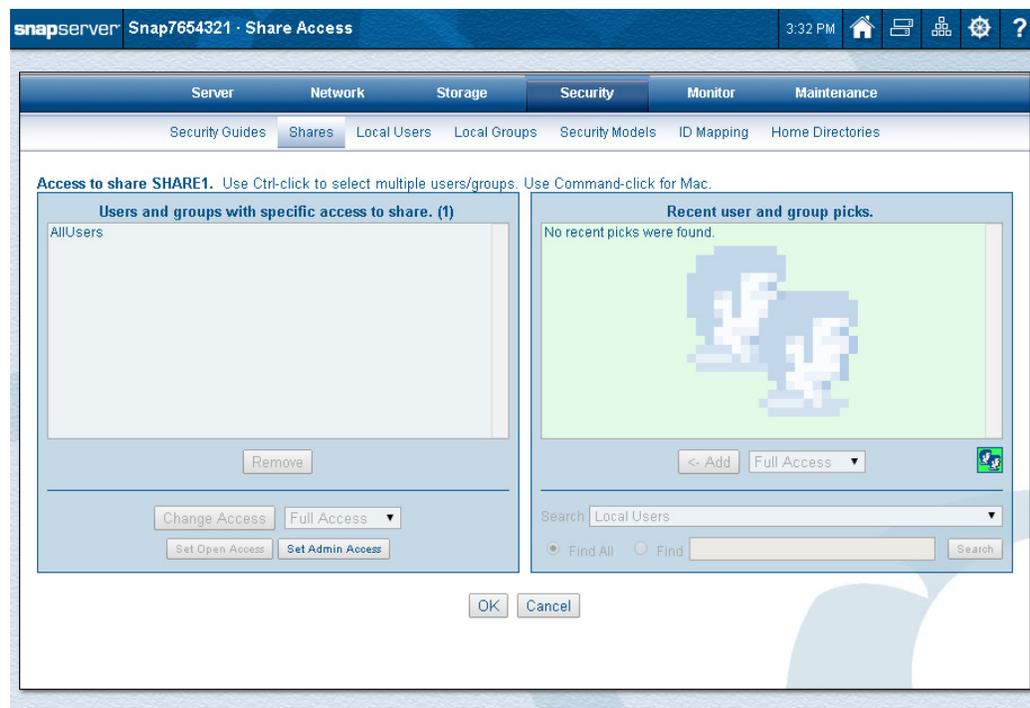
d. Select one or more **names** in the list.

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 again to return to the normal search box.



3. To **remove** share access permissions for a user or group:

- Select one or more **users or groups** in the left box.
- 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 finally click **Change Access**.
5. To quickly specify either Open or Admin-only **access** for the entire share, click either **Set Open Access** or **Set Admin Access**.
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.

To configure NFS access, click the link shown in the **NFS Access** column for the share you want to configure. 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 top screenshot shows the 'Shares' page for Snap7654321. It features a table with columns: Share, Volume, Path, Access, NFS Access, Protocols, and Attributes. The 'NFS Access' column for 'SHARE1' is circled in red. Below the table, there is an 'Important Security Note' and a link to 'View online help for more information'.

The bottom screenshot shows the 'NFS Share Access' page for Snap7654321. It includes a section for 'Options' with checkboxes for 'SnapServer default options' (checked) and 'Read-only (ro,async)'. There is an 'NFS Host' input field and an 'Add Host' button. Below this is a text area for 'NFS access (exports) for share SHARE1:' containing the text: `*(rw,insecure,async,root_squash,no_all_squash)`. At the bottom, there are 'OK' and 'Cancel' buttons.

The NFS access text box is a window into the client access entries in the *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 **Shares** page.

NOTE: The syntax used in this file is equivalent to standard Linux exports file syntax. If the server detects any errors in syntax, a warning message appears. You can choose to correct or ignore the error warning.

The Exports File Default Options. The SnapServer default setting provides read-write access to all NFS clients.

```
*(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 SnapServer default.
no_root_squash	no_root_squash means that if root is logged in on your client 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 Exports File Default Options on page 191](#)) 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**).

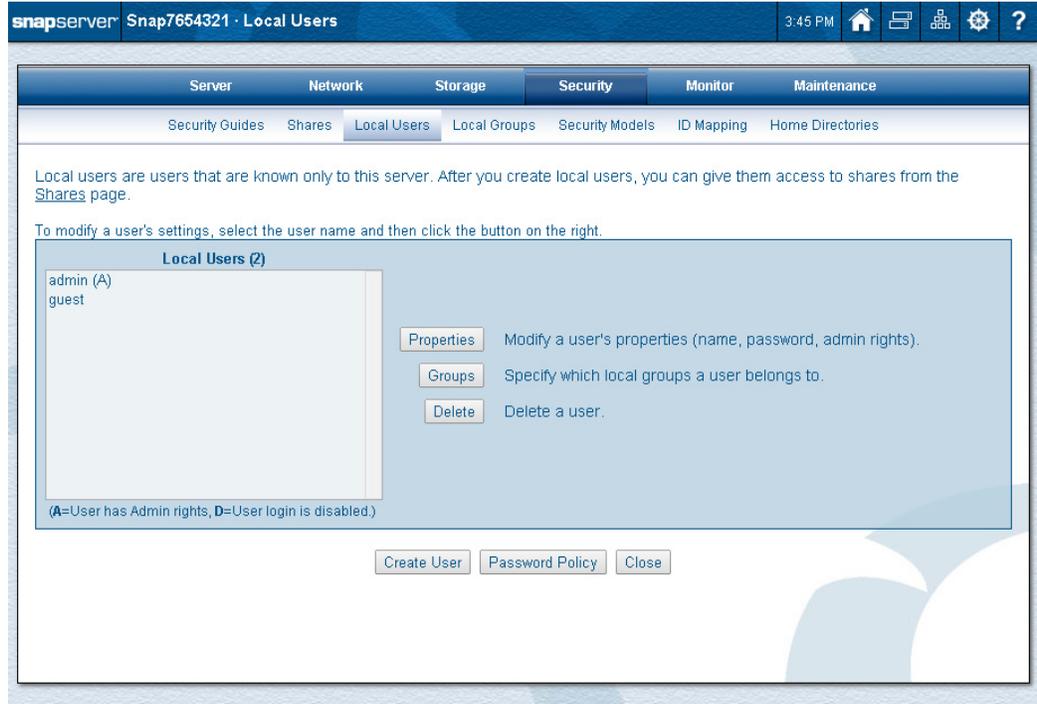
These entries can be independently removed, added, and modified on each NFS-enabled share.

Using the Add Host Option. 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 provides all the options to manage local users. Local users are users that are known only to the server being accessed. Each SnapServer 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 **Create** to create a new user on this server. Enter the user data, select any special options, and click **Create User** again.

The screenshot shows the 'Create Local User' page in the SnapServer web interface. The page title is 'Snap7654321 · Create Local User'. The navigation menu includes 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'Security' menu is expanded, showing 'Security Guides', 'Shares', 'Local Users', 'Local Groups', 'Security Models', 'ID Mapping', and 'Home Directories'. The main content area contains the following form:

To create a new user, specify a name, password and user ID (UID).

Name:

Full Name: (optional)

Password:

Confirm Password:

User ID (UID):

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:

To Create a Local User

1. On the **Local Users** page, click **Create User**.
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.
Confirm Password	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 175 .

Option	Description
Disable User Login	Check this box to disable the user login. The user's information will remain in the system, but login rights are denied. The user login can be re-enabled by clearing the box. This box 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. Check this box to exempt this user from password expiration and character requirement policies.
Grant Admin Rights To This User	Check this box 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

Highlight a user and click **Properties** to open the **Local User Properties** page to make changes to the user's full name, password, or user ID (UID). Note that the UID cannot be changed for the built-in admin user.

The screenshot shows the 'Local User Properties' dialog box in the SnapServer web interface. The dialog has a title bar with 'snapserver Snap7654321 · Local User Properties' and a status bar with '9:00 AM' and navigation icons. The main content area has tabs for 'Server', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under the 'Security' tab, there are sub-tabs for 'Security Guides', 'Shares', 'Local Users', 'Local Groups', 'Security Models', 'ID Mapping', and 'Home Directories'. The 'Local Users' sub-tab is active. The form contains the following fields and options:

- Name: Fred
- Full Name: Fredrick Sandstone (optional)
- Password: (empty)
- Confirm Password: (empty)
- User ID (UID): 18001
- Disable user login
- Exempt this user from password expiration and character requirements. (highlighted with a red box and an arrow pointing to it from the text 'Only shown if Password Policy enabled')
- Grant admin rights to this user. (A local user with admin rights will be able to access this Web Management Interface.)

At the bottom of the dialog are 'OK' and 'Cancel' buttons.

To Edit Local User Properties

1. On the **Local Users** page, highlight the user you want to edit and click **Properties**.
2. On the **Local User Properties** page that opens, enter or change any of the **information**:

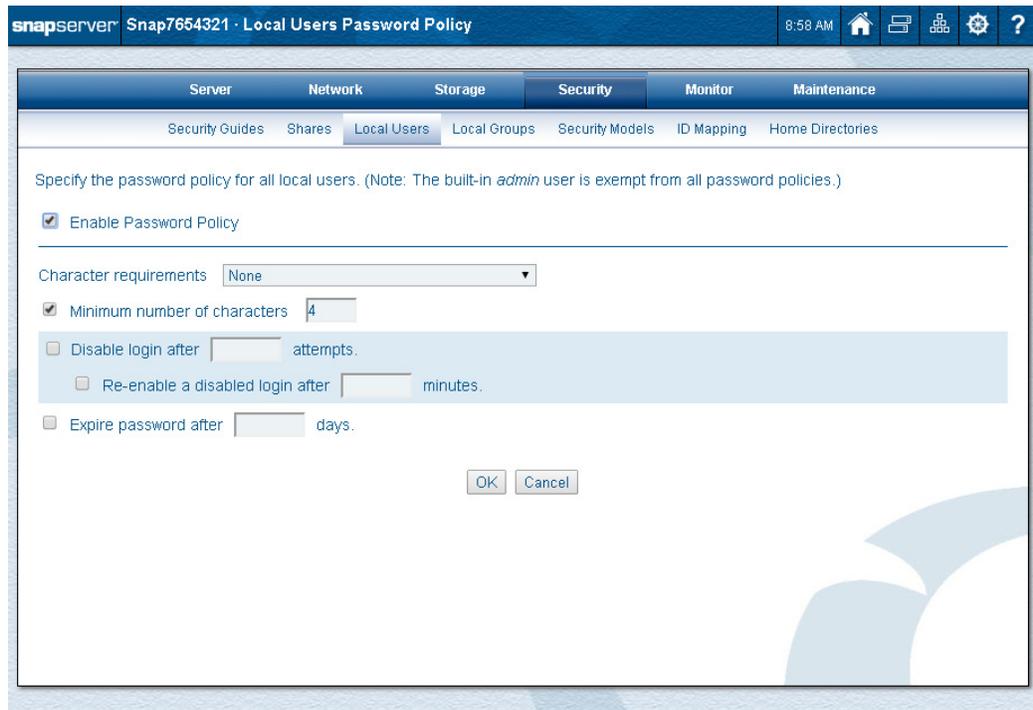
Option	Description
Name	NOTE: Cannot be modified. Instead, delete and recreate the user with the same UID if you need to change the user name.
Full Name	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
Password	Passwords are case-sensitive. To keep the existing password, leave this field blank.
Password Verify	Type the chosen password again for verification. To keep the existing password, leave this field blank.
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 175 . 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.
Disable User Login	Check this box to disable the user login. The user's information will remain in the system, but login rights are denied. The user login can be re-enabled by clearing the box. This box 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	NOTE: This box is only visible if Password Policy is enabled. Check this box to exempt this user from password expiration and character requirement policies.
Grant Admin Rights To This User	Check this box to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click **OK**.

Local 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 on page 194](#) for procedures to set password policy for local users. Also, the built-in *admin* user is automatically exempt from all password policies.

Use **Password Policy** to make changes to the local user password settings.



To Set Password Policy for Local Users

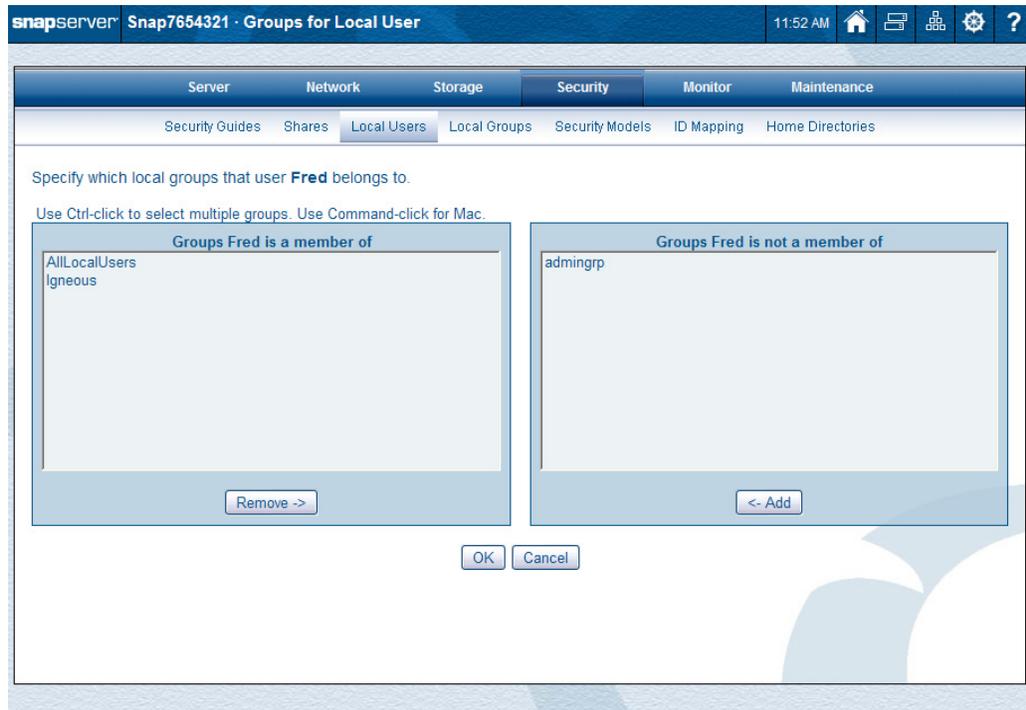
1. On the **Local Users** page, click **Password Policy**.
2. On the **Local Users Password Policy** page, check the **Enable Password Policy** box.
3. Enter the following **information**:

Option	Description
Character Requirements	Select the alpha/numeric/special character requirements for the password from the drop-down list.
Minimum Number of Characters	Check this box to enable the policy, then enter the minimum number of characters required for the password.
Disable Login After <i>n</i> Attempts	Check this box 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 box 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 box and enter a time period after which the system will allow the user to log in again. This saves the administrator from having to manually re-enable the user.
Expire Password After <i>n</i> Days	Check this box 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://<server_name>/changepassword">http://<server_name>/changepassword .

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

Assign User to Group

Use the **Groups for Local User** page (**Security > Local Users > Groups**) to make changes to a local group membership.

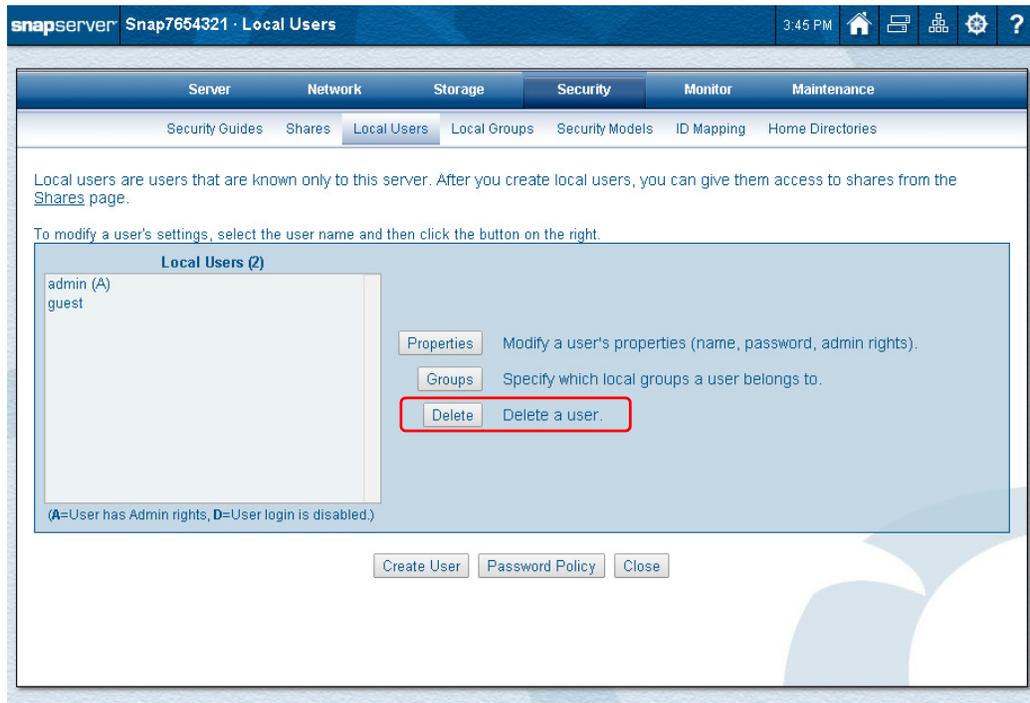


Add or Remove Users from Groups

1. On the **Local User** page, select a **user**.
2. Click **Groups**.
The group settings for the selected user are shown.
3. To make a **change**:
 - To add the user to a group, from the list on the **right**, select a **group name** and click **<-Add**.
 - To remove the user from a group, from the list on the **left**, select the **group name** and click **Remove->**.
4. Click **OK** to save your changes.

Delete Local User

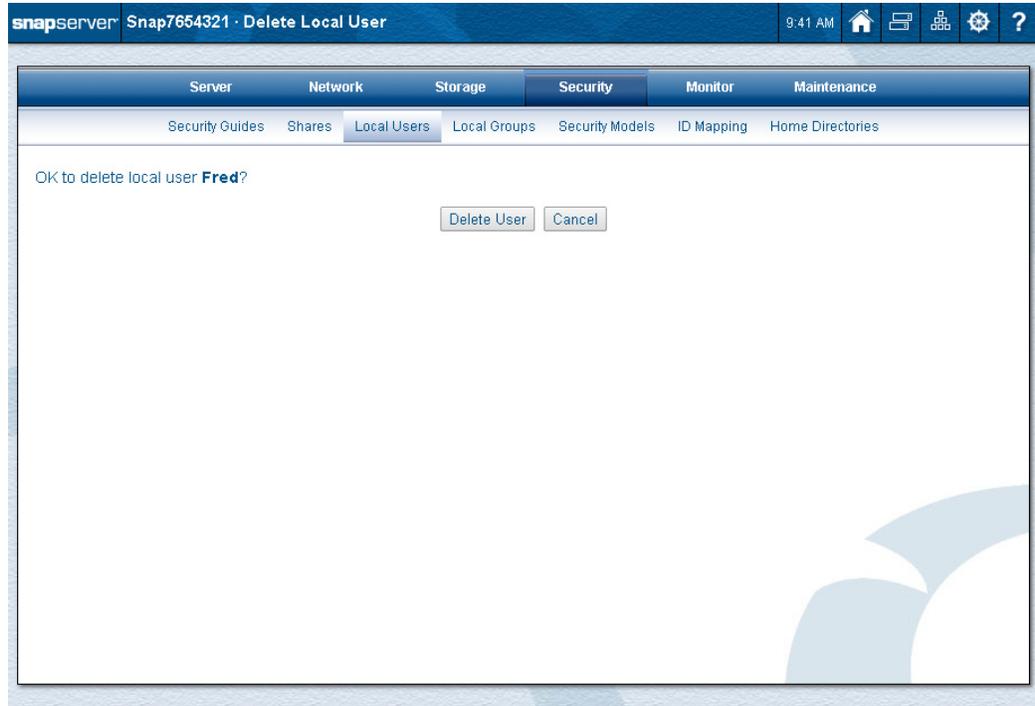
On the **Local Users** page, use the following process to remove a user.



To Delete a Local User

1. On the **Local Users** page, select the user to be deleted.
2. Click **Delete**.

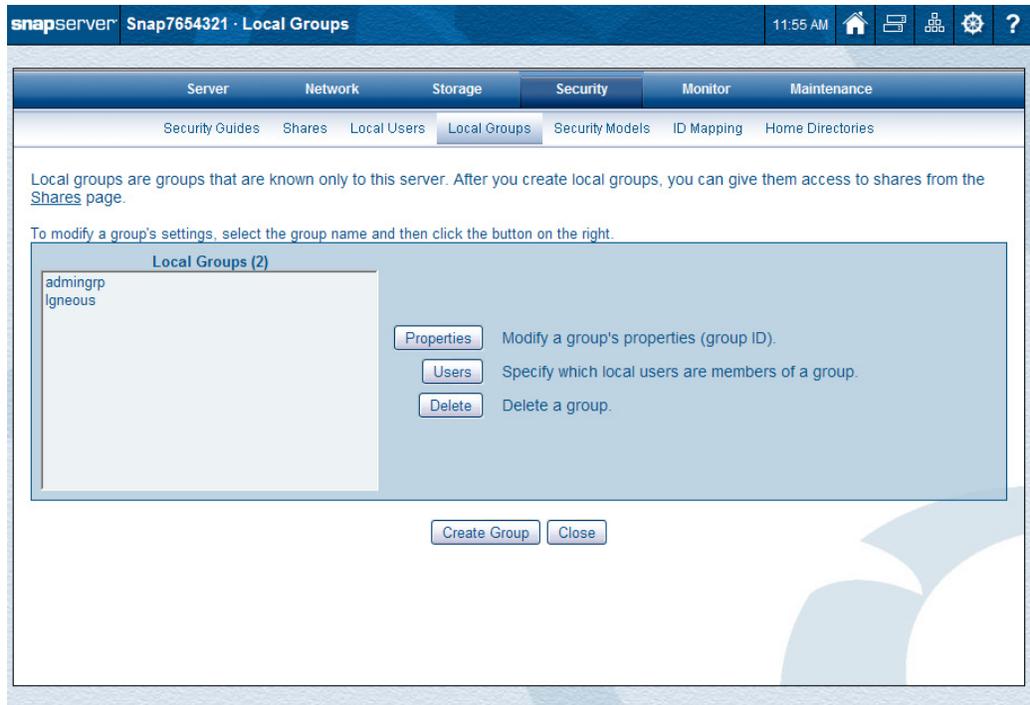
The confirmation page is displayed.



3. Click **Delete User** to delete the selected user.

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 SnapServer comes with one predefined group (**admingrp**).



Create New Group

Use **Create** to create a new group on this server. Options include the group name and changing the Group ID (GID).

To Create a New Local Group

1. On the **Local Groups** page, click **Create Group** to access the **Create Local Group** page.

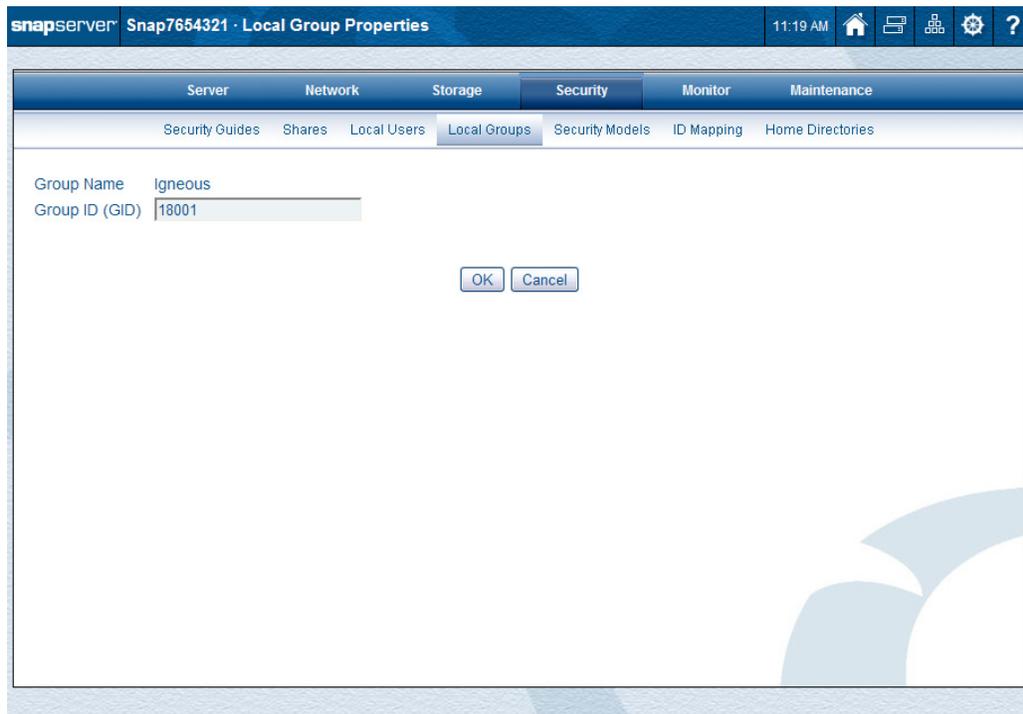
2. Enter the following **information**:

Option	Description
Group Name	Use up to 31 characters (alphanumeric and the underscore only).
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 175 .

3. Click **Create Group** when finished.
4. The **Users for Local Group** page is displayed, allowing you to immediately add users to your new group.
5. Click **OK** when you are finished adding users.

Edit Group Properties

Use **Properties** to open the **Local Group Properties** page to make changes to the options there.



To Edit Local Group Properties

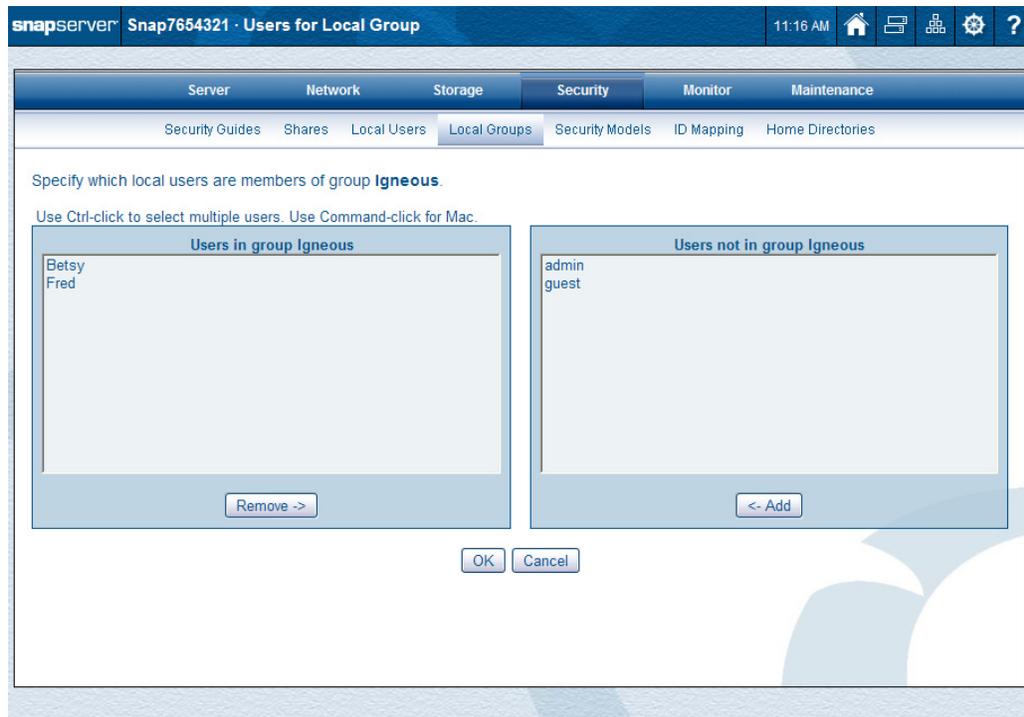
1. On the **Local Groups** page, select the group you want to edit and click **Properties**.
2. On the **Local Groups Properties** page that opens, you can change the **GID**. For information on available UID ranges, see [User and Group ID Assignments on page 175](#).

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 the membership of a local group.

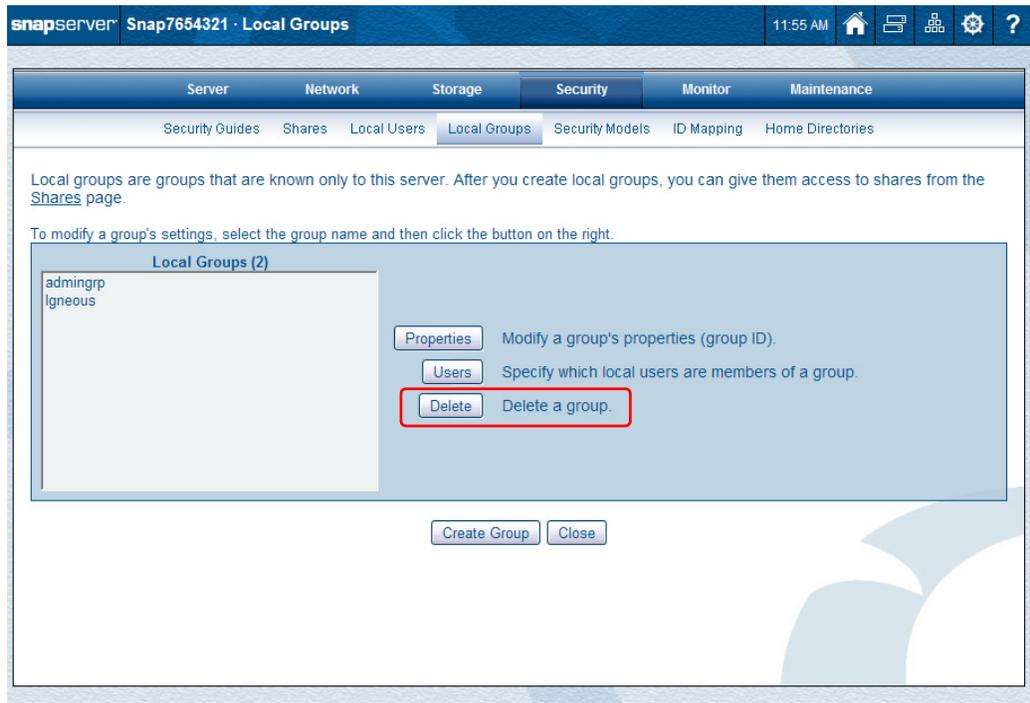


Add or Remove Group Users

1. On the **Local Groups** page, select a group name and click **Users**.
2. To make a **change**:
 - To add the user to a group, from the list on the **right**, select a **user name** and click **<-Add**.
 - To delete the user from a group, from the list on the **left**, select the **user name** and click **Remove->**.
3. Click **OK** when finished.

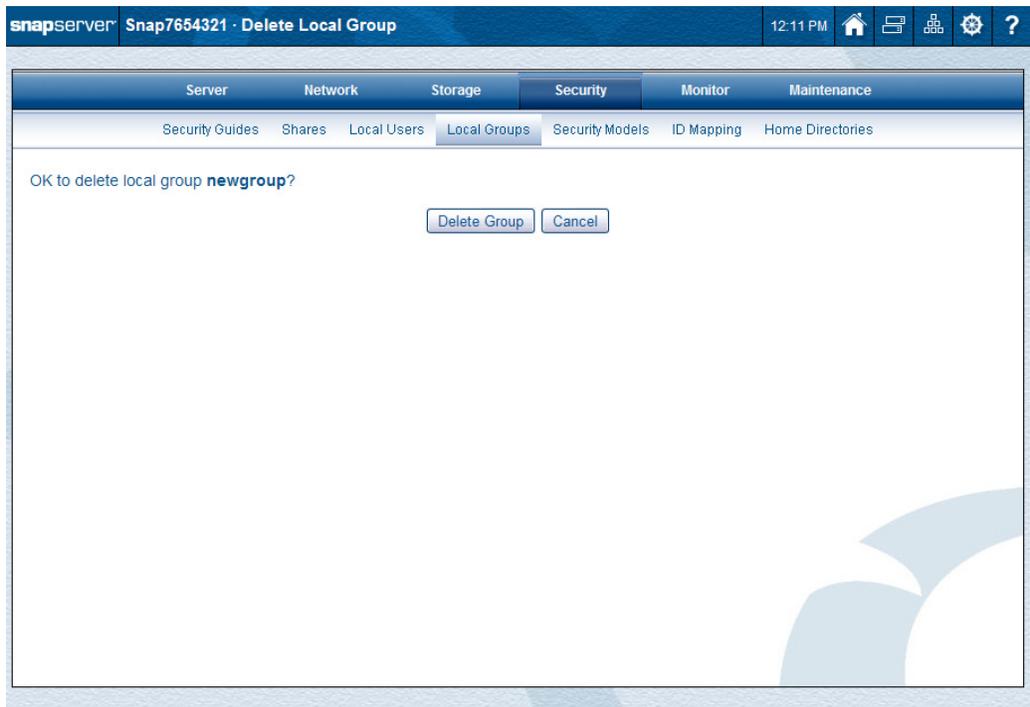
Delete Group

On the **Local Groups** page, use the following process to remove a group.



Delete a Group

1. On the **Local Groups** page, select the **group** to be deleted and click **Delete**. The delete confirmation page is displayed.



2. Click **Delete Group** to delete the selected group (or **Cancel** to cancel the deletion).

Security Models

There are three file-level security models that can be used by a SnapServer:

- **Windows/Unix**
- **Windows**
- **Unix.**

For Traditional RAID, the security model can be configured on volumes and the folders created in the root of the volumes. For DynamicRAID, the security model can only be configured on the volumes.

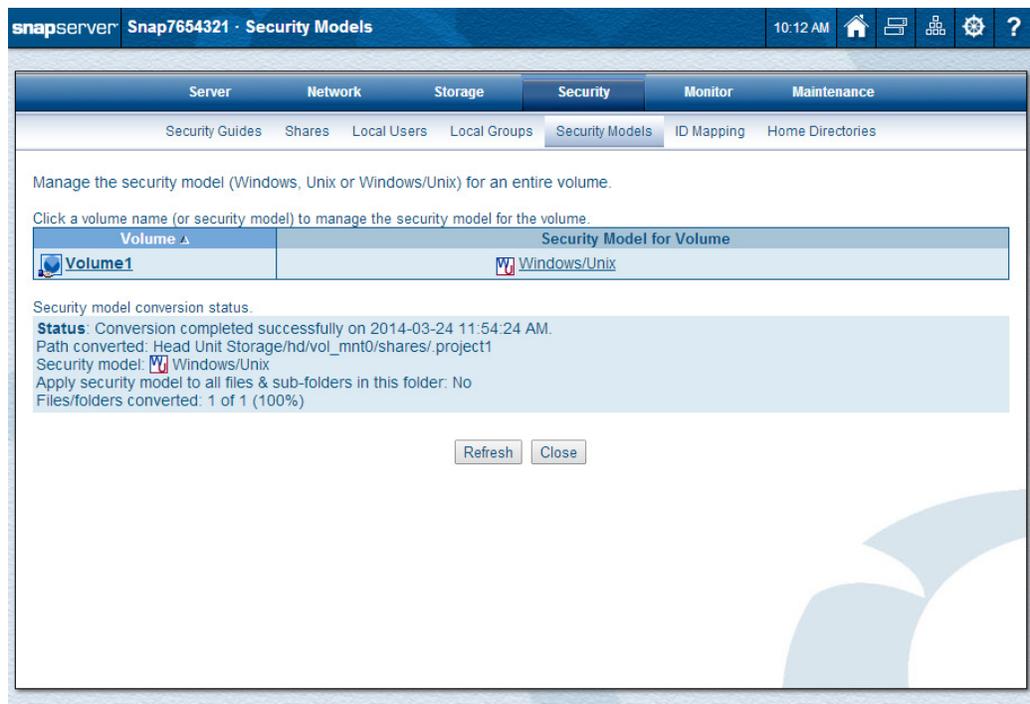
The security model determines the rules regarding which security personality is present on files and folders created by the various protocols and clients, and whether the personality of files and folders can be changed by changing permissions.

NOTE: 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/Unix** security may have a Unix personality).

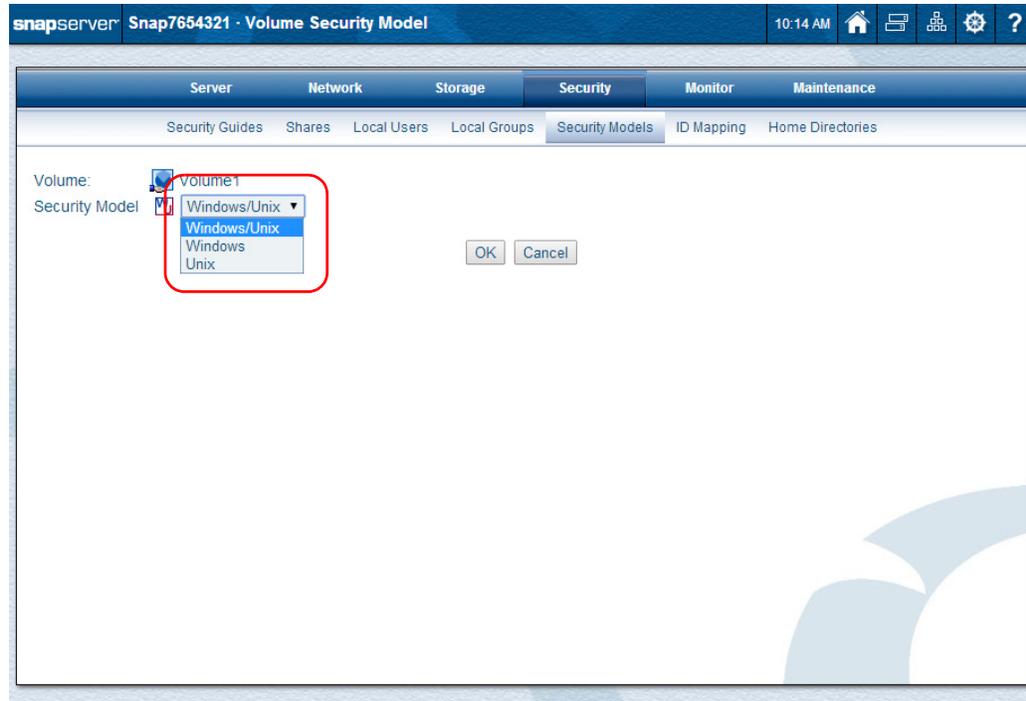
For more information about security models, see [Appendix B, Security and Access](#).

Managing Volume Security Models

1. Select **Security > Security Models**.



2. Click the **Security Model for Volume** name (**Windows/Unix**, **Windows**, or **Unix**). Clicking the **Volume** name does the same thing.
3. From the drop-down list, select the **security model type** desired and click **OK**.



4. At the confirmation message, click **Apply Security Model**.

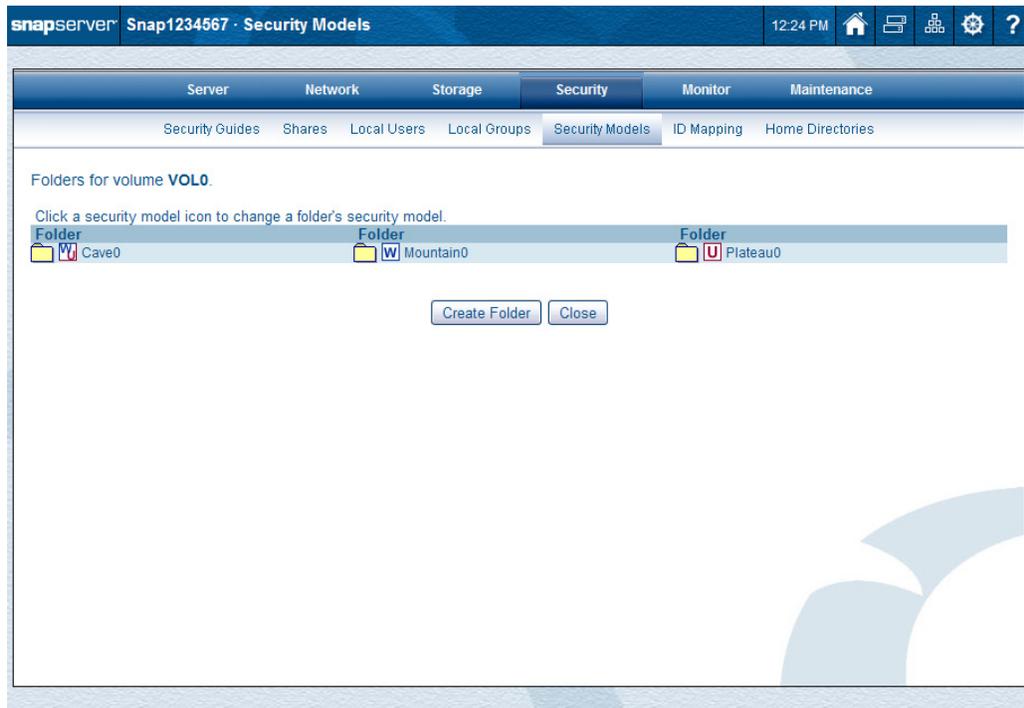
If there are files and directories under the volume, you are prompted whether you want to recursively apply the change. This resets permissions on all files and directories to make them accessible by all users, and configured for the Windows personality (Windows and Windows/UNIX security models) or UNIX personality (UNIX security model). When done, the main page displays a conversion status.

Managing Folder Security Models in Traditional RAID

NOTE: This is only available with Traditional RAID.

1. Select **Security > Security Models**.

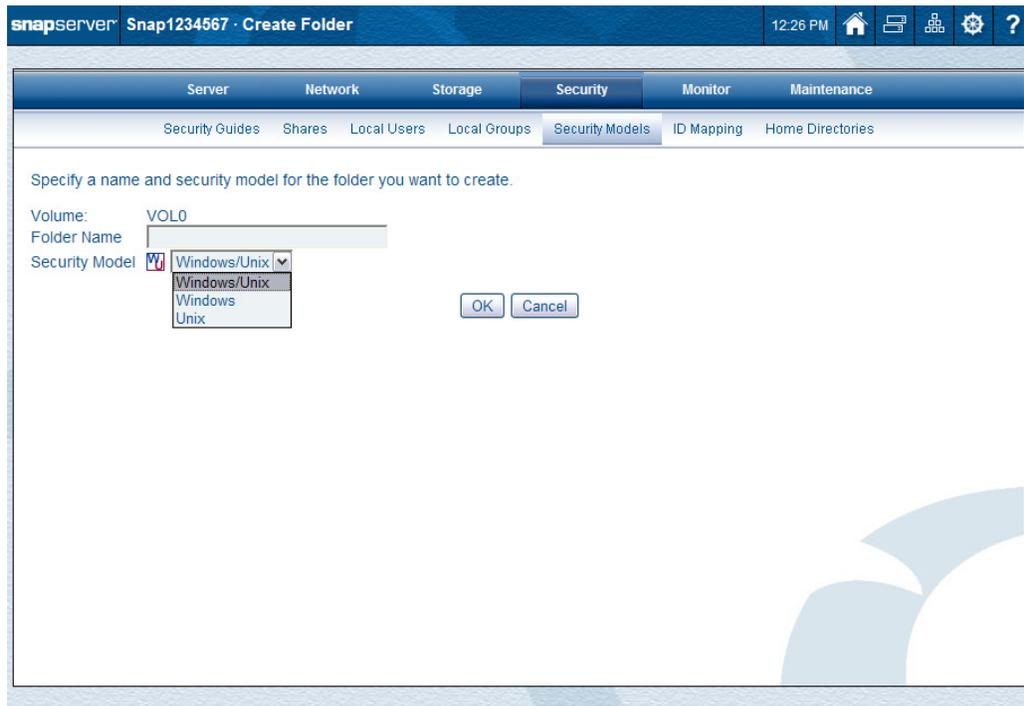
2. Click the **volume name**.



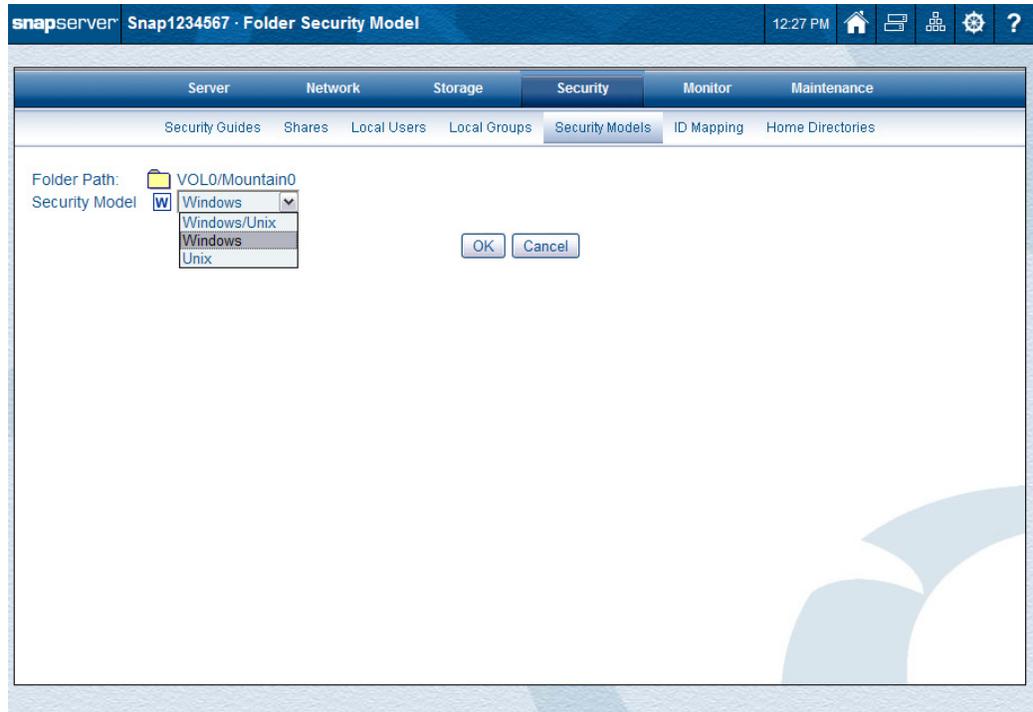
3. At the **Security Models** folder page, do one of the following:

- **Create a new folder** with a specific security model:

Click **Create Folder**, enter the **folder name**, select the **security model type** from the drop-down list, and click **OK**.



- **Change the security model** of a folder:
Click the security model **icon** (**W/U, W, or U**) of the folder, select the **security model type** from the drop-down list, and click **OK**. At the confirmation message, click **Apply Security Model**.

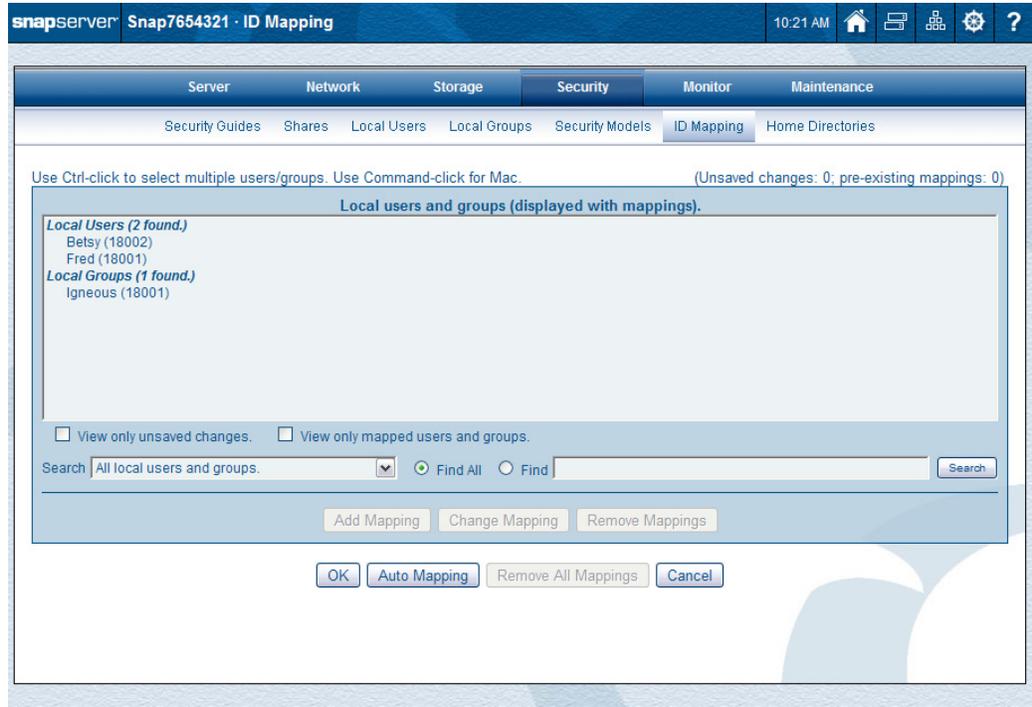


NOTE: If there are files and directories under the volume, you are prompted whether you want to recursively apply the change.

4. At the results page, click **Close**.

ID Mapping

ID mapping allows users and groups that exist on Windows domains to share user and group IDs with local, LDAP, or NIS users and groups. This results in the same permissions and quota consumption applying to both users and groups in an ID-mapped pair.



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, LDAP, or NIS user or group from the displayed list on the default page. You can then use **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.

Options 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:

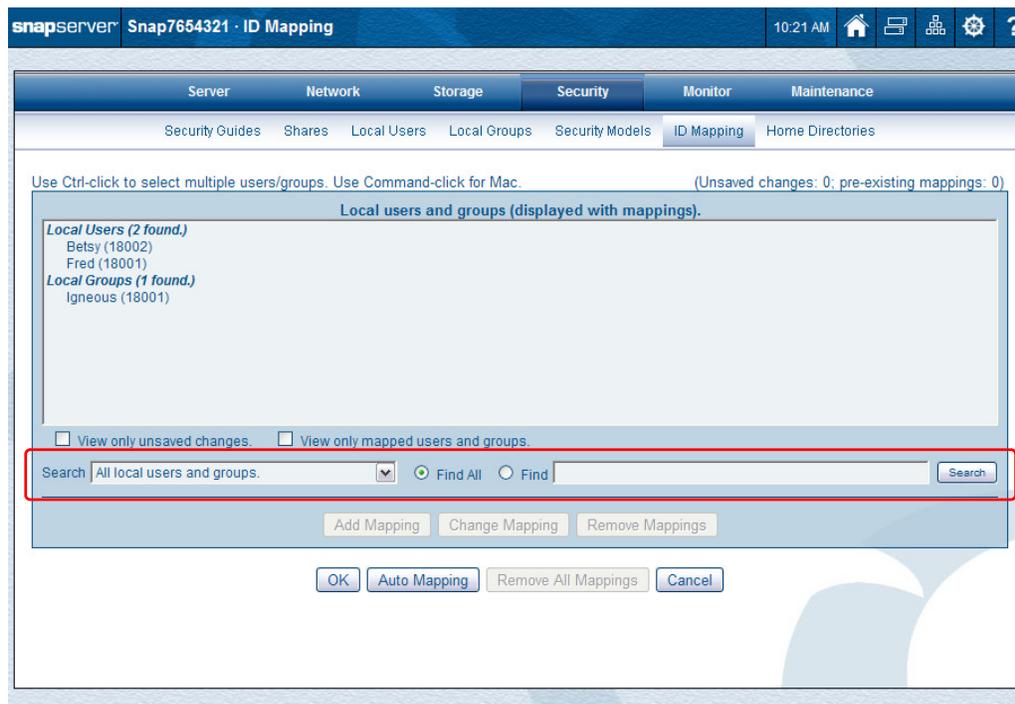
- Check **View only unsaved changes** to display only mapping changes that have not yet been applied.
- Check **View only mapped users and groups** to display only local, LDAP, or NIS users and groups that have been mapped to a Windows domain user or group.

Add Mapping

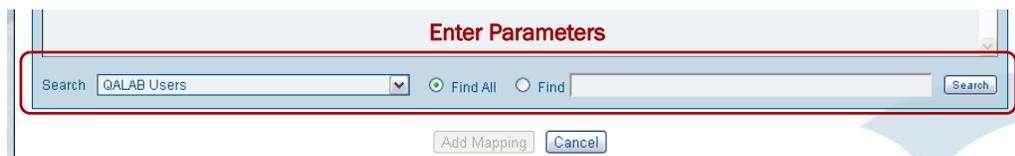
NOTE: Adding or changing an ID mapping requires that the server be joined to a Windows Active Directory domain.

Follow this procedure to map a user or group:

1. If the desired user or group to be mapped to does not appear in the **ID Mapping** page list, use the **search option** to locate it.



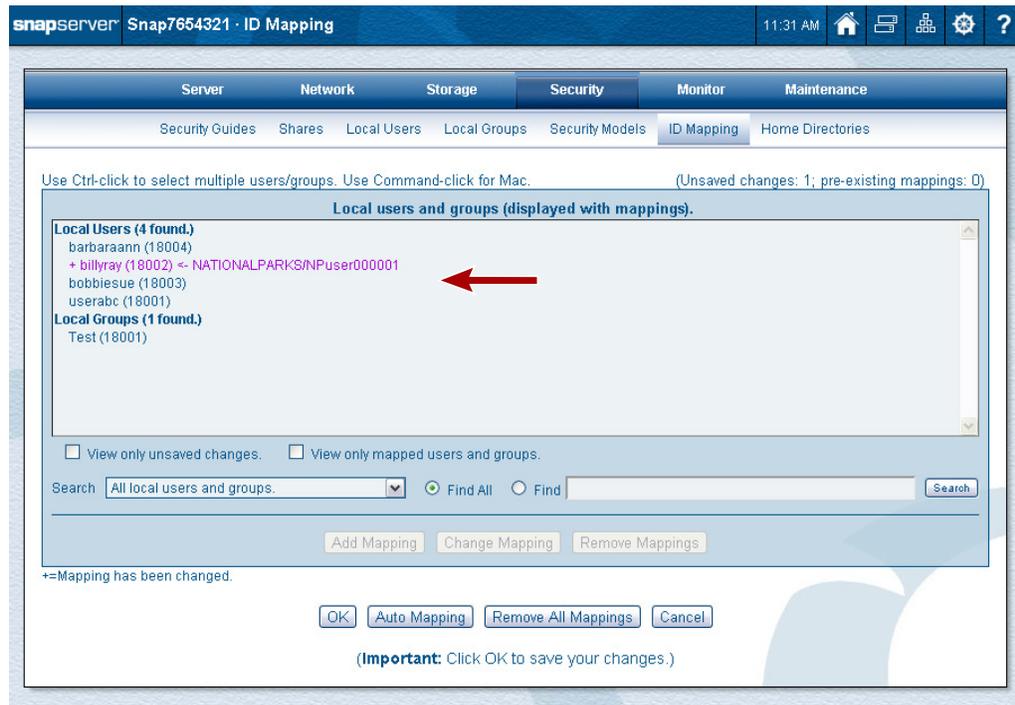
- a. At the bottom of the list, using the **Search** drop-down list, select the local, LDAP, or NIS user or group **list** to be searched.
 - b. Select **Find** and enter the **search string** (or select **Find All**).
Enter the exact **name** (or a string with a wild card "*" before or after) as the search string.
 - 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, to find the user/group you want to map to, select the Windows domain **user or group list**, the scope of the search, enter a search string if needed, and click **Search**.
 - To search for a specific user or group, use either **Find All** or a **Find** search string (wild card "*" before or after is 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 **Find** search string (using the first few letters 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/group** to which you want to map the local, LDAP, or NIS user, and click **Add Mapping**.

The mapping result is shown on the default page with the users/groups that were mapped in purple with a plus (+) in front of their name.



NOTE: To display only changes that have not yet been applied, check the **View only unsaved changes** box. To display only local or NIS users/groups that have been mapped to a Windows domain user or group, check **View only mapped users and groups** box.

5. Repeat [Steps 1–4](#) to add **additional mappings**.
6. Click **OK** to save changes (or **Cancel** to reset).
7. When done with all your mappings, click **OK** to activate them.
8. At the confirmation page, click **Save Changes**.
9. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

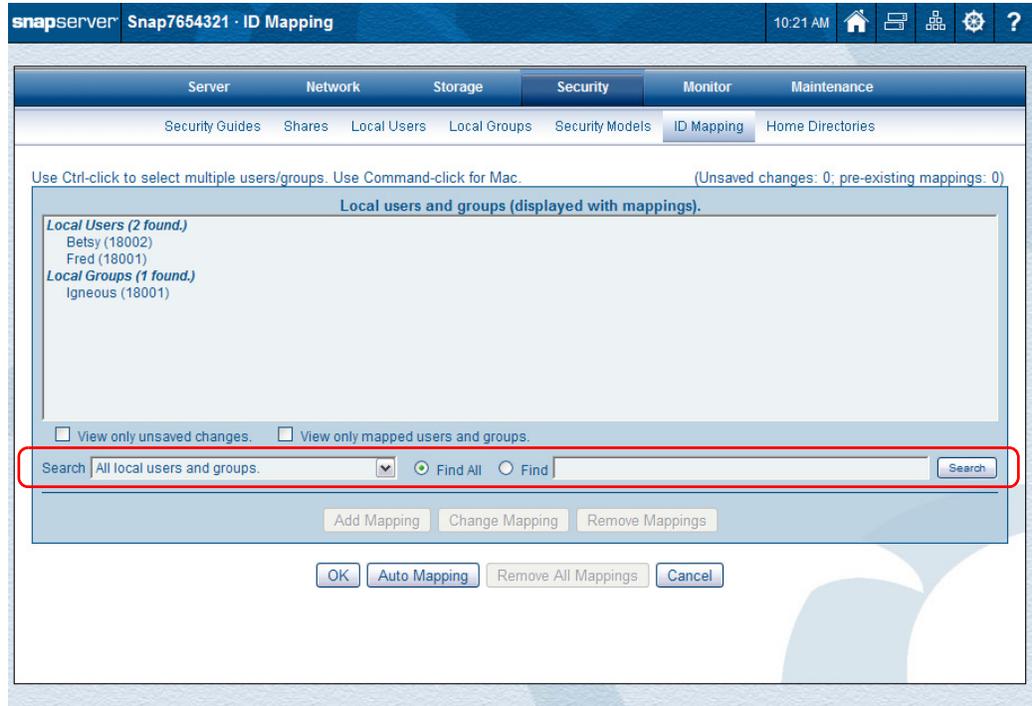
 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

See [Filesystem Updates on page 221](#) for more details.

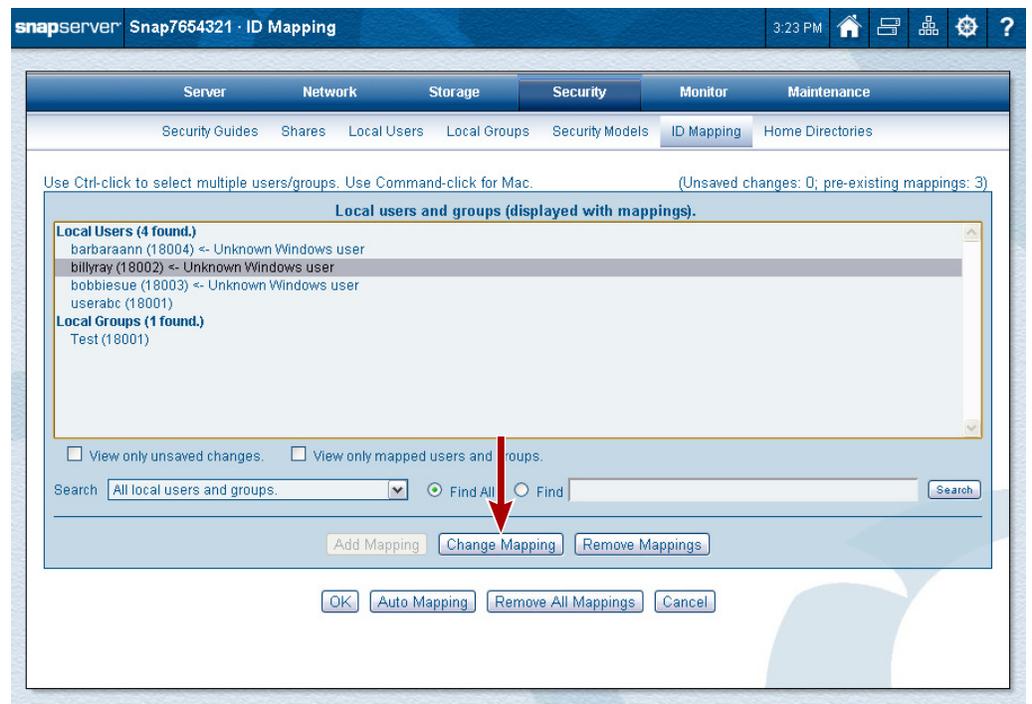
Change Mapping

To map an already mapped local, LDAP, or NIS user or group to a different Windows domain user or group, follow these steps:

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, LDAP, or NIS user or group **list** to be searched.
 - b. Select **Find** and enter the **search string** (or select **Find All**).
Enter the exact **name** (or a string with a wild card "*" before or after) as the search string.
 - c. Click **Search** to display any matches.
2. Select a mapped **user/group** to be changed and click **Change Mapping**.



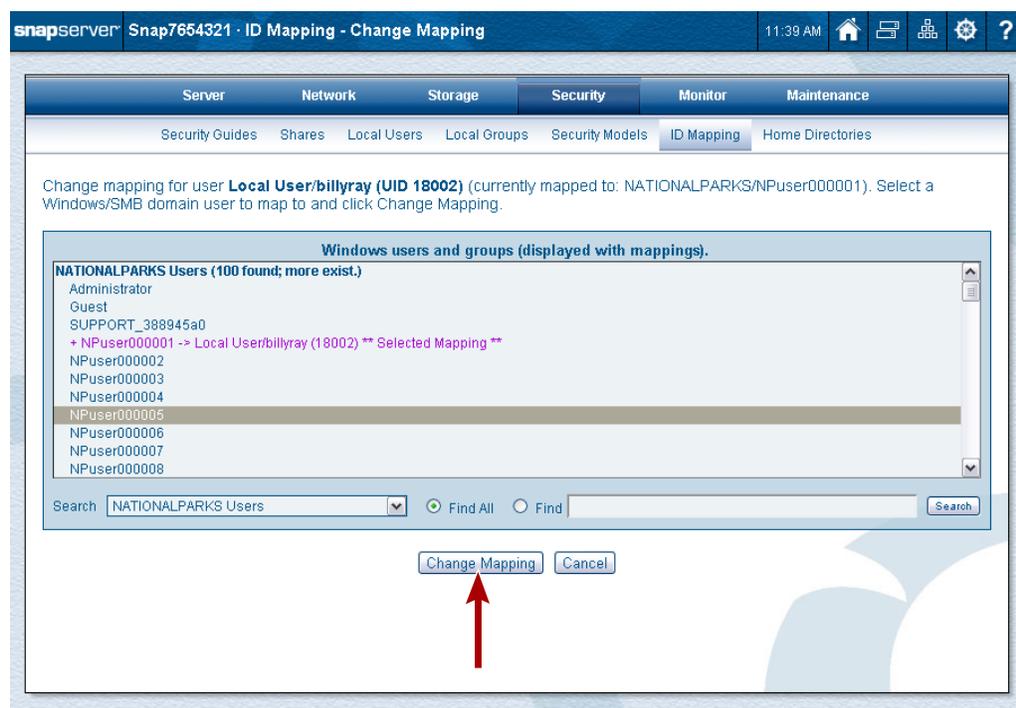
- At the **Change Mapping** page, to find the user/group you want to map to, select the Windows domain **user or group list**, the scope of the search, enter a search string if needed, and click **Search**.
 - To search for a specific user or group, use either **Find All** or a **Find** search string (wild card "*" before or after is 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 **Find** search string (using the first few letters 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.
- From the search results, select the Windows domain **user/group** you want to re-map the local, LDAP, or NIS user to and click **Change Mapping**.



5. Repeat [Steps 1–4](#) until all **changes** are made.
6. Click **OK** to save changes (or **Cancel** to reset).
7. When done with all your mappings, click **OK** to activate them.
8. At the confirmation page, click **Save Changes**.
9. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.



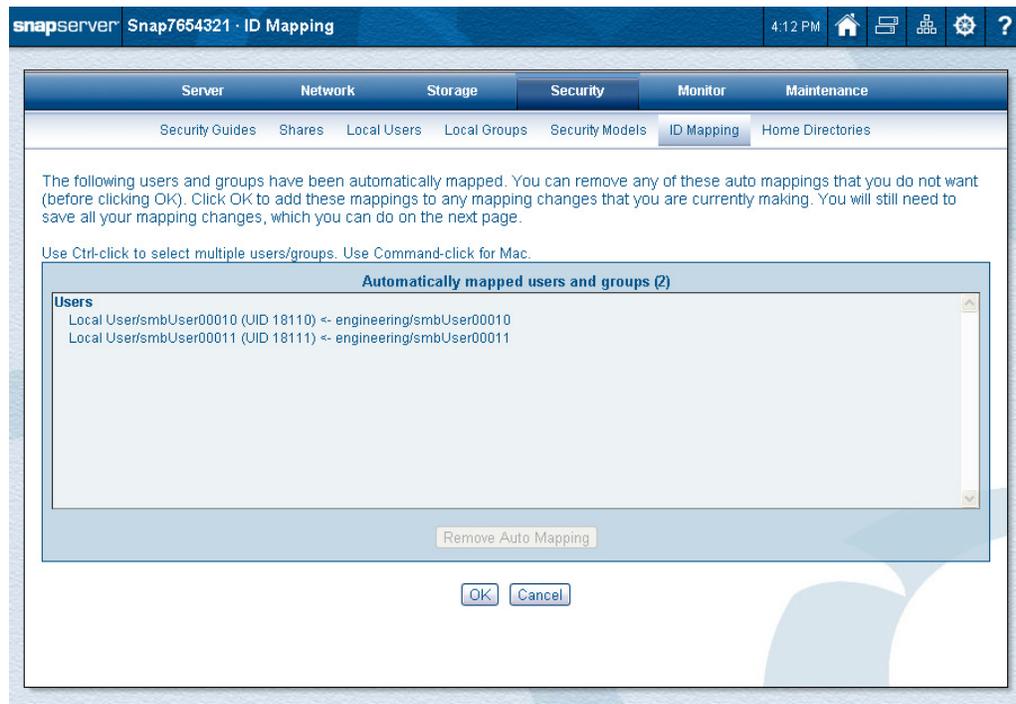
IMPORTANT: Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

See [Filesystem Updates on page 221](#) for more details.

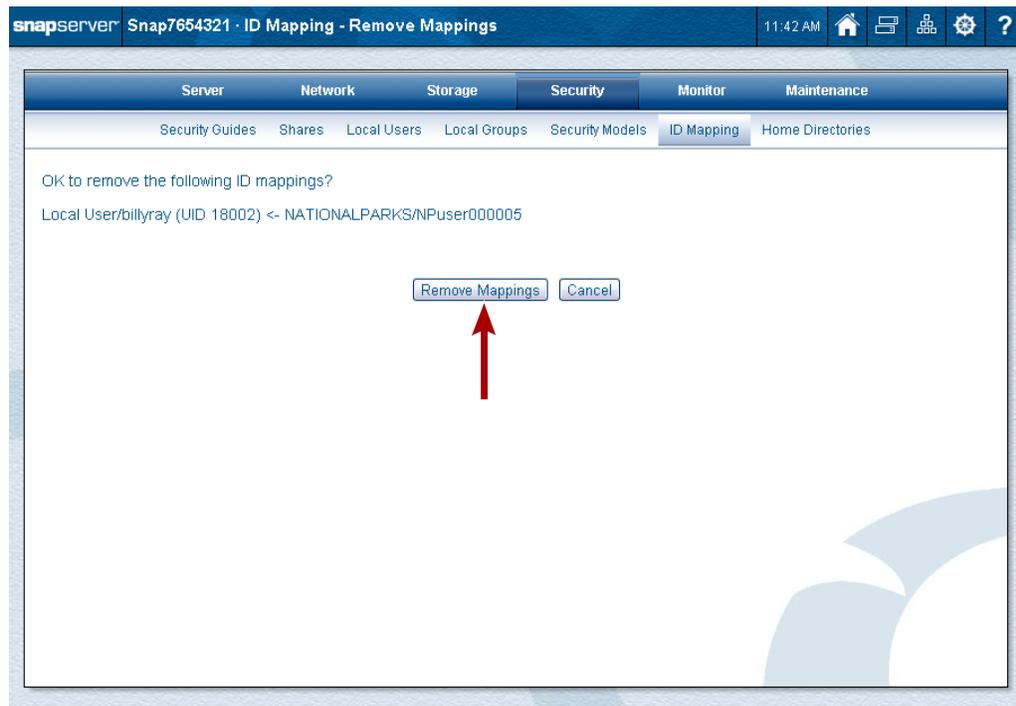
Auto Mapping

Auto mapping generates a list of ID mappings for Windows users and groups that have the same name as your local, LDAP, or NIS users and groups (local has precedence over LDAP and NIS).

1. Click **Auto Mapping** to generate a **list** of Windows domain users/groups that have the same name as your local, LDAP, or NIS users and groups.
Domain, local, LDAP, and NIS user/group lists are compared. The matches are automatically queued. Users and groups already mapped are not affected.
2. At the **Auto Mapping** confirmation page, click **View Auto Mappings** to display a page summarizing your changes.



- At the summary page, verify the **mappings** and remove (**Remove Auto Mapping**) any users or groups you do not want to map.



- When the list is correct, click **OK** to save changes (or **Cancel** to reset).
- When done with all your mappings, click **OK** to activate them.
- At the confirmation page, click **Save Changes**.
- At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

 **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.

See [Filesystem Updates on page 221](#) for more details.

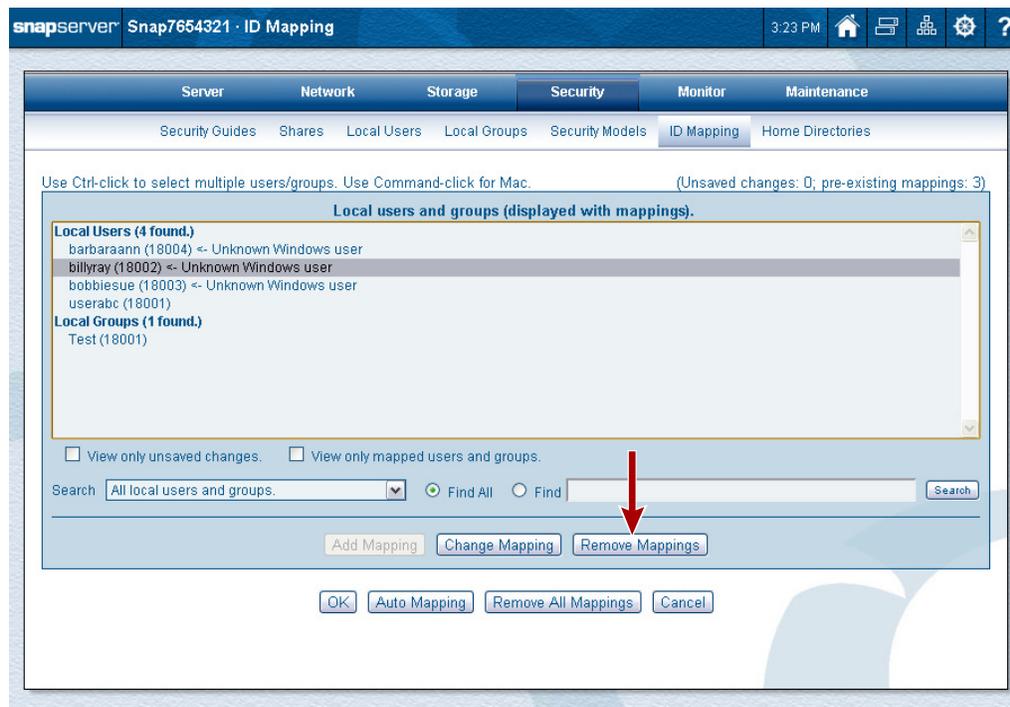
Remove Mappings

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 on page 209](#). You also have the option to update the filesystem after removing the ID mappings.

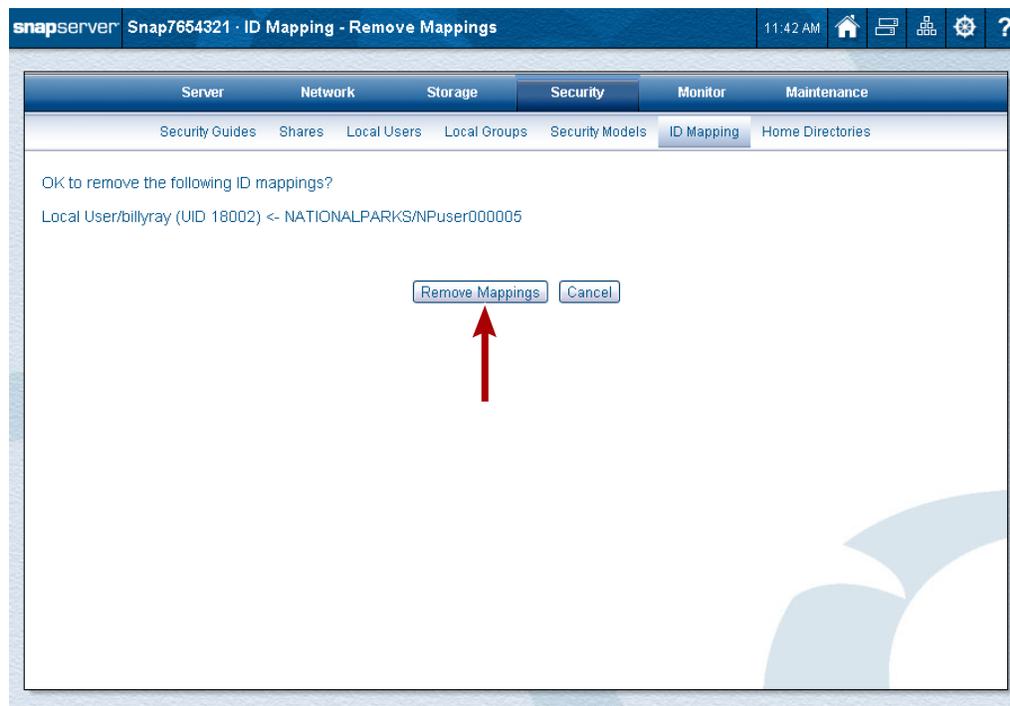
Remove Individual Mappings

- At the default **ID Mapping** page, select one or more **users/groups** you wish to unmap. To make it easier to find mappings for removal, check **View only mapped users and groups** to display only local, LDAP, or NIS users or groups that have been mapped.

2. Click **Remove Mappings**.



3. At the confirmation page, verify the **users/groups** listed and click **Remove Mappings**.



The selected mappings are removed and the default page is displayed with the users/groups that were unmapped in purple with a plus (+) in front of their name.

4. Click **OK** to save changes (or **Cancel** to reset).
5. When done with all your mappings, click **OK** to activate them.

- At the confirmation page, click **Save Changes**.
- At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

 **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.

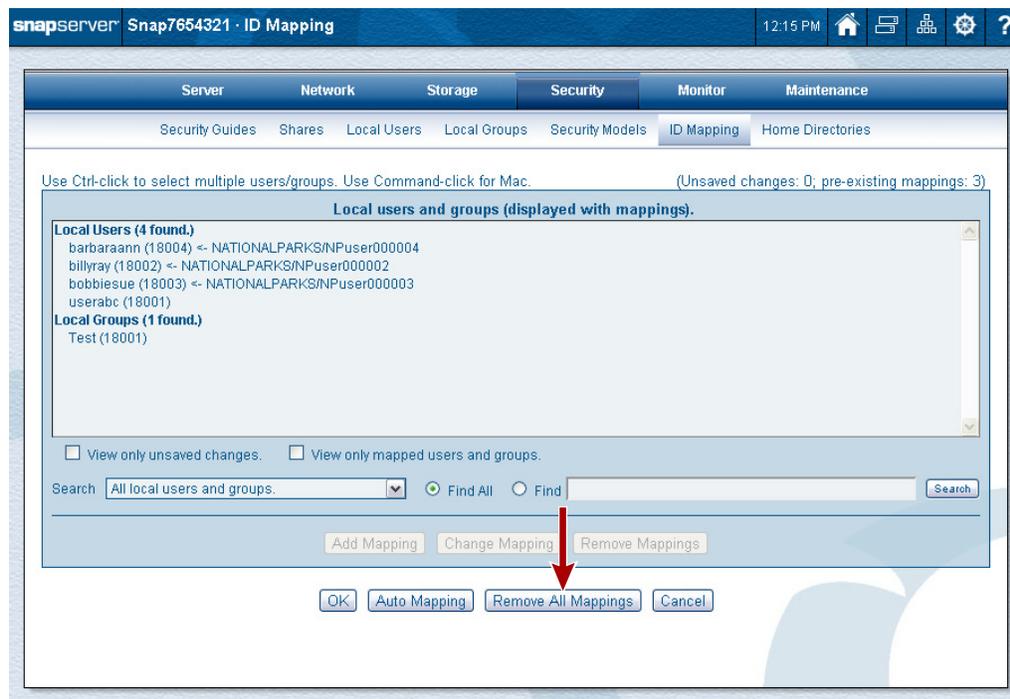
See [Filesystem Updates on page 221](#) for more details.

Remove All Mappings

The **Remove All Mappings** option allows you to remove **all** ID mappings on the server. If there are no mappings, the button is grayed out.

- At the default **ID Mapping** page, click **Remove All Mappings**.

If needed, 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, LDAP, or NIS users/groups that have been mapped to a Windows domain user or group.



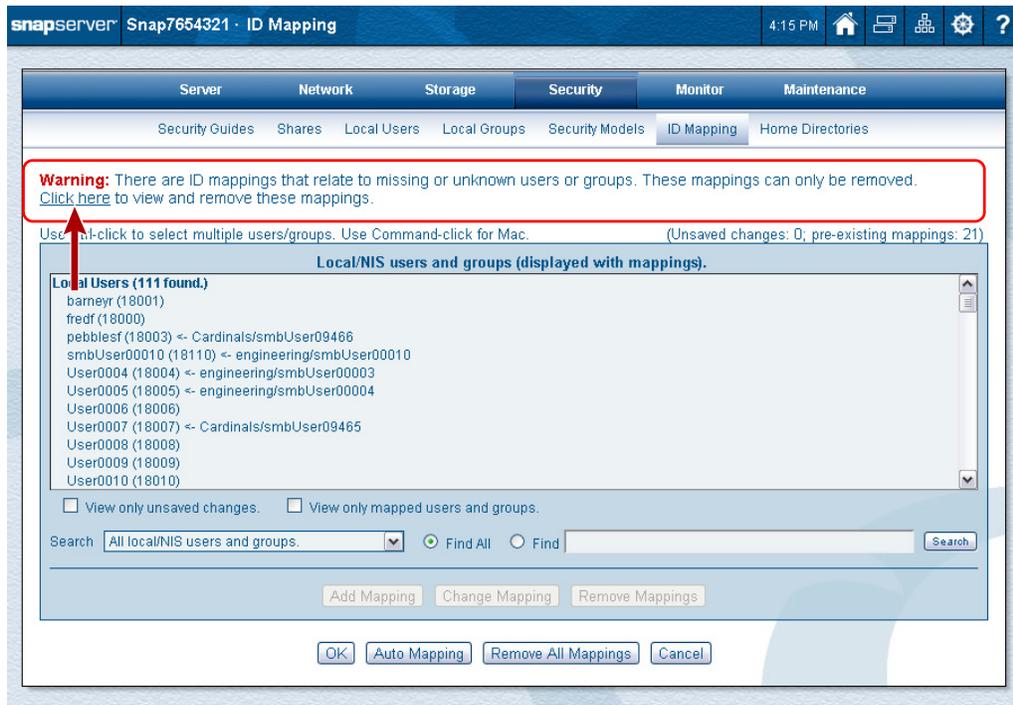
- At the confirmation page, click **Remove Mappings**.
All the mappings are removed and the default page is displayed with the users/groups that were unmapped in purple with a plus (+) in front of the names.
- Click **OK** to save changes (or **Cancel** to reset).
- When done with all your mappings, click **OK** to activate them.
- At the confirmation page, click **Save Changes**.
- At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

 **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.

See [Filesystem Updates on page 221](#) for more details.

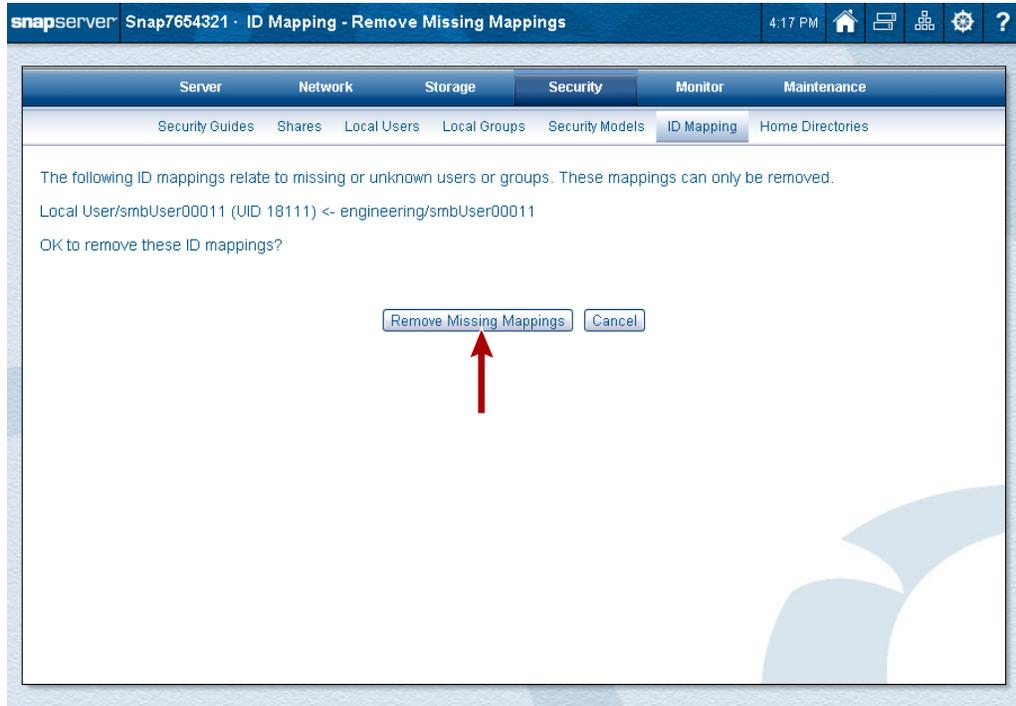
Remove Missing ID Mappings

If the server 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:

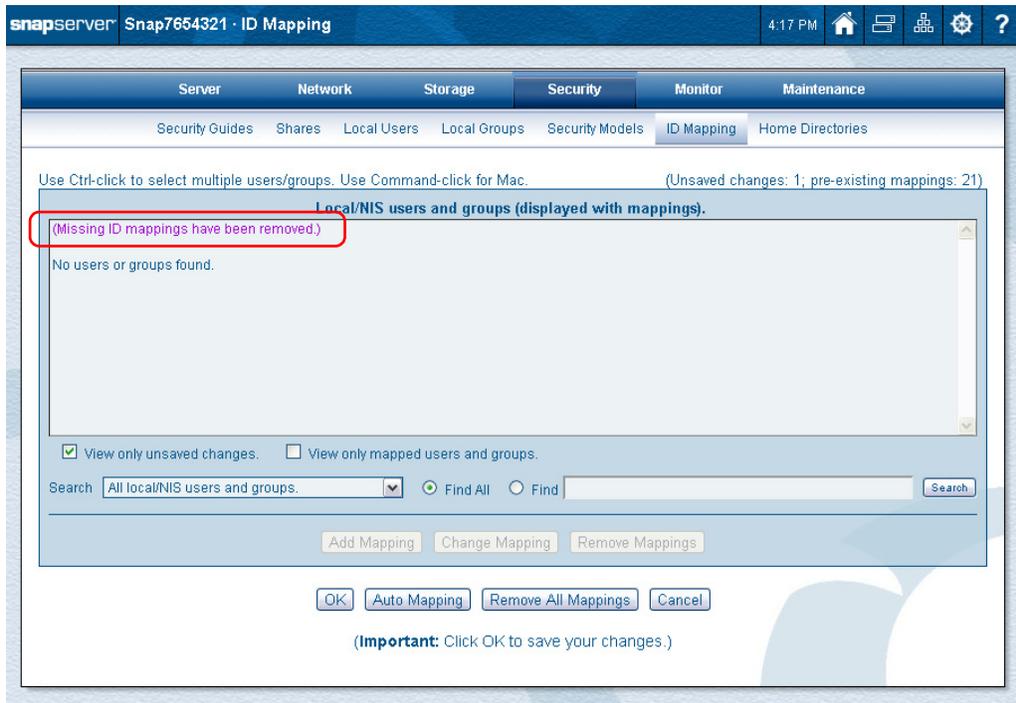


The screenshot shows the SnapServer web interface for the 'ID Mapping' page. At the top, there is a navigation bar with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. Below this is a sub-navigation bar with links for Security Guides, Shares, Local Users, Local Groups, Security Models, ID Mapping, and Home Directories. A red box highlights a warning message: "Warning: There are ID mappings that relate to missing or unknown users or groups. These mappings can only be removed. Click here to view and remove these mappings." Below the warning, there is a list of local users and groups, including barneyr (18001), fredf (18000), pebblef (18003), smbUser00010 (18110), User0004 (18004), User0005 (18005), User0006 (18006), User0007 (18007), User0008 (18008), User0009 (18009), and User0010 (18010). At the bottom of the page, there are buttons for Add Mapping, Change Mapping, Remove Mappings, OK, Auto Mapping, Remove All Mappings, and Cancel.

1. Click the **Click here** link in the warning message to display the **Remove Missing Mappings** page.



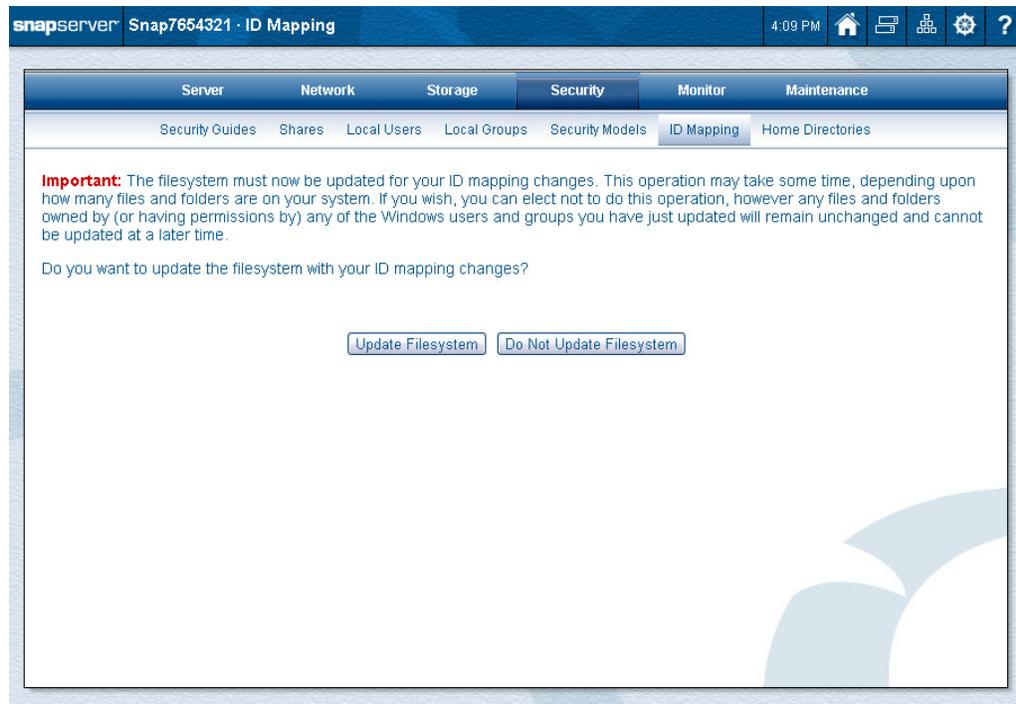
2. Click **Remove Missing Mappings** to clear the missing mappings from the system. A confirmation is shown on the **ID Mapping** main page.



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.

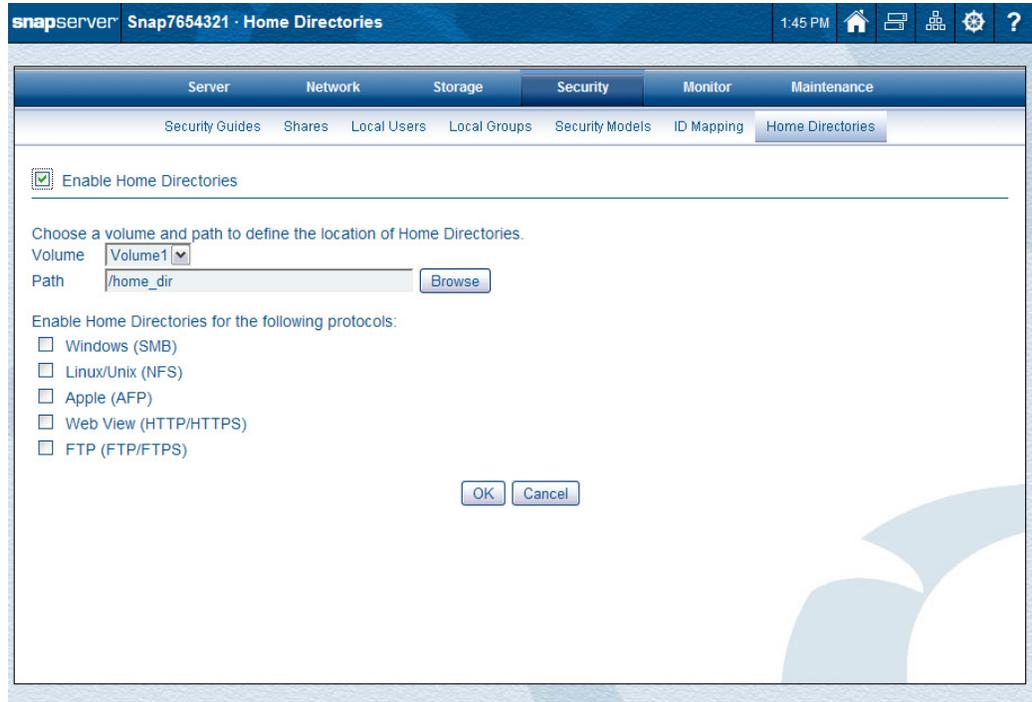


If you choose **Update Filesystem**, UID and GID ownership on files and SIDs in ACLs are updated to reflect the ID mapping operation.

 **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 **Security > Home Directories**), 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 are visible inside the root, but the user's home directory is 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 are directed to the domain user's home directory. In some cases, data in the local user's home directory is 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 are 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 are not copied to the domain user's home directory.

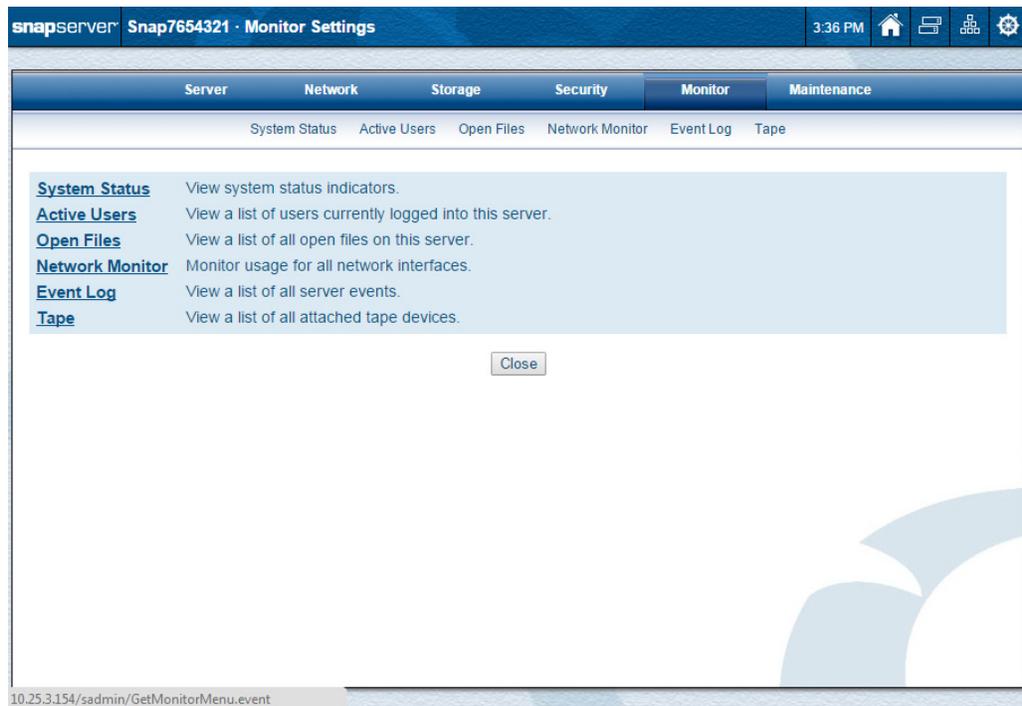
Configure Home Directories

Check or complete the following fields and click **OK**.

Field	Description
Enable Home Directories	Check to enable Home Directories for local users and activate the options. 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.



Topics in System Monitoring:

- [System Status](#)
- [Active Users](#)
- [Open Files](#)
- [Network Monitor](#)
- [Event Log](#)
- [Tape](#)

System Status

Use the **System Status** page (**Monitor > System Status**) to assess the hardware status and key information of the server.



The following status fields are displayed for the server (head unit). Any critical messages are displayed in a **red** font.

Field	Description
Server Name	Name of the server. The default server name is Snapnnnnnnn, where nnnnnnn is your server number (for example, Snap1123578).
Server Model	Server model name/number.
OS Version	The version of GuardianOS currently loaded on the SnapServer.
Server Number	Number derived from the MAC address of the Ethernet 1 port, used as part of the default server name.
Serial Number	Unique number assigned to the server.
Uptime	The amount of time the server has been up (since the last reboot) in “days:hours:minutes” format.
Memory	Amount of system RAM.
CPU	The type of central processing unit (CPU). If more than one CPU exists, each is listed separately.
Ethernet 1	Provides details on the server’s primary Ethernet connection.
Ethernet 2	If it exists, shows details on the server’s secondary Ethernet connection.
Ethernet n	If an optional Ethernet card is installed, shows details on the server’s other Ethernet connections are shown.

Field	Description
Ambient Temp.	The temperature of the space inside the chassis.
CPU Temp.	Current CPU temperature.
Power Supply	The status of power supply modules
Fan Status	The status of fan modules.

The following status fields will be displayed for each expansion unit.

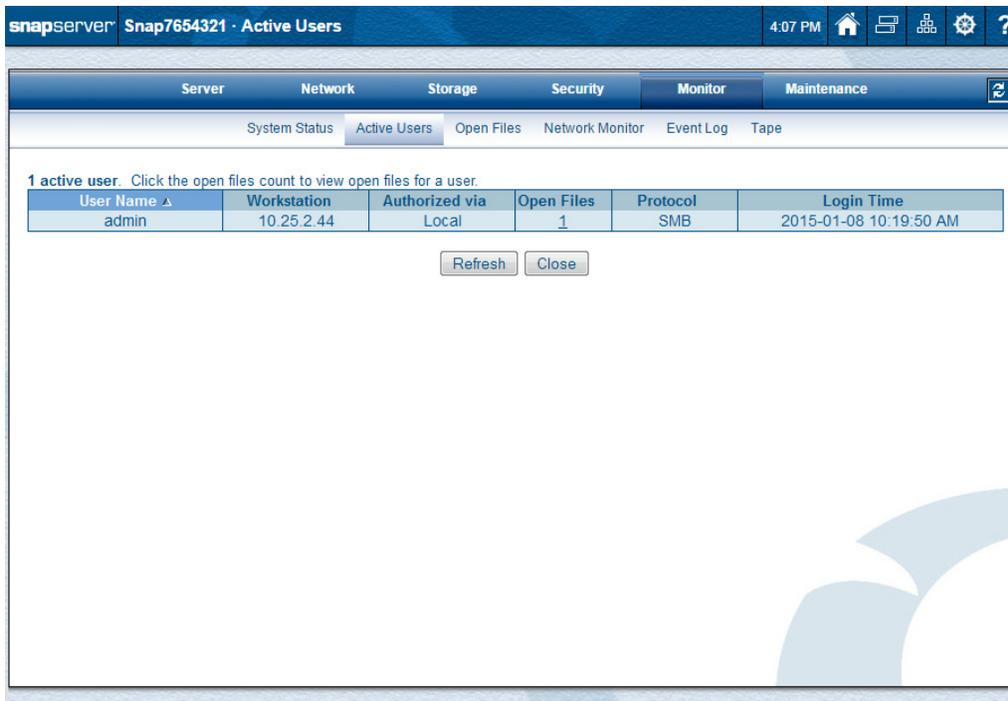
Field	Description
Expansion Unit	EXTN1, EXTN2, etc.
Expansion Model	SnapExpansion, etc.
Serial Number	The serial number of the expansion unit
Ambient Temperature	The temperature of the space around the expansion unit.
Power Supply	The status of the power supply
Fan Status	The status of fan modules.

Click **Refresh** to update the information. Click **Close** to return to the main **Monitor** page.

Active Users

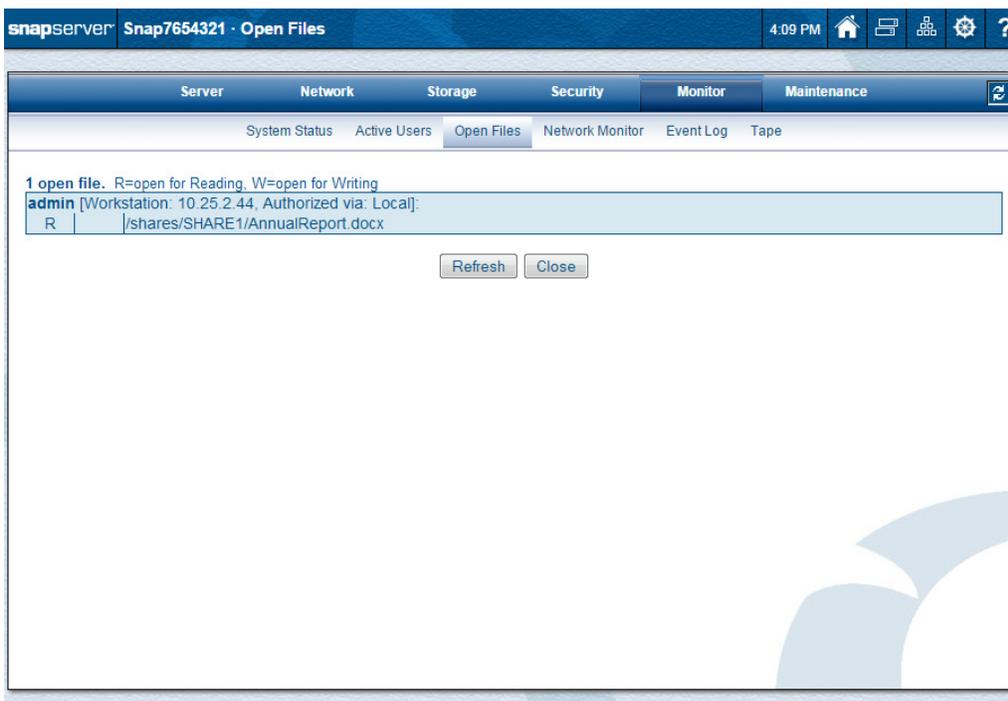
Use this page to view read-only details on the active users logged on to the server. Information available on this page includes user names of all active users, their workstation names, authorization, the number of open files they have on the share, the protocol, and when they logged on. Columns can be sorted in ascending or descending order by clicking the column head.

NOTE: Active users are not displayed for HTTP or NFS.



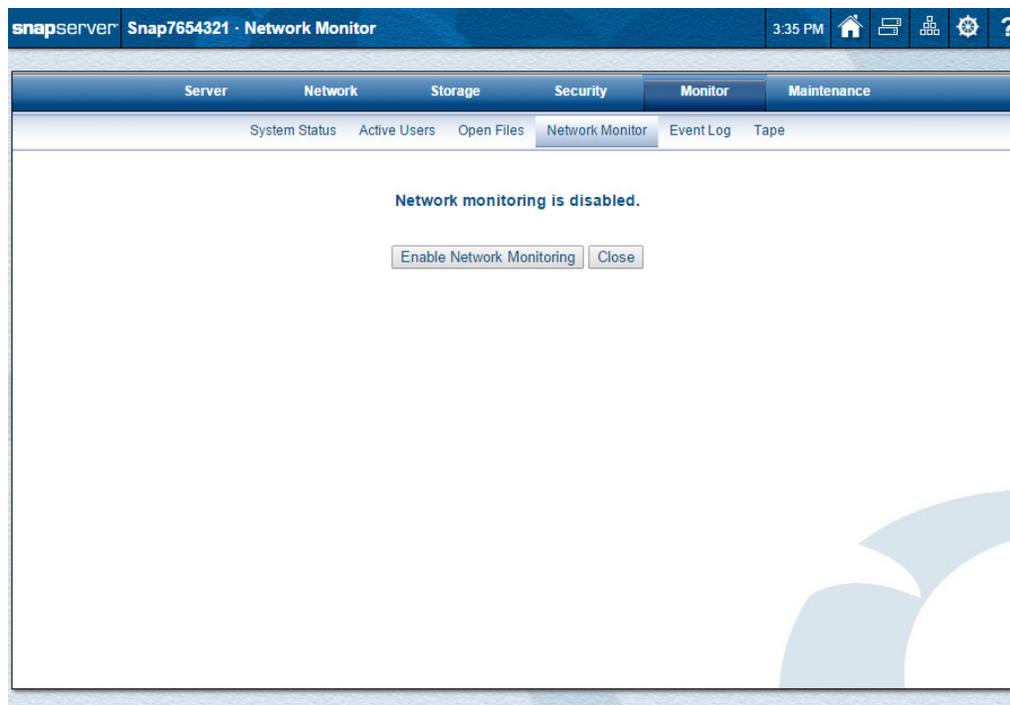
Open Files

Use this page to view read-only details on the open files in use on this server.



Network Monitor

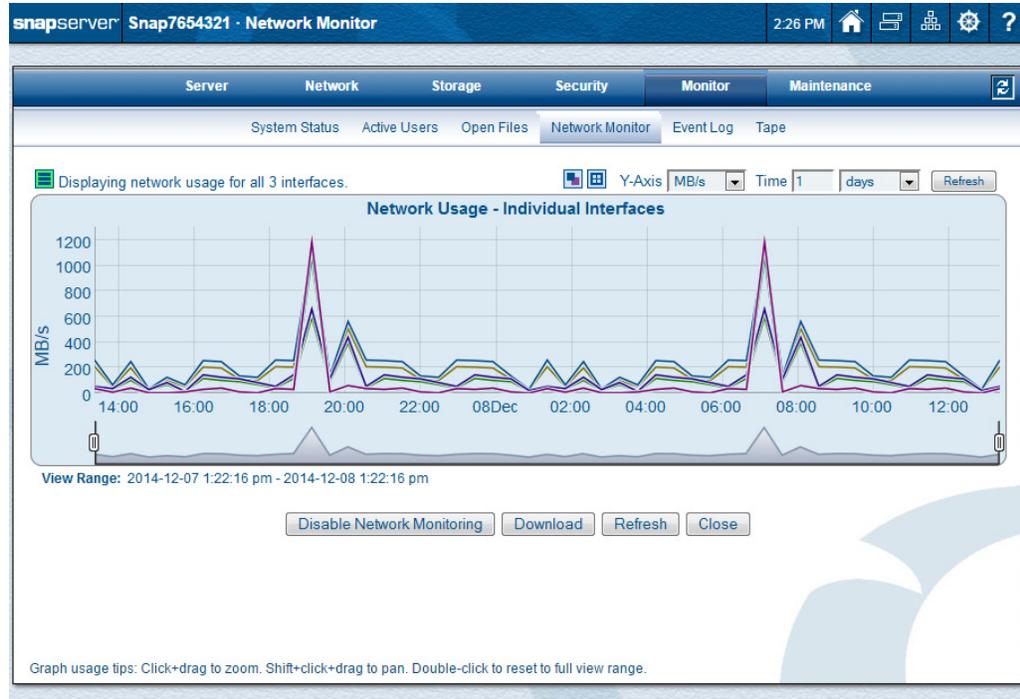
This feature can be used to monitor network utilization. Monitoring is disabled by default. Go to **Monitor > Network Monitor** and click **Enable Network Monitoring** to turn it on.



NOTE: When using Internet Explorer, due to Internet Explorer browser limitations, the Network Monitor feature only works when using Internet Explorer 9 or later. An error message is displayed for earlier versions of the browser.

View Usage

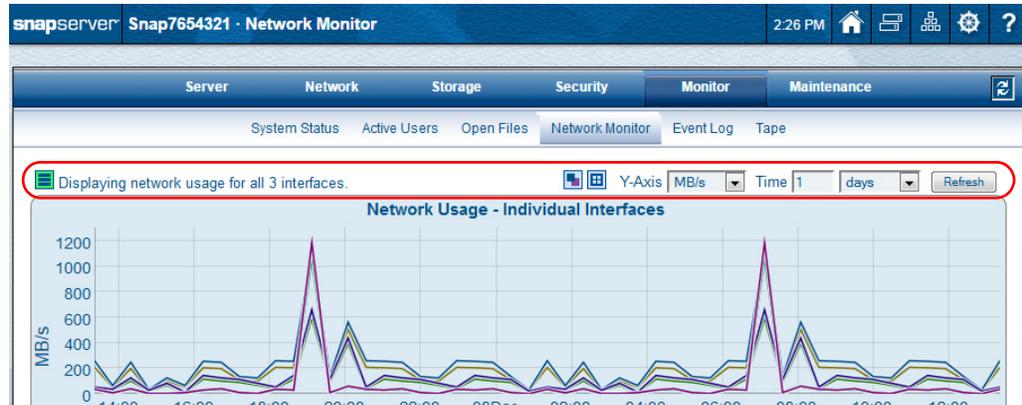
You can go to **Monitor > Network Monitor** to view graphs showing current usage, total throughput, or combined usage for a user-configurable time period. The data is refreshed every 15 seconds.



NOTE: You can manually refresh the data by using the **Refresh** button above the graph, the **Refresh** button at the bottom of the page, or the **Refresh** icon (🔄) at the right corner of the tab bar.

Icons and options are located above the graph to modify the view:

- The three icons are used to change the information presented in the graph:
 -  or  – Network interfaces (Selected Units or All) selected for graphing. Note that if only one interface exists, the All Interfaces icon is shown.
 -  or  – Throughput (Individual or Total) of input and output for the selected interfaces.
 -  or  – Network activity (Individual Interface or All Interfaces).
- Units of measure for the Y-axis are selectable (**Percent, MB/s, or GB/s**).
- Time interval of the data can be adjusted to be shown as a period of time (1 to 999 **minutes, hours, or days**).



Icons / Options	Description
<p>Select Interfaces</p>	<p>Click this icon to select the network interface. Choose either individual interfaces or Select All Interfaces. You can select multiple interfaces by using Ctrl+Click.</p> <ul style="list-style-type: none"> • When blue, individual interfaces are selected. • When green, all interfaces are selected. <p>Interface selection is not available when Combined Usage () is enabled. To display the Select Interfaces icon, turn Combined Usage off.</p>
<p>Total Throughput</p>	<p>This icon controls whether the graph represents individual (separate input and output) throughput or total combined (input plus output) throughput.</p> <ul style="list-style-type: none"> • When blue, the option is not active and the numbers reflect the individual input and output usage. • When green, the option is active and the numbers reflect the total throughput (combined input and output). <p>Click the icon to enable the option (green) or disable it (blue).</p>
<p>Combined Usage</p>	<p>This icon controls whether the graphs represent network usage for individual interfaces or combined network usage of all interfaces:</p> <ul style="list-style-type: none"> • When blue, the option is not active and the numbers shown reflect network activity for individual interfaces. • When green, the option is active and the numbers shown reflect network activity for combined interfaces and the Select Interfaces icon disappears. <p>Click the icon to enable the option (green) or disable it (blue).</p>
<p>Y-Axis Display Options</p>	<p>Use this drop-down menu to set the unit of measurement of the Y-Axis to be either Percent of network usage, MB/s (megabytes per second), or GB/s (gigabytes per second).</p>
<p>Time Options</p>	<p>Controls the overall time range represented in the graph. Enter a value from 1 to 999 and use the drop-down menu to select the time interval of minutes, hours, or days.</p>

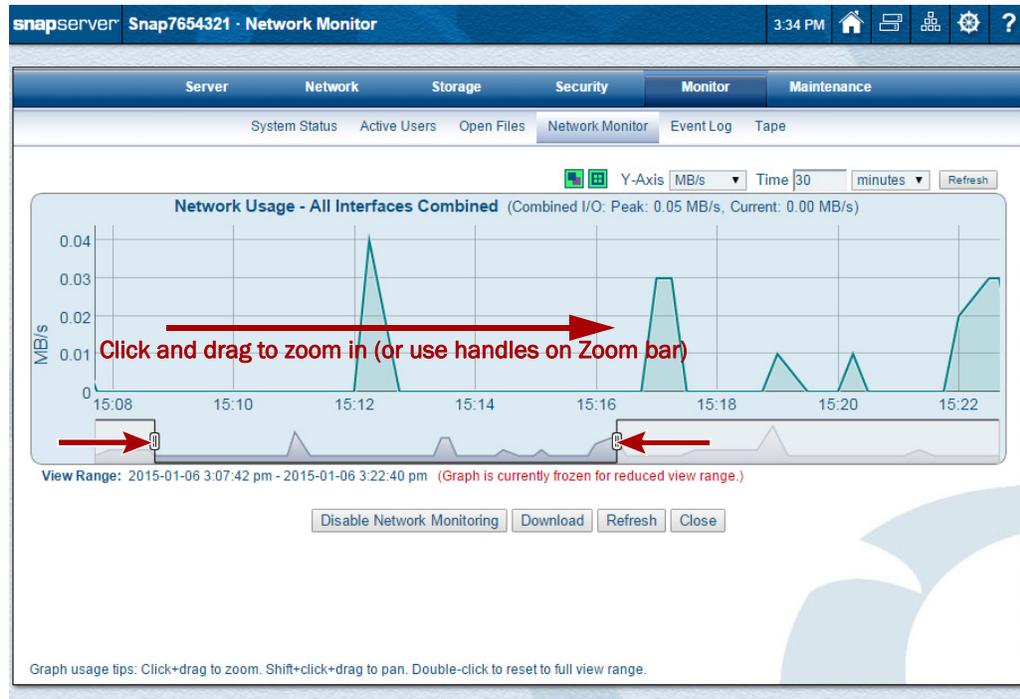
When you mouseover the usage bars, tool tip messages are displayed under the graph titles for that bar. Depending on the type of bar (Combined, Input, or Output), the message shows information about the interface.

The buttons at the bottom of the page let you disable network monitoring, download the network usage logs, manually refresh the data, or close the page.

Graph Options

Below each graph is a gray Zoom Bar that can be used to show a specific time range. When zoomed in, the graphs are frozen and not updated.

- You can scale the magnification of the graph by either clicking and dragging horizontally within the graph area or using the handles at the sides of the Zoom Bar.
- To pan and view any time period within the specified overall time range in more detail, Shift-click and drag the graph or click and drag the ends of the Zoom bar horizontally.
- To reset the zoom level and restore automatic updates, double-click within the graph area.



Download Usage Records

To download the record displayed as a CSV file, click **Download**. Depending on your browser, the file is saved or a dialog box asks you to determine the location of the downloaded file.

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:

Color	Icon	Entry Type
Red	E	Error (E)
Yellow	W	Warning (W)
(no color)	I	Informational or Unclassified (I)

Filter the Log

Edit the following fields as appropriate, then click **Refresh**.

Option	Description
Severity	Select the type of alerts and information you want to view.
Display Last <i>n</i> Days	Enter the number of days' worth of entries you want to view.
Most Recent First	Check this box to start the list with the most recent entry; deselect to start the list with the oldest entry.

Tape

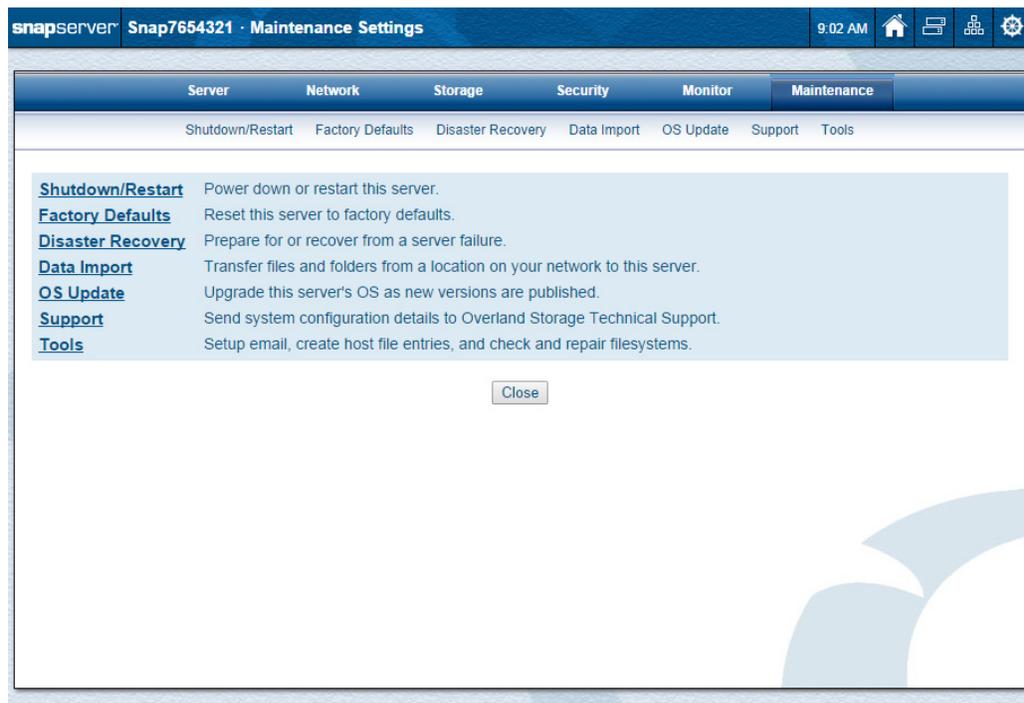
Use the **Tape Monitor** page (**Monitor > Tape**) to view read-only details on the SCSI and USB tape devices attached.

Device Model	Device Type	Device name	CX	Bus	ID	LUN
HP MSL G3 Series	Medium Changer	/dev/sg33	SCSI	7	12	1
HP Ultrium 5-SCSI	Sequential-Access	/dev/sg32	SCSI	7	12	0

The following table describes the fields:

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 server 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 options used to maintain this SnapServer. There is also a **Tools** submenu of special, related options.

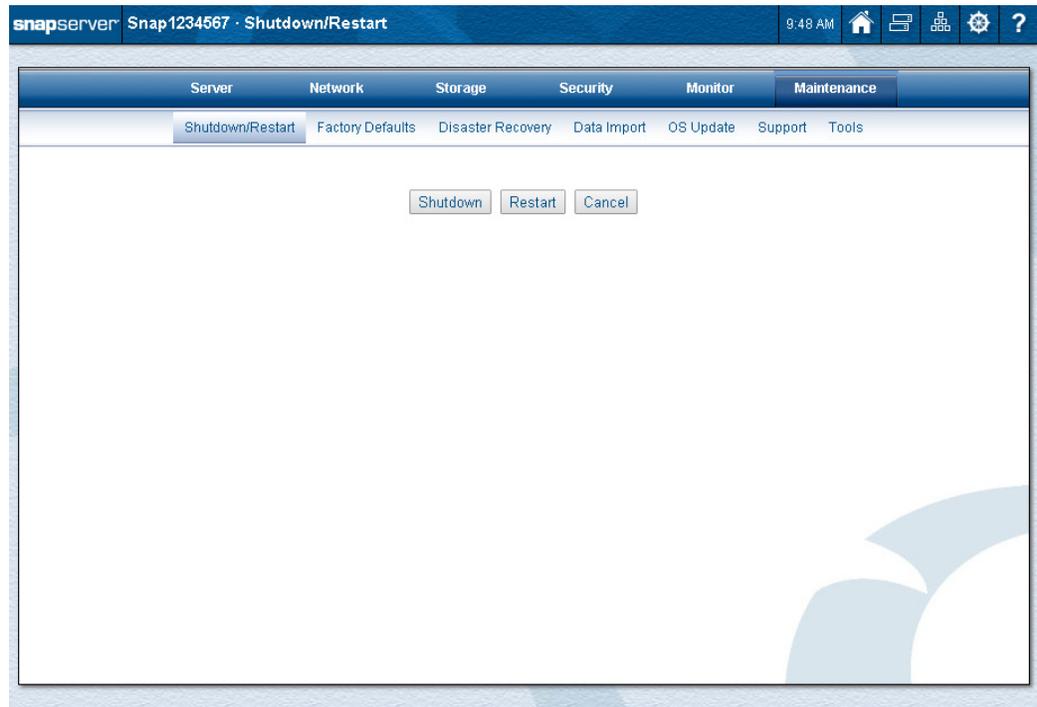


Topics in Maintenance

- [Shutdown and Restart](#)
- [Factory Defaults](#)
- [Disaster Recovery](#)
- [Data Import](#)
- [OS Update](#)
- [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 SnapServer 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.

The Power button on the front of the server can be used to power on or power off (in an emergency) a 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.

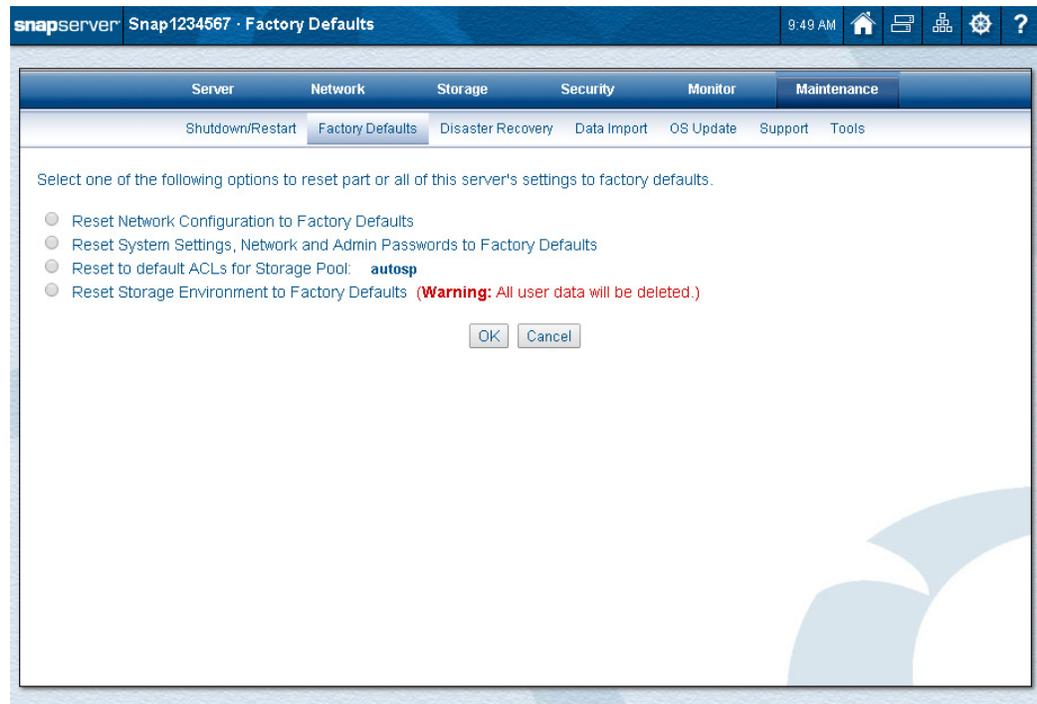
NOTE: SnapServers have a persistent power state. When a physical loss of power occurs, the server returns to the same state it was in when the power went out. Therefore, if the server is powered down prior to a power loss, it will remain powered down when the power is restored, and if it was powered up prior to a power loss, it will power back on when 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/Unix 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.

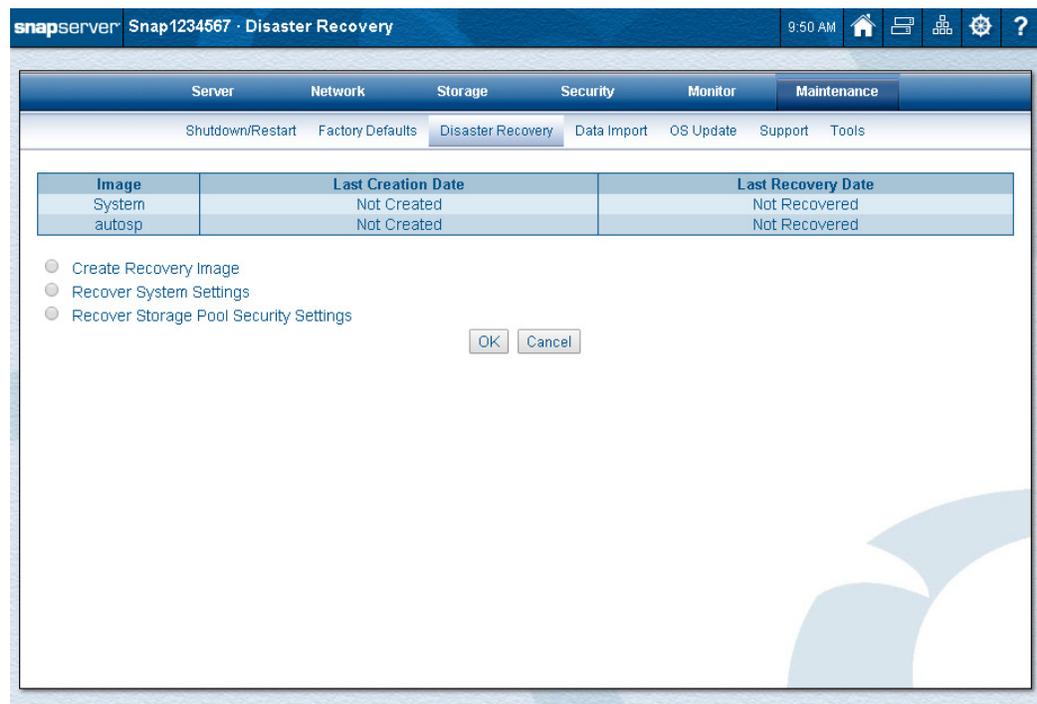
 **CAUTION:** ALL USER DATA WILL BE DELETED on the SnapServer when this option is selected. A confirmation page will be displayed and the admin password must be entered before the process begins.

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 configuration, such as network and RAID configurations, together with 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 \(Appendix A\)](#) for information on backup options.



Disaster recovery also encompasses 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 these procedures:

- Reinstalling the SnapServer operating system (GuardianOS).
- Restoring the SnapServer 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.

More on Disaster Recovery:

- [Backing Up Server and Volume Settings](#)
- [SnapDRImage File and Volume-Specific Files](#)
- [System Settings Recovery](#)
- [Volume and Storage Pool Security Settings Recovery](#)
- [Replacing or Cloning a 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 [Initial Setup and Configuration on page 15](#).

SnapDRImage File and Volume-Specific 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, 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 is 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 RAIDs; (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. Click **Create Recovery Image** 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 **Download Recovery Image** to download the SnapDRImage file to a safe location on another server or backup medium. (See [SnapDRImage File and Volume-Specific Files on page 239](#) 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: Back up volume-specific files with scheduled data backups.

Ensure the `.os_private` directory on each volume is included in your backup configuration so the volume-specific files are written to tape as part of your regular volume backup procedures.

System Settings Recovery

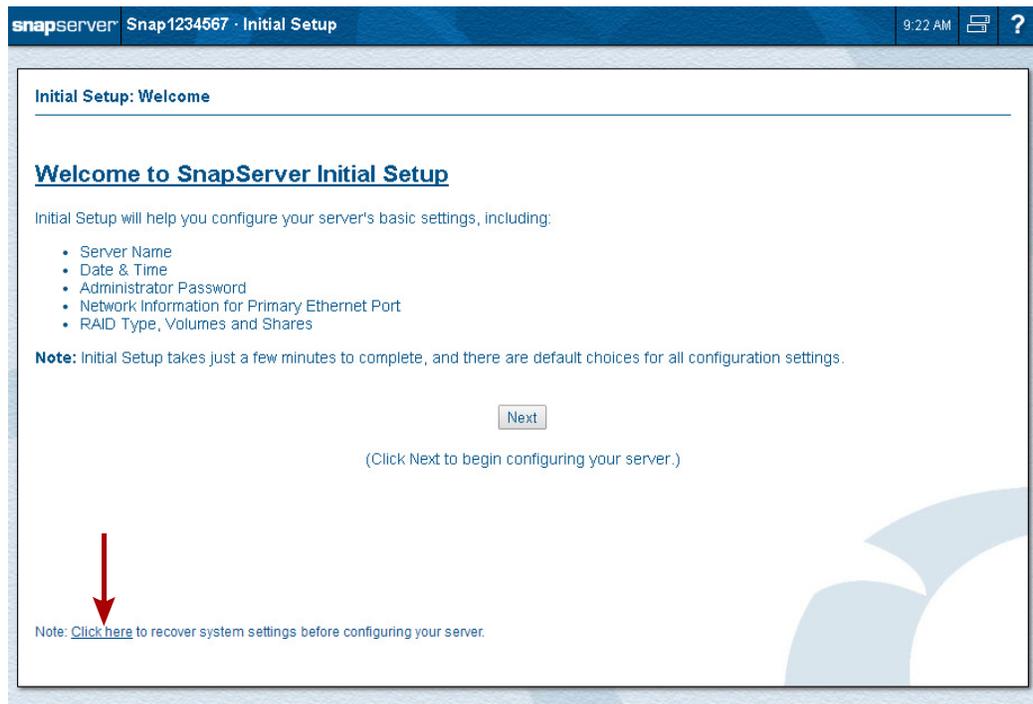
System settings include all network, user, share, and storage configuration, and can only be performed on an uninitialized server being used as a replacement or clone. If system settings must be restored to a configured server, contact technical support to perform a fresh install of the OS to put it back in the uninitialized default state.



CAUTION: A fresh install of the OS and return to the uninitialized default state will destroy all existing data on the server.

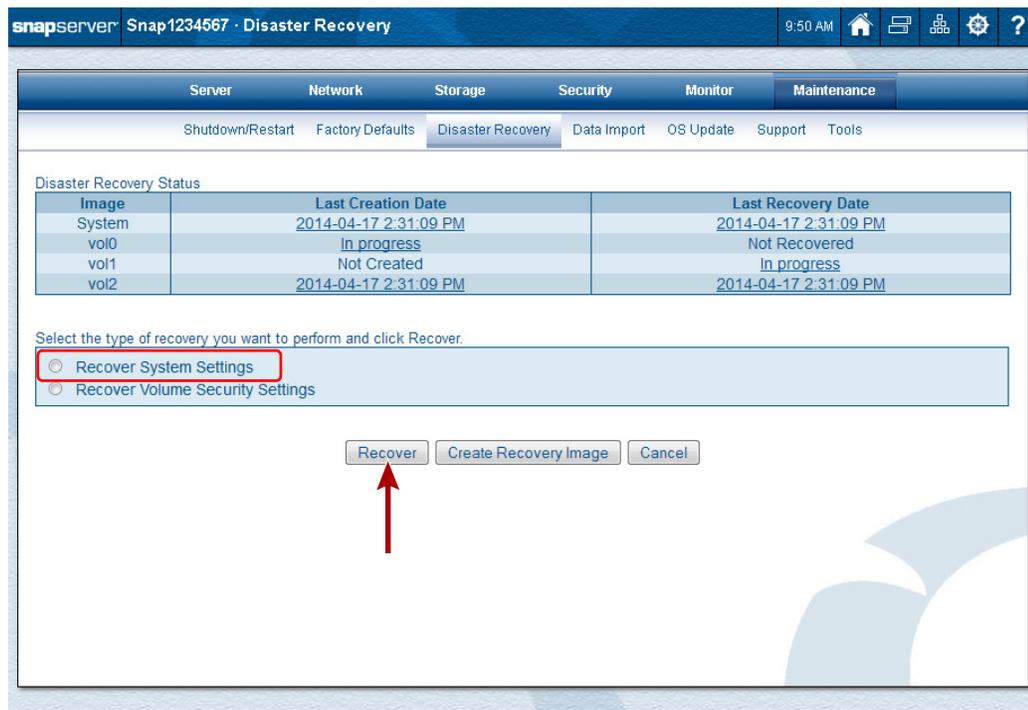
Recovering System Settings

1. Click the **link** on the **Initial Setup Welcome** page:



This link launches the **Disaster Recovery** page.

- At the **Disaster Recovery** page, select the **Recover System Settings** option and click **Recover** to open the Server Recovery 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 is not interrupted which might result in a loss of data.

- At the **Server Recovery** page, use **Browse** 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**. After recovery completes, the server restarts.
- After the server restarts, log in to **Administration**, navigate to **Network > Windows/SMB**, and, if necessary, rejoin the Windows domain.

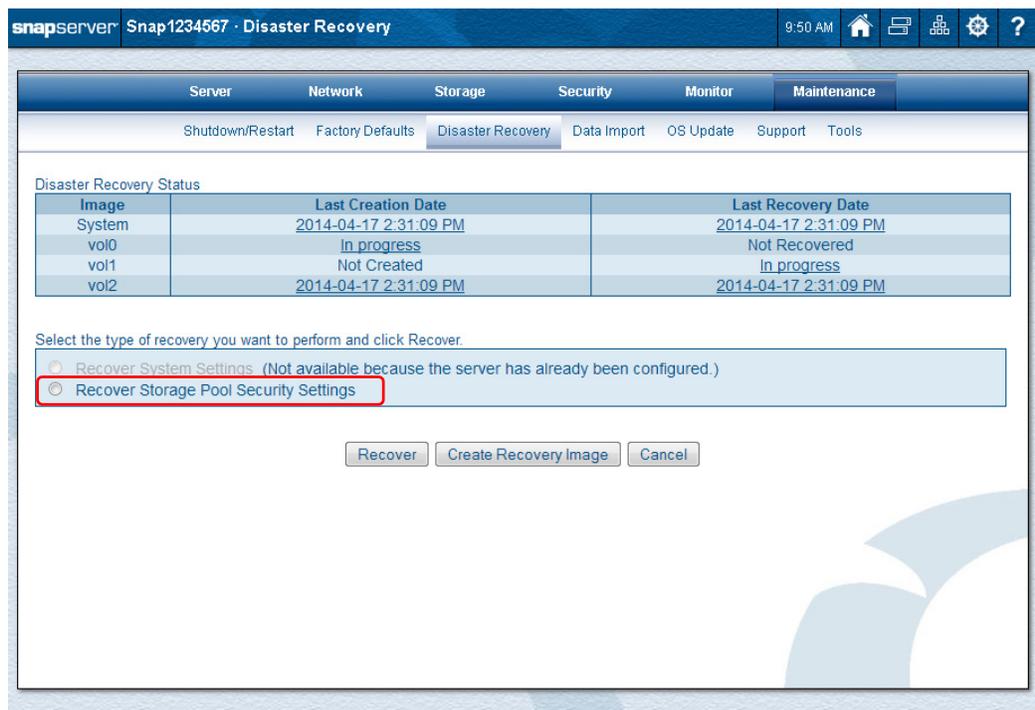
The server is now available for additional configuration, data restore from tape backup, and volume or storage pool security recovery. If any configuration operations failed, view the recovery log on the **Disaster Recovery** page.

Volume and Storage Pool Security Settings Recovery

Volume and storage pool security settings include file system security and quota settings for each volume, and is only available on configured servers with volume, data, and user configuration in place. Recovery requires the backup.acl, backup.qta.groups, and backup.qta.users files to be present in the .os_private directory on each volume or storage pool on which you wish to recover security settings.

Recovering Security Settings on a Volume or Storage Pool

- If necessary, restore **user data** from tape backups to **each** volume or storage pool you want to recover.
 - Traditional RAID** – Ensure the correct backup.acl, backup.qta.groups, and backup.qta.users files matching the data are placed in the .os_private directory on each of the volumes.
 - DynamicRAID** – The files need to go into .os_private on the **first** volume of the storage pool. When you create a Disaster Recovery Image (DRI), the files for the whole storage pool go in the first volume; and when you restore, you need to restore the files back to that first volume's .os_private.
- Connect to the **Administration** page and navigate to **Maintenance > Disaster Recovery**.



- Select the **security settings** option and click **Recover**.
DynamicRAID shows **Recover Storage Pool Security Settings** while Traditional RAID shows **Recover Volume Security Settings**. For the remaining steps, references to *storage pools* are used for DynamicRAID systems and *volumes* for Traditional RAID systems.
- Select the **storage pool/volume** you want to restore.
The creation date of the recovery file on a storage pool/volume indicates when the recovery files were generated. Storage pools/volumes that do not have recovery files in .os_private appear as unavailable.
- Click **Recover** to start the operation and follow the onscreen instructions.
- After recovery completes, check the **recovery log** for the storage pool/volume on the **Disaster Recovery** page if there were any errors.
- Repeat **Steps 3–6** to recover any **additional** storage pools/volumes.

Replacing or Cloning a Server

Disaster Recovery combined with restore from backup can be used to recover configuration to a replacement server or to clone a server's configuration to another server.



IMPORTANT: When recovering configuration to a server replacing a failed server, Overland Storage strongly recommends that you contact a technical service representative before proceeding.

When recovering configuration, any third-party license keys you have not purchased through Overland Storage are 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. You will also need to reschedule snapshots and reconfigure CA Antivirus.

Replacing or Cloning a Server

1. Recover **system settings** as described in [Recovering System Settings on page 241](#).
2. Recover **volume or storage pool security settings** as described in [Recovering Security Settings on a Volume or Storage Pool on page 243](#).
Be sure to restore the backup.acl, backup.qta.groups, and backup.qta.users files matching the user data to the .os_private directory on each volume you want to restore.
3. If necessary, reconfigure or reschedule the following items:
 - Reconfigure your **Snap EDR** settings.
 - Reconfigure your **BitTorrent Sync** settings.
 - Reconfigure your **CA Antivirus** settings.
 - Reschedule your **snapshot** times.

Data Import

Use the **Data Import** page (**Maintenance > Data Import**) to import (migrate) data from another SnapScale cluster, SnapServer, or other computer that supports CIFS or NFS (v2, v3, or v4) to this server.

Windows/SMB Page:

The screenshot shows the SnapServer Data Import interface for the Windows/SMB protocol. The page title is "Snap1234567 · Data Import" and the time is 12:56 PM. The navigation menu includes Server, Network, Storage, Security, Monitor, and Maintenance. The sub-menu includes Shutdown/Restart, Factory Defaults, Disaster Recovery, Data Import, OS Update, Support, and Tools. The main content area is titled "Use Data Import to copy or move files and folders from a location on your network (Source) to this server (Target)." The "Source:" section has a "Network Protocol" dropdown set to "Windows (SMB)" (highlighted with a red box), with a note "(specifies how to communicate with host)". Below this are input fields for "Auth. Name", "Auth. Password", "Host", "Share", and "Path", each with a "Browse" button. The "Target (This SnapServer):" section has a "Volume" dropdown set to "VOL0" and a "Path" input field with a "Browse" button. The "Options:" section includes an "Import Type" dropdown set to "Copy (source data is maintained)", and three checkboxes: "Include all sub-folders (if source path specifies a folder)" (checked), "Overwrite existing target files and folders (that have identical names as the source files and folders)" (checked), and "Preserve file/folder permissions" (unchecked). There is also an unchecked checkbox for "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." At the bottom right are buttons for "Start Import", "View Log", and "Close".

NFS Page:

The screenshot shows the SnapServer Data Import interface for the NFS protocol. The page title is "Snap7654321 · Data Import" and the time is 11:27 AM. The navigation menu includes Server, Network, Storage, Security, Monitor, and Maintenance. The sub-menu includes Shutdown/Restart, Factory Defaults, Disaster Recovery, Data Import, OS Update, Support, and Tools. The main content area is titled "Use Data Import to copy or move files and folders from a location on your network (Source) to this server (Target)." The "Source:" section has a "Network Protocol" dropdown set to "NFS" (highlighted with a red box), with a note "(specifies how to communicate with host)". Below this are input fields for "User Name" (with a note "(Snap local, NIS or LDAP user)"), "Host", "Export", and "Path", each with a "Browse" button. The "Target (This SnapServer):" section has a "Volume" dropdown set to "Volume1" and a "Path" input field with a "Browse" button. The "Options:" section includes an "Import Type" dropdown set to "Copy (source data is maintained)", and three checkboxes: "Include all sub-folders (if source path specifies a folder)" (checked), "Overwrite existing target files and folders (that have identical names as the source files and folders)" (checked), and "Preserve file/folder permissions" (unchecked). There is also an unchecked checkbox for "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." At the bottom right are buttons for "Start Import", "View Log", and "Close".

If an error is encountered during the import (for example, a file or folder is locked and cannot be imported), the 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 main page and recreates the import. With the exception of the password, all fields will still be populated with the specifications of the last import 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, review [Preserving Permissions on page 249](#).

- Verify imported data

NOTE: If you elect to verify imported data, all data is 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 has 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 NFS v2/v3 import, consider using the user root and configuring the NFS export on the source to `no_root_squash` for the IP Address of the server for the duration of the import.

NOTE: Only one import job can run at a time.

To create a data import job, perform the following procedure:

1. On the **Data Import** page, complete the required **information** for both the source and target.

Option	Description
<i>Source:</i>	
Network Protocol	<p>Protocol that the server uses to connect to the source server. Use the drop-down list to select:</p> <p>NOTE: If you are importing via SMB, SMB must also be enabled on the target server (enable at Network > Windows/SMB).</p> <ul style="list-style-type: none"> • Windows (SMB) – Uses SMB for Windows with the source data on a Windows root directory (default option). • NFS – Uses NFS v2/v3 for Unix/Linux-based servers or a GuardianOS server with source data on a Unix root directory.

Option	Description
Auth. Name & Auth. Password / User Name	<ul style="list-style-type: none"> For the Windows (SMB) network protocol, enter both the Auth. Name and Auth. Password (Windows user name and password to log in to the source server over SMB). For the NFS network protocol, enter the User Name (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 server you are importing data from.
Share/Export	Specify the share (Windows) or export (NFS) on the source server containing the data you want to import. NOTE: Wildcards are not supported when specifying the source share to import.
Path	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. NOTE: Wildcards are not supported when specifying the path to import.
<i>Target (This SnapServer):</i>	
Volume	Specify the volume where you want the data imported.
Path	Specify the path to the directory 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 the Verify Imported Data option.
Include All Sub-folders	If the folder you select for import contains subfolders, selecting this option imports all files and folders underneath this folder (default is checked). NOTE: If disabled, <i>only</i> the files in this folder are imported.
Overwrite Existing Target Files & Folders	If any files/folders on the target have identical names with files/folders on the source, checking this option overwrites those files/folders during import (default is checked).
Preserve File/Folder Permissions	Selecting this option retains the source permissions when the files/folders are imported to the target (unchecked by default). NOTE: Before selecting this option, review Preserving Permissions on page 249 .

Option	Description
Verify Imported Data	<p>Selecting this option causes all source data to be read twice, once to write to the target and once to perform a binary comparison with the data written (default is unchecked).</p> <p>If enabled and, if the Import Type is Move, files on the source are only removed after verification. Otherwise, files are removed during the process of copying them to the target. If you select to move the files rather than copy them, it is strongly recommended that you enable the Verify Imported Data option.</p> <p>If a file mismatch occurs during verification, the target file is moved to a <code>data_import_verify_failures</code> 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 do not re-copy files that have already been copied and verified).</p> <p>NOTE: Depending upon how much data is being imported, verifying imported data can be a lengthy process.</p>
Email Notification	<p>Clicking the email notification link takes you to the Email Notification page (for more information, see Email Notification on page 258).</p> <p>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.</p>

- Once you have completed the import information, click **Start Import** to begin the import. You can see the progress of the import, the estimated time until completion, and the import log on the secondary **Data Import** page.
- When the import is complete, click **View Log** 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 **Stop Import** on the **Data Import** secondary page. If a file was in the process of being copied, the partially-copied file on the target is removed.

Recreating an Import Job

The **Data Import** log records all errors that occurred during import. You can import just the files and folders that were not imported during the original job due to an error condition (for example, the file was locked).

- Review the **Data Import errors log** and correct all error conditions (such as unlocking a locked file).
- Reopen the **Data Import** page. All fields (except the password) from the last import will still be visible on the page.

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 all problematic files from the first import are deleted from the target so they can be re-imported.

NOTE: If an import failed, it is strongly recommended that you enable the **Verify imported data** option for the re-importation.

- Enter your password and click **Start Import** to run the import again.

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 Personality Directory

If you are importing from a Windows server (or other type of server that follows the Windows security model) to a Windows personality directory, permissions are retained exactly as they exist on the source. However, as is the case when moving files with permissions between Windows servers, permissions for users who are unknown on the target are retained but not enforced. This includes permissions for:

- Local users on the source machine.
- Domain users for domains unknown to the server (for example, trusted domains, if the server is not configured to support trusted domains).
- Certain built-in Windows users and groups.

Importing from a Unix Security Model to a Unix Personality directory

If you are importing from a Unix server to a Unix personality directory, Unix permissions for UIDs/GIDs are copied exactly from source to target; thus, identities of the users and groups are 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 security model target, Unix permissions are retained and the security personality on the resulting files and directories will be Unix.

However, when importing from a Windows source to a Unix security model target, 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 SnapServer or SnapScale Cluster

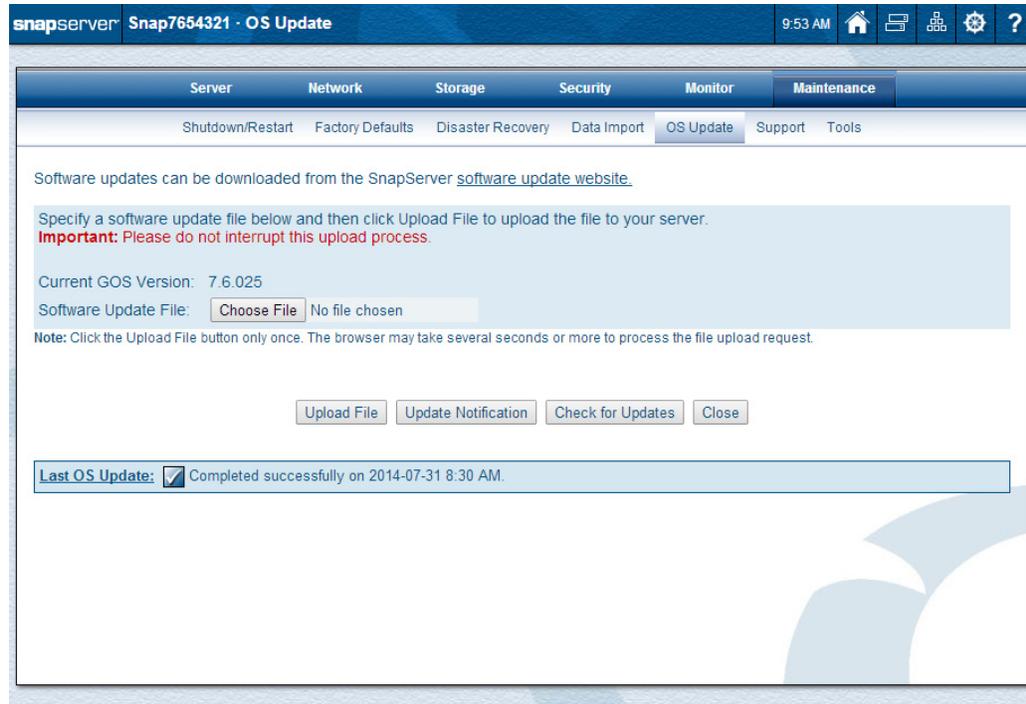
When importing from a different SnapServer or SnapScale cluster, it is recommended that you maintain the same security model on the target that you have on the source.

- If your source uses a Windows security model and has permissions assigned to Windows domain users, use a Windows (SMB) connection for import. Windows permissions are retained exactly as they are on the source, with the same enforcement limitations for unknown users as for importing from Windows servers (see [Importing from a Windows Security Model to a Windows Personality Directory on page 249](#)).
- If your source server or cluster uses a Unix security model and has permissions assigned to local or NIS users, use an NFS connection for import.

NOTE: Local users who have Unix permissions on the source are not created on the target with the same UIDs.

OS Update

Use this page to install updates to GuardianOS and other installed software, and to configure your system to automatically check for updates.



Information about the last GuardianOS update is listed at the bottom of the page and shows the basic information about the update.



CAUTION: Do not interrupt the update process. You may severely damage the server if you interrupt a software update operation before it is complete.

Update the GuardianOS



IMPORTANT: It is highly recommended that all active iSCSI users be disconnected before continuing.

1. Click **Check for Updates**.

If an update is available, follow the instructions on the page to download it.

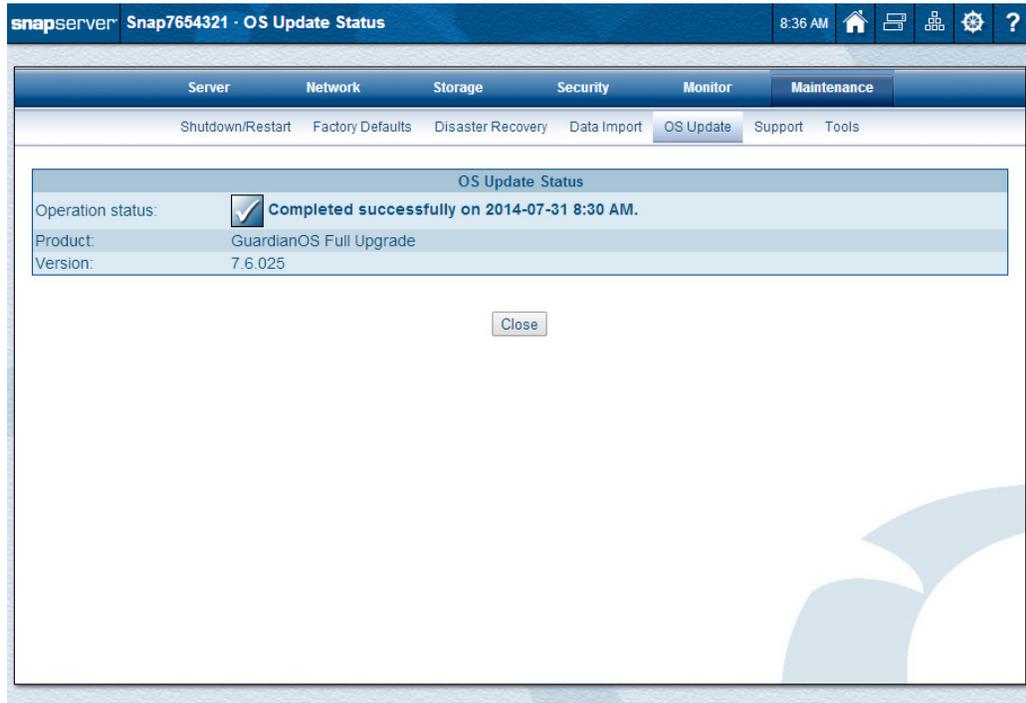
NOTE: If the server does not have access to the Internet, download the latest GuardianOS image (.gsu) or other software package from the [Overland Storage website](#) to a computer on the same network that the server can access.

2. On the **OS Update** page, click **Browse** (or **Choose File**, depending upon your browser), locate the file to be uploaded, and select it.

3. Click **Upload File** to start the upload to the server.

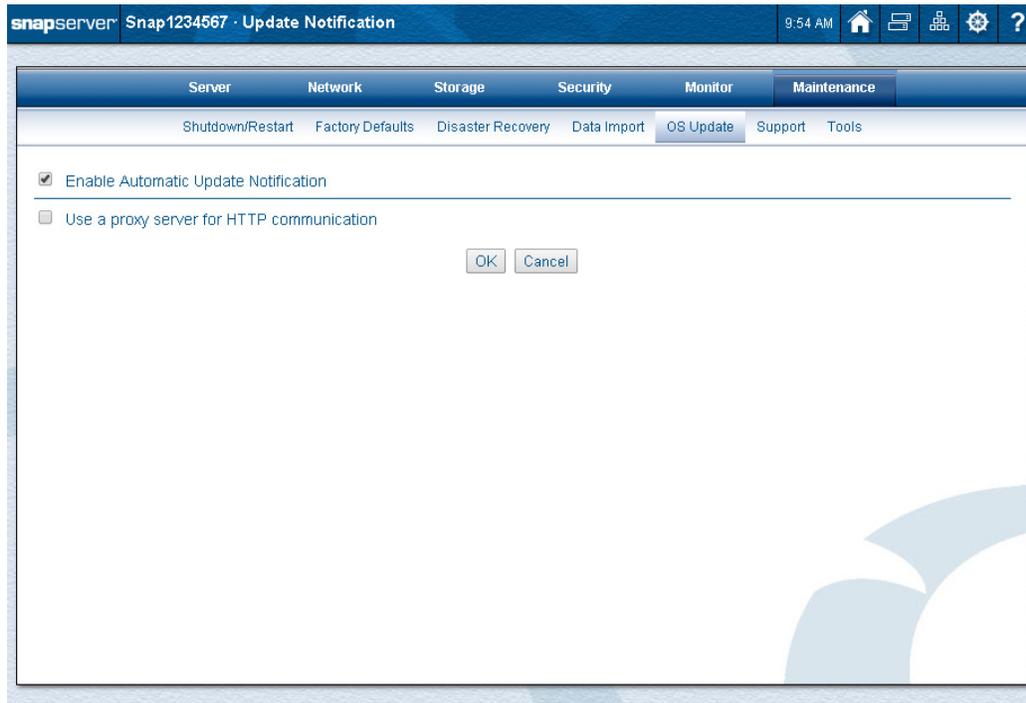
Only click the button once. Some browsers show the percent of the upload progress in their bottom status bar. The SnapServer uploads the software package and then prompts you to reboot the server to perform the update. Click **Restart for Update** (or click **Cancel** to abort the update).

After an upgrade and reboot, the **OS Update Status** page displays the success or failure of the last update performed.



Update Notification Option

You can configure GuardianOS to display an alert when updates are available for the server.

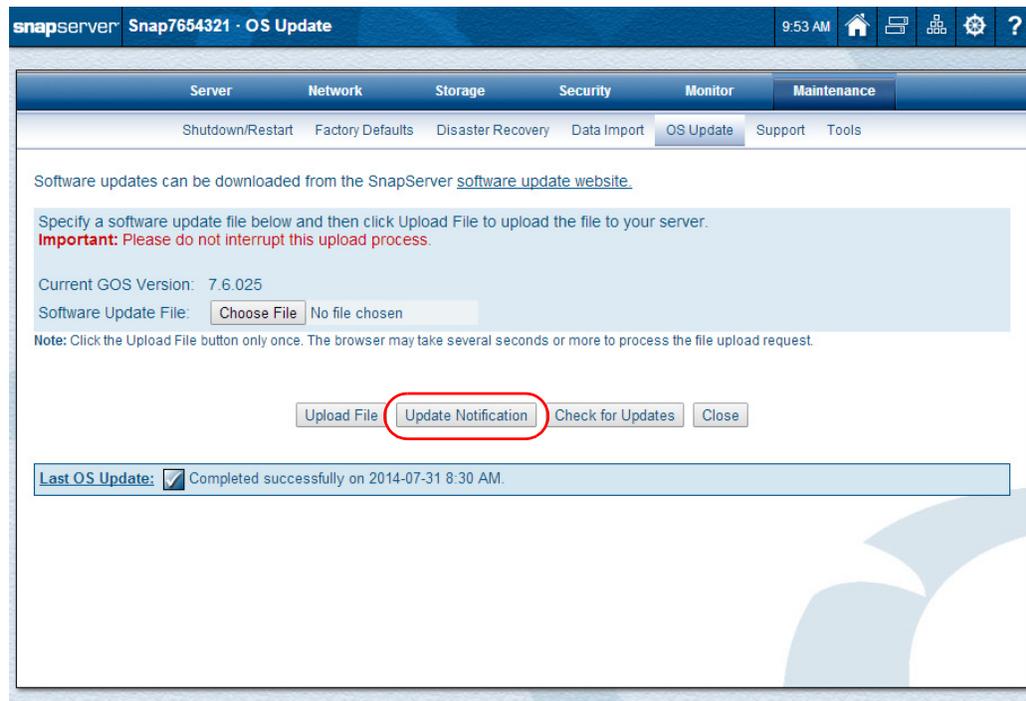


When enabled, **Update Notification** checks weekly for updates that are applicable to the server. If updates are available, a banner alert is displayed just below the menu bar on all Web Management Interface pages.

NOTE: You can choose to hide the banner by clicking either the *Remind me later* or *Hide this message* link on the banner. For *Remind me later*, the Web Management Interface displays the banner after the next check for updates; for *Hide this message*, the banner is hidden for the update in question until a later version is released.

Configuring Update Notification

1. Go to **Maintenance > OS Update** and click **Update Notification**:



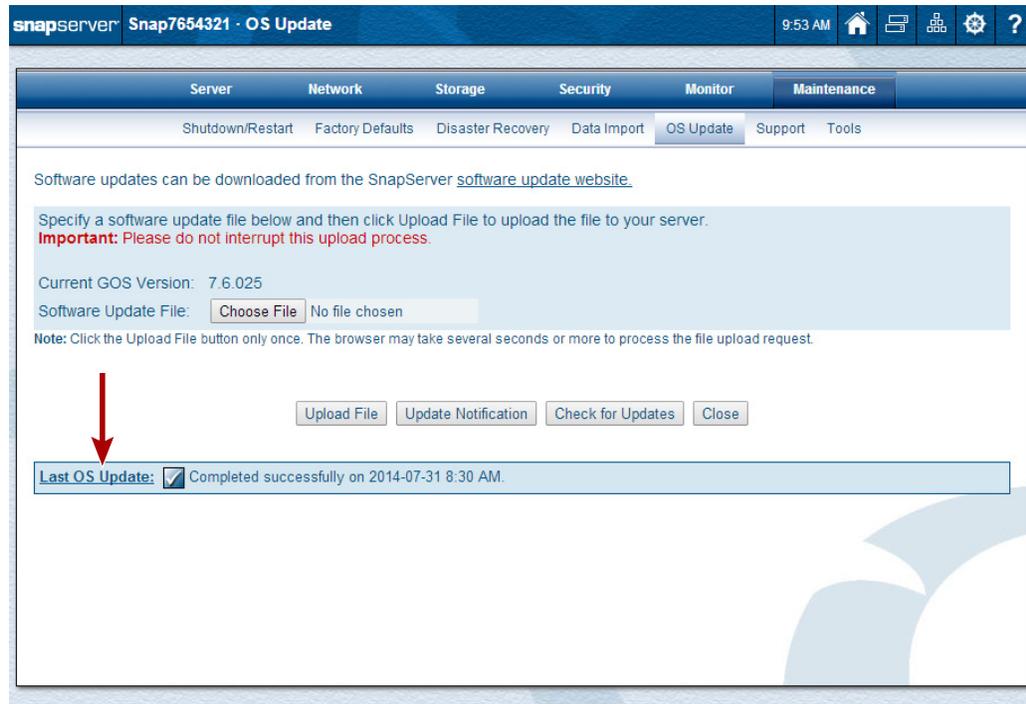
2. Check the **Enable Automatic Update Notification** box.

The screenshot displays the SnapServer web interface for the 'Update Notification' configuration. The top navigation bar includes tabs for Server, Network, Storage, Security, Monitor, and Maintenance. The 'Maintenance' tab is active, and the 'OS Update' sub-tab is selected. The main content area shows two checked checkboxes: 'Enable Automatic Update Notification' and 'Use a proxy server for HTTP communication'. Below the second checkbox are input fields for 'Proxy Host' and 'Proxy Port'. At the bottom of the form are 'OK' and 'Cancel' buttons.

3. If your environment requires using a **proxy server** for external web-based communication:
 - a. Check the **Use a proxy server for HTTP communication** box.
Additional proxy options are displayed.
 - b. Complete the **Proxy Host** and **Proxy Port** fields.
4. Click **OK**.

Last OS Update

At the bottom of the **OS Update** page is a **Last OS Update** link and information. Click this link to view a detailed status of the last update applied to the server.

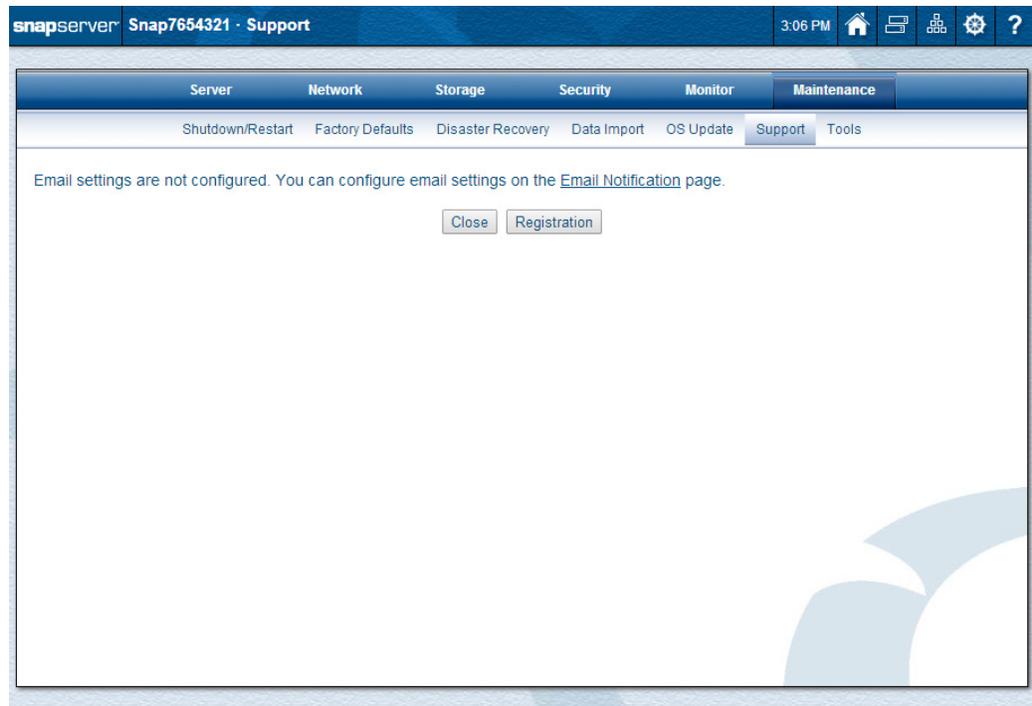


Support

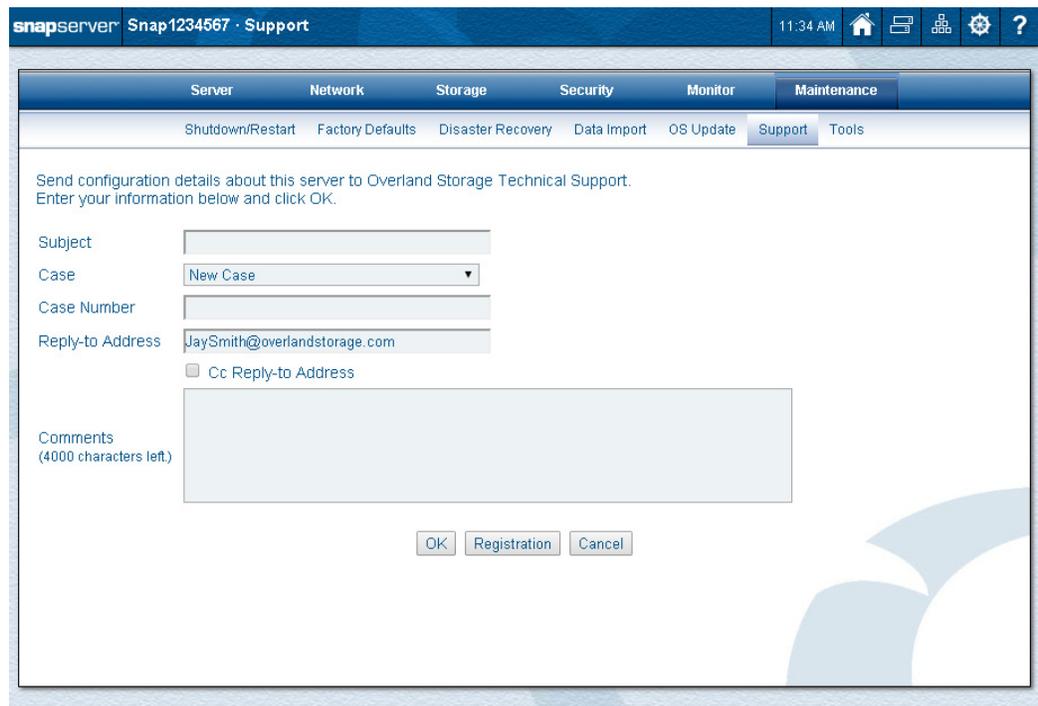
The **Support** page provides an easy way to contact Overland Technical Support, and transmit system logs and files that contain information useful for troubleshooting purposes.



IMPORTANT: The **Support** page is not accessible until you have configured **Email Notification** in the **Tools** submenu.



Once email is configured, the **Support** page is available with your contact information entered:

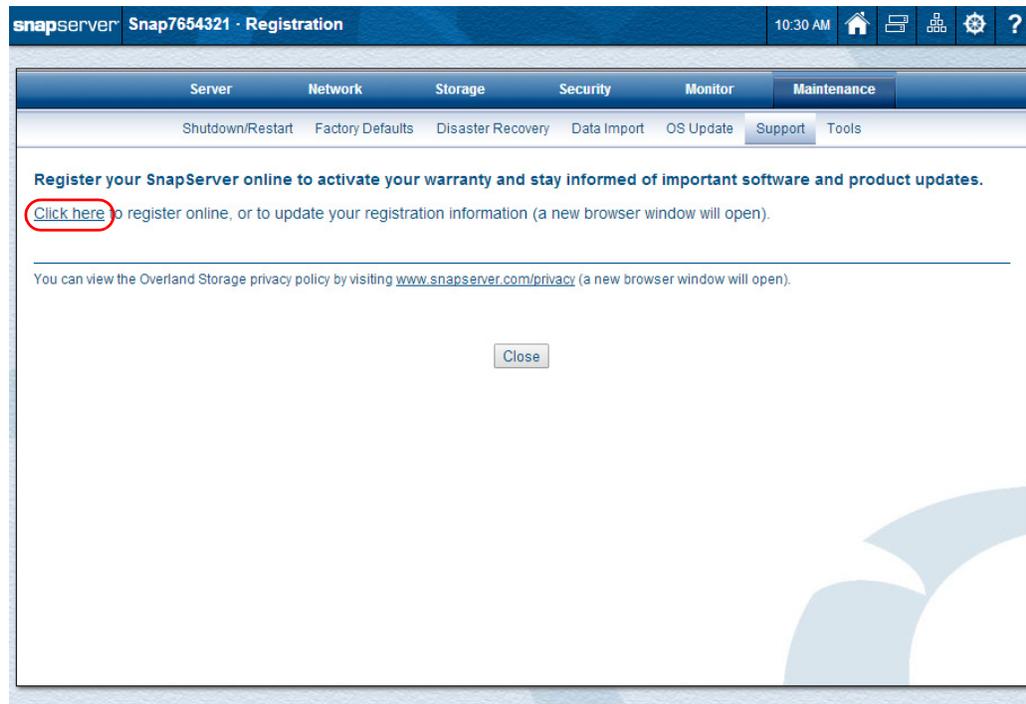


Registering Your Server

The very first time you start your server, a **Registration Reminder** page appears. 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>.

If you skipped the registration during setup, to register the server now, click **Registration** on the **Maintenance > Support** page:



To Register Your Server

NOTE: To use this feature, access to the Internet is required.

To register your server to activate its warranty support, you can either:

- Click the link on the initial **Registration Reminder** page.
- Go to **Maintenance > Support** and click **Registration**.

Click the **Click Here** link to launch the Overland Storage Support website and register online.

1. At the [Site Login](#), enter your **e-mail address** and **password**, and click **GO**.

If you are not yet a member, follow the **New Member** link to get set up.

E-mail:

Password:

Remember Me

GO >

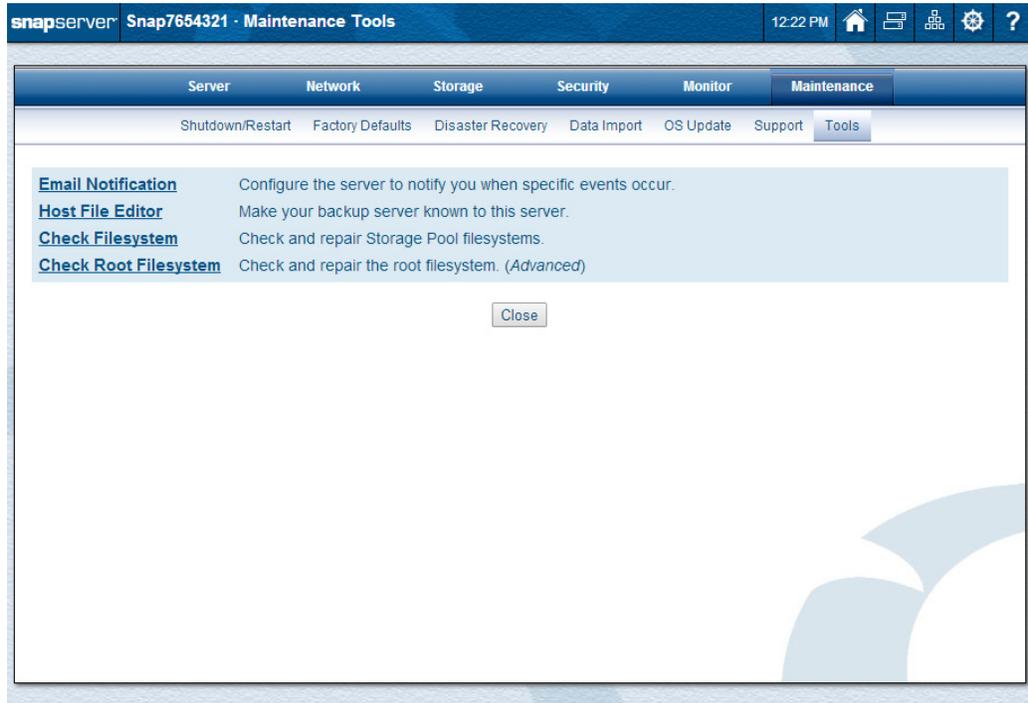
[Forgot your password?](#) [New member?](#)

2. At the **Confirm Automated Product Registration** page, enter the **date**, **reseller**, and **product site**.
3. Click **Confirm** to complete the process.

Once you have registered, you will receive a confirmation email to complete the registration.

Maintenance Tools

The **Tools** option provides a submenu of general-purpose maintenance options and features.



Email Notification

To configure the server to send email alerts in response to system events or activate Overland support, navigate to **Maintenance > Tools > Email Notification**.

The screenshot shows the 'Email Notification' configuration page in the SnapServer web interface. The page is titled 'Snap7654321 · Email Notification' and is part of the 'Maintenance' section under 'Tools'. The configuration options are as follows:

- Enable Email Notification
- SMTP Server: (Host name or IP address)
- SMTP Port: (Port number for SMTP server)
- Use Authenticated SMTP
- Use Secure Connection
- Email Address of Sender:
 - Use default: Snap7654321@devnet.myoverland.net
 - Use specific:
- Email Addresses of Recipients:
 - (optional)
 - (optional)
 - (optional)
- Send email notification for the following events:
 - Server shutdown/restart
 - RAID Set event
 - Volume is full
 - Hardware event
 - Printing event
 - Administrative operation event
 - License event
- Send a test email to listed email addresses upon saving settings.

Buttons for 'OK' and 'Cancel' are located at the bottom of the form.

To set up email alerts, you need the SMTP server's IP address and the email address of each recipient (up to four) who is to receive the alert.

Configuring Email Notification

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

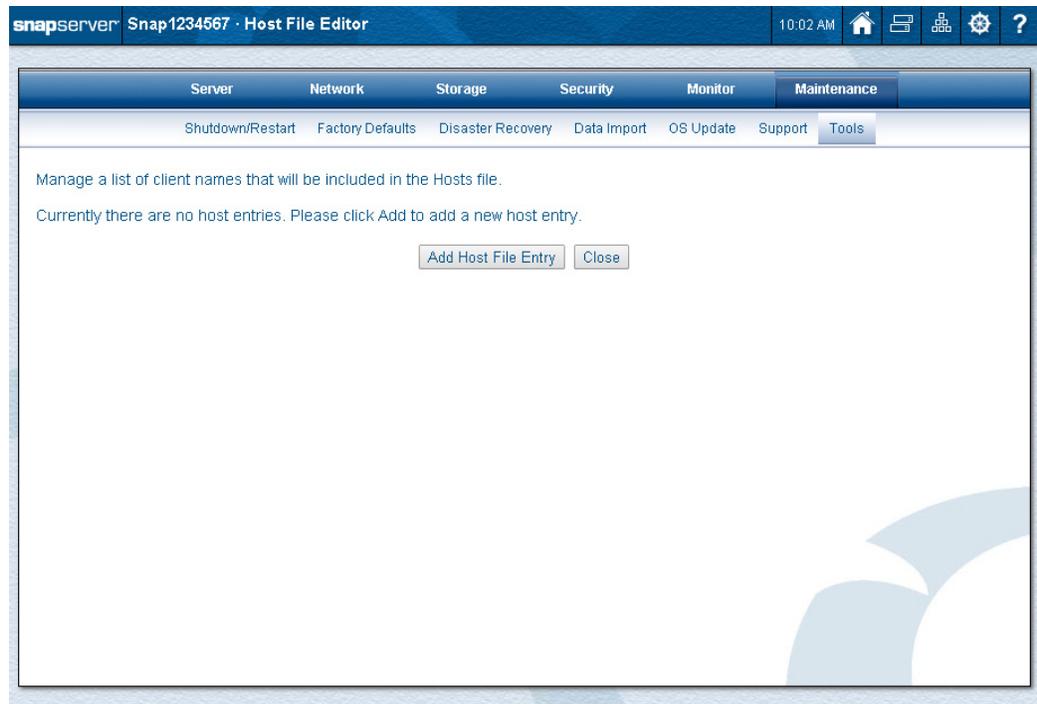
Option	Description
Enable Email Notification	To enable email notification, check the Enable Email Notification box.
SMTP Server	Enter a valid SMTP server IP address or host name.
SMTP Port	Enter a port number for the SMTP server or accept the default. The default is 25.
Use Authenticated SMTP	Check this box to authenticate 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. The types of methods supported (in order) are CRAM-MD5, LOGIN, and PLAIN.
Use Secure Connection	Check this box to encrypt emails from the server. STARTTLS and TLS/SSL encryption protocols are supported.

Option	Description
Email Address of Sender:	<p>Choose one:</p> <ul style="list-style-type: none"> The default address (<i>server_name@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>server_name@ipaddress</i>). Specify a specific sender.
Email Addresses of Recipients	Enter the email addresses to receive the notifications. One address is required but as many as four email addresses can be entered.
Send Email Notification	<p>Check the boxes next to the events you wish to be notified about:</p> <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 disable the SnapExtension.
Send a Test Email	To verify your settings, check Send a test email to listed email addresses upon saving settings , then click OK .

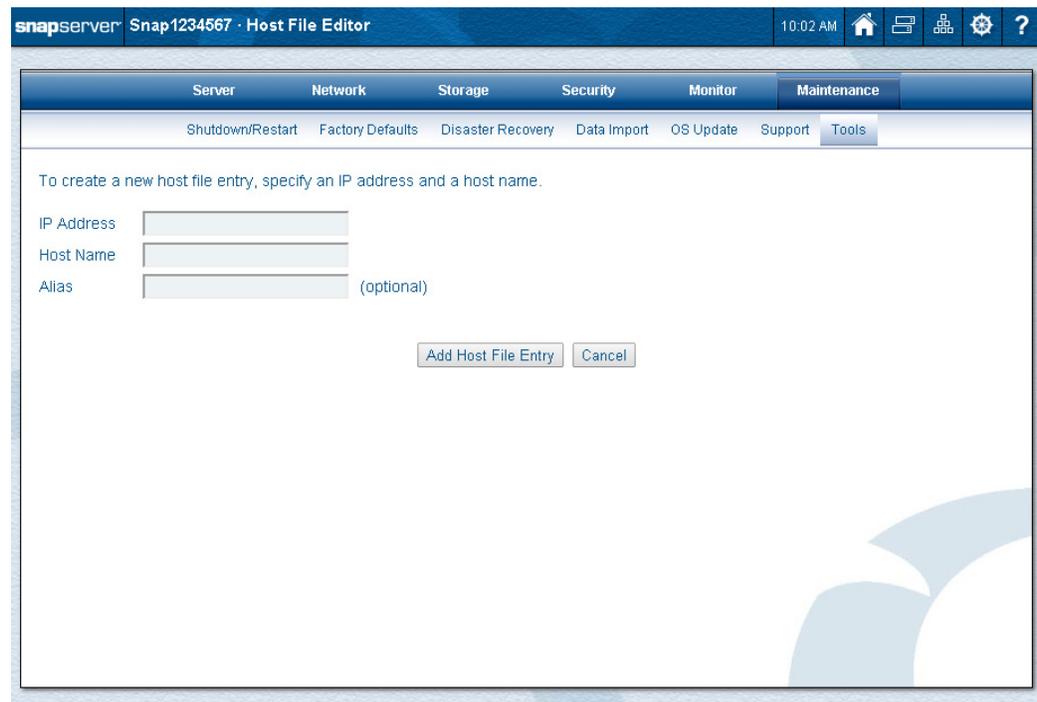
If **Send a Test Email** is checked, when you save your changes, an email is sent to all configured email recipients.

Host File Editor

Use this page to identify external hosts in the hosts file for the SnapServer. 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 table below, and then click **Add Host File Entry** again.



Use this table to complete the options shown:

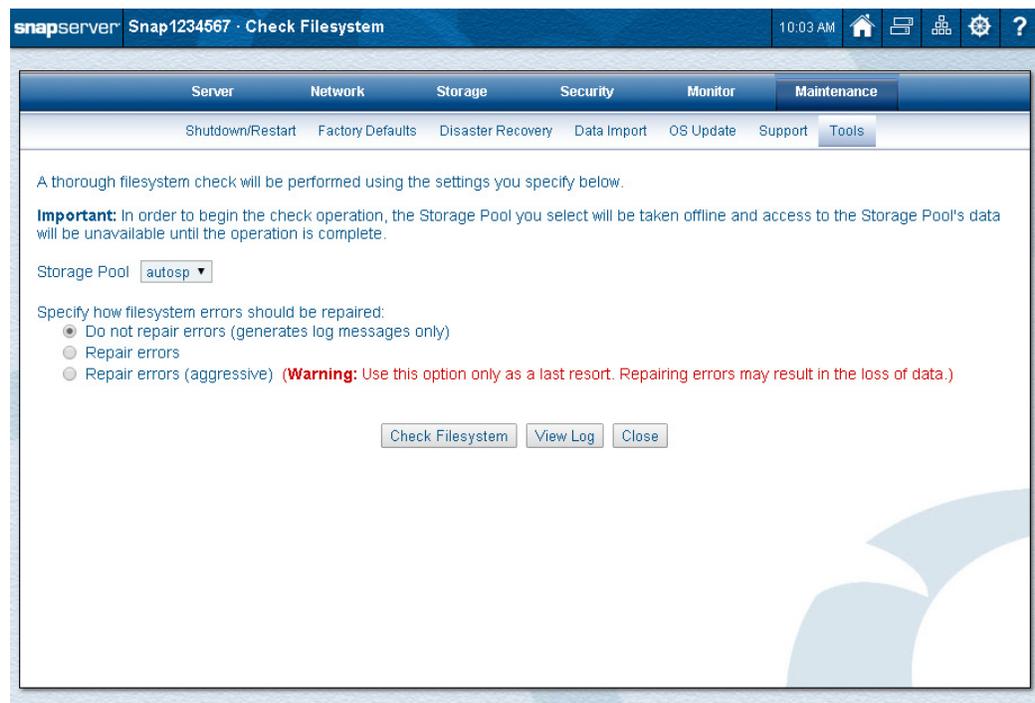
Option	Description
IP Address	The IP address of the external host.
Host Name	Enter the fully qualified hostname for the external host, using the format: <i>myserver.mydomain.com</i> . NOTE: Some applications may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.
Alias (optional)	Enter an optional abbreviated address for the external host, using the format: <i>myserver</i> . NOTE: Some applications 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.

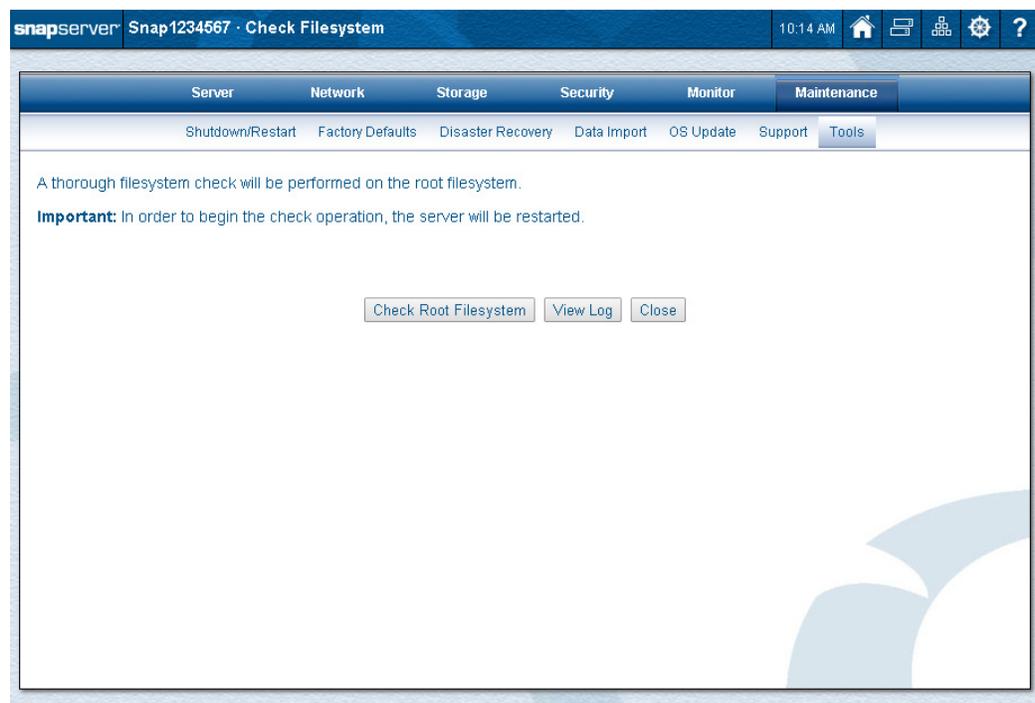
1. In **Maintenance > Tools**, click **Check Filesystem**.
2. From the drop-down list, select the **volume** (Traditional RAID) or **storage pool** (DynamicRAID) to be checked.
3. Choose the **type** of repair operation:
 - **Do not repair errors (generates log messages only)** – Checks for errors, but 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 **View Log** 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 **Check Root Filesystem**.
3. Click **Check Root Filesystem** again on the confirmation screen.
A reboot is required and takes place automatically.
4. After the server reboots, to view a log of the results, click **View Log**.

The GuardianOS site map (⚙️) provides links to a majority of 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 are the focus of this chapter.

With the exception of **Mgmt. Interface Settings**, these options are also directly navigable from the various menus in the Web Management Interface. Also the **Home**, **Snap Finder**, **SnapExtensions**, **Site Map**, and **Help** options are accessible from any page by clicking their respective icon in the top right corner of the page (see the table in [Web Management Interface on page 32](#)).

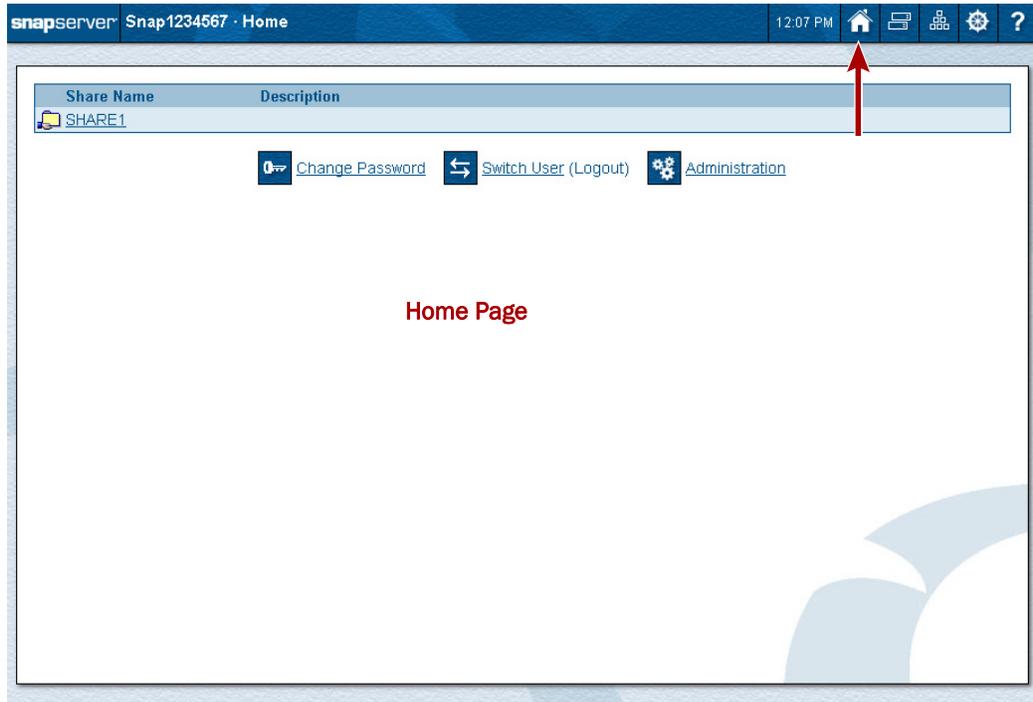
snapserver						
Server	Network	Storage	Security	Monitor	Maintenance	Misc.
Server Name	Information	Storage Pools	Security Guides	System Status	Shutdown/Restart	Administration
Date/Time	TCP/IP	Volumes	Shares	Active Users	Factory Defaults	Home
SSH	Windows/SMB	> Create Volume	> Create Share	Open Files	Disaster Recovery	SnapExtensions
UPS	Apple/AFP	Snapshots	Local Users	Network Monitor	Data Import	Snap Finder
Printing	NFS	> Create Snapshot	> Create Local User	Event Log	OS Update	> Snap Finder Properties
	LDAP/NIS	> Snapshot Schedules	> Password Policy	Tape	> Update Notification	BitTorrent Sync
	FTP	iSCSI	Local Groups		> Check for Updates	Change Password
	SNMP	> Create iSCSI Disk	> Create Local Group		> OS Update Status	Mgmt. Interface Settings
	Web	> VSS/VDS Access Control	Security Models		Support	
	ISNS	Disks	ID Mapping		> Registration	
		RDX QuikStor	Home Directories		Tools	
					> Email Notification	
					> Host File Editor	
					> Add Host	
					> Check Filesystem	
					> Check Root Filesystem	

Topics in Misc. Options

- [Home Pages](#)
 - [Home Page](#)
 - [Administration Page](#)
- [SnapExtensions](#)
 - [BitTorrent Sync](#)
 - [Snap EDR](#)
- [Snap Finder](#)
 - [Edit Snap Finder Properties](#)
- [Change Password](#)
- [Management Interface Settings](#)

Home Pages

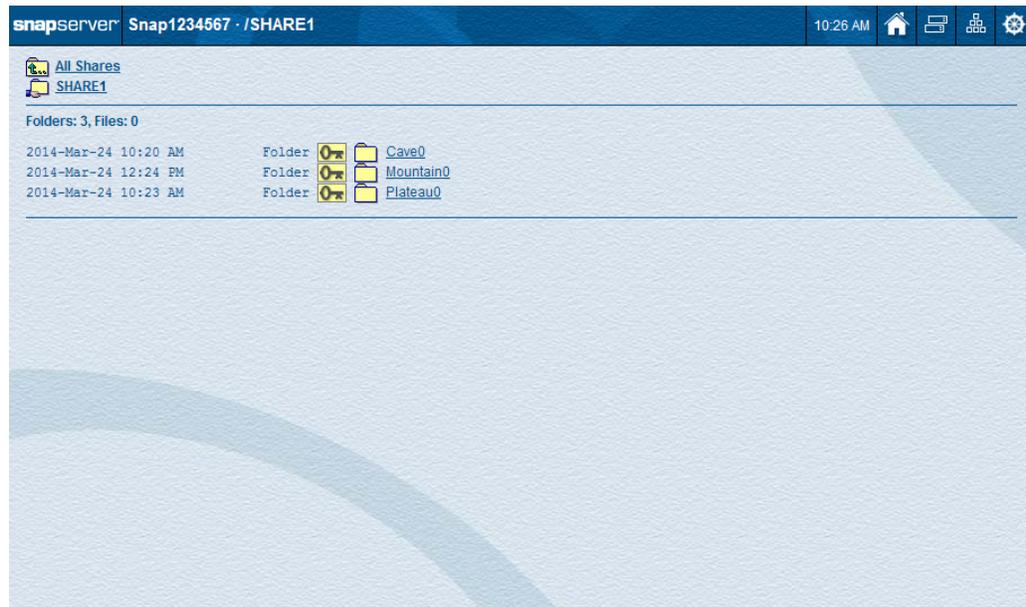
When you first launch the Web Management Interface, the **Home** page is displayed showing any existing shares and three options. Once logged in using the **Administration** link, you can switch between the **Home** page and the **Administration** page using the Home page (🏠) icon on the button bar.



Home Page

The Web Management Interface **Home** page displays a list of all shares and three basic options. Users can navigate the share structure to locate and view or download files without logging in but they cannot modify or upload files.

Click a share name to see a list of files:

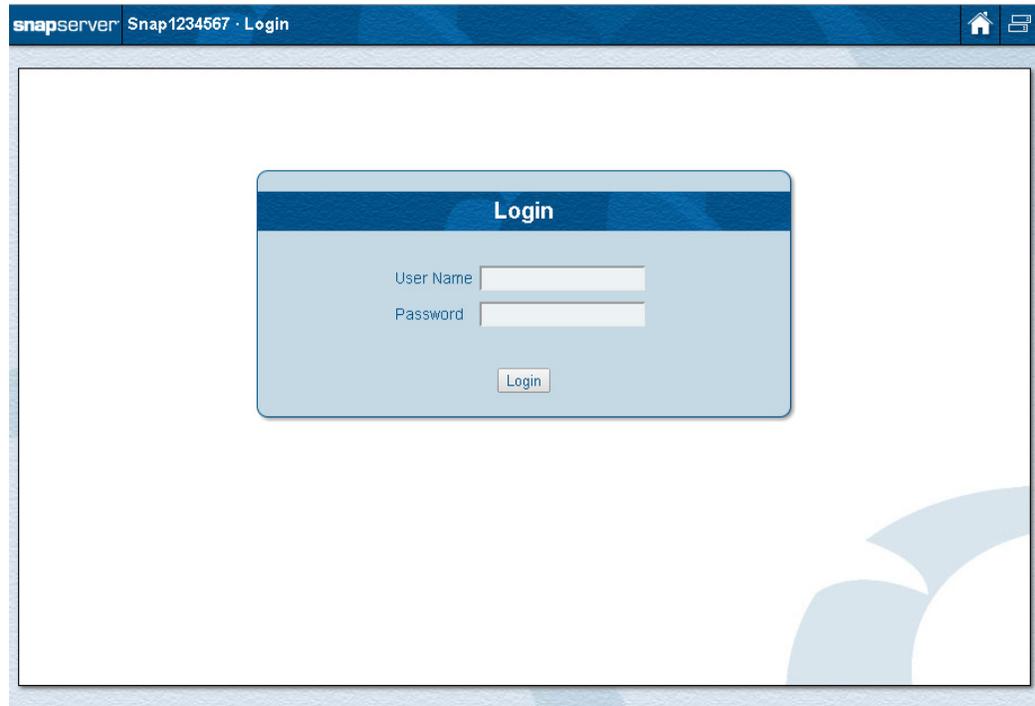


For users with admin rights, a key icon () appears next to the file/folder on the share. Clicking this icon displays a popup box with security information about the file/folder.

This page also provides three key administrative function links:

- **Change Password** () – 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 on page 276](#).

- **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 **Administration** page (see [Administration Page on page 267](#)). 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**).

Administration Page

The **Administration** page is accessible by clicking either the **Administration** link in the Site Map or the Administration (⚙️) or Home page (🏠) icons on the **Home** page. If web root is enabled, it can also be accessed directly by entering the address:

```
http://<server_name>/sadmin
```

in a web browser where *<server_name>* is the unique server name in the format Snapnnnnnnnnn. 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 by accessing online help. The tabs at the top provide access to the various functions and features of the GuardianOS.

The screenshot shows the SnapServer Administration interface. At the top, there is a navigation bar with tabs for Server, Network, Storage, Security, Monitor, and Maintenance. The main content area displays the Server Status box. On the right side of this box, the text "Auto-refresh is: OFF" is visible. Below the status box are "Refresh" and "Close" buttons. A link at the bottom of the page reads "Click here to find out what's new in GuardianOS 7.5 and this Web Management Interface."

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 (or **Refresh** below the Server Status Box) to manually refresh the information.

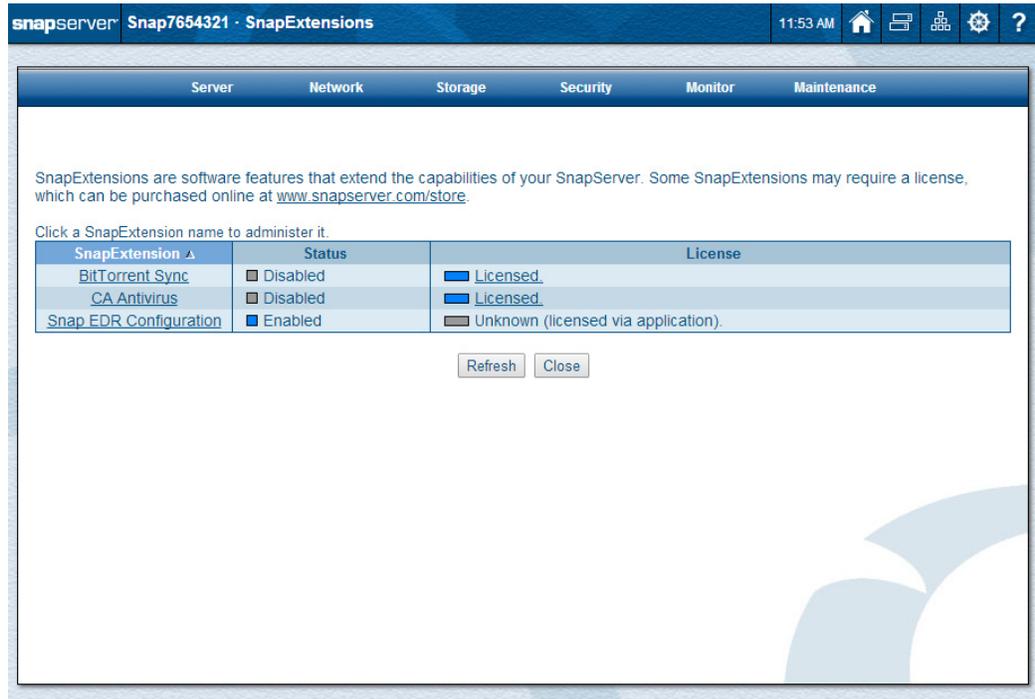
This screenshot is similar to the previous one, but the "Auto-refresh is: ON" text is now displayed with a red arrow pointing to it. Additionally, an auto-refresh icon (🔄) is now visible in the top right corner of the main content area, with a red arrow pointing to it from the right edge of the screenshot.

From the **Administration** page, clicking  takes you to the **Home** page.

SnapExtensions

The SnapExtensions icon  opens the SnapExtensions page. This page is used to manage the SnapExtensions installed on your SnapServer.

NOTE: Mouseover the icon to display a popup menu with direct access to SnapExtensions that are both installed and enabled.



SnapExtensions are software features that extend the capabilities of your SnapServer. Some SnapExtensions may require a license, which can be purchased online at www.snapserver.com/store.

Click a SnapExtension name to administer it.

SnapExtension 	Status	License
BitTorrent Sync	<input type="checkbox"/> Disabled	<input checked="" type="checkbox"/> Licensed
CA Antivirus	<input type="checkbox"/> Disabled	<input checked="" type="checkbox"/> Licensed
Snap FDR Configuration	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Unknown (licensed via application)

If any SnapExtensions are installed, you can click the SnapExtension name in the left column of the table to display the management page for that extension.

BitTorrent Sync



CAUTION: BitTorrent Sync bypasses share and file security. Be sure to only share data that is intended to be accessible by any user with the folder secret.

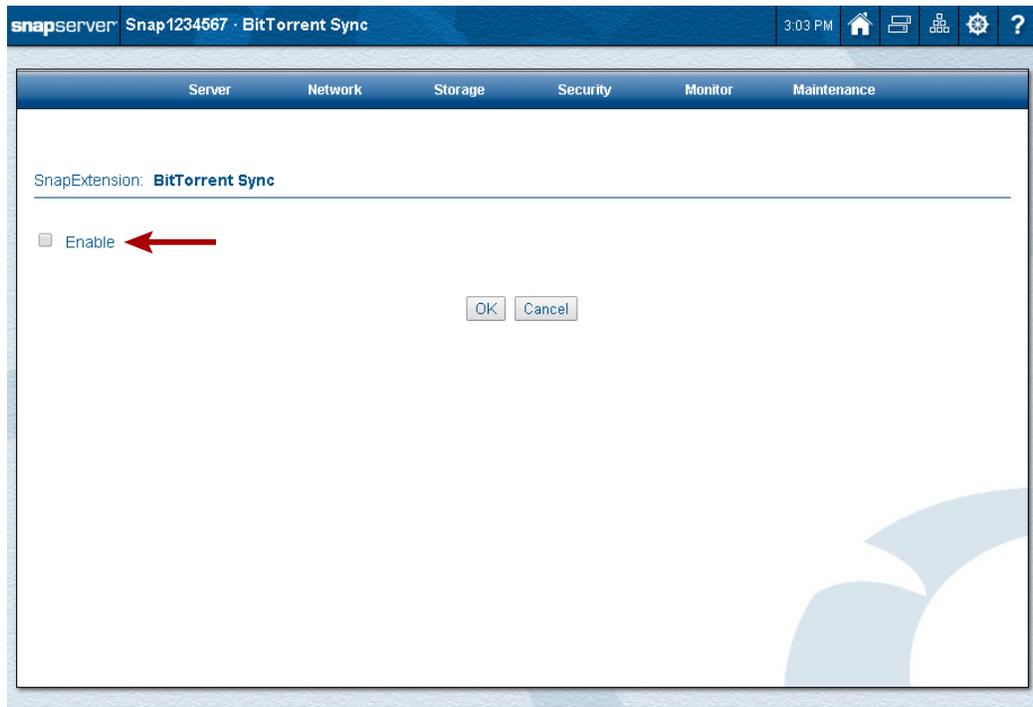
NOTE: Cookies must be enabled on your browser for BitTorrent Sync to work.

BitTorrent Sync (BTSync) is a SnapExtension that is preloaded on SnapServer. It lets you share and sync an unlimited number of files and folders of any size across multiple platforms. For more information, visit <http://www.bittorrent.com/sync>.

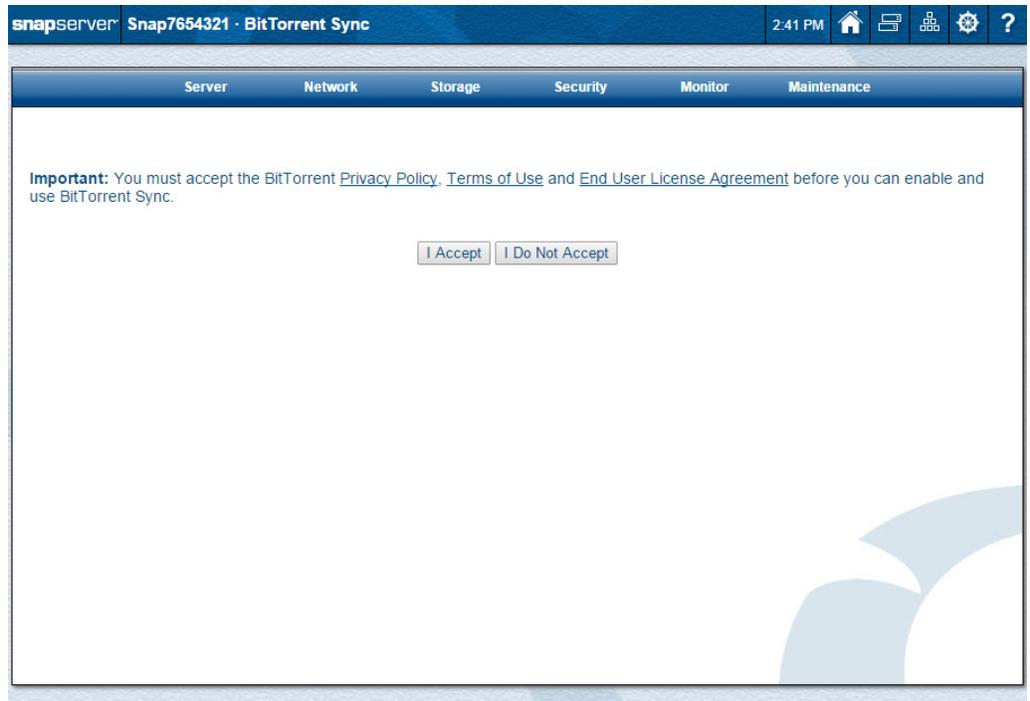
To use BitTorrent Sync, it must first be enabled:

1. On the **SnapExtensions** page, click the **BitTorrent Sync** name in the table to access the configuration page.

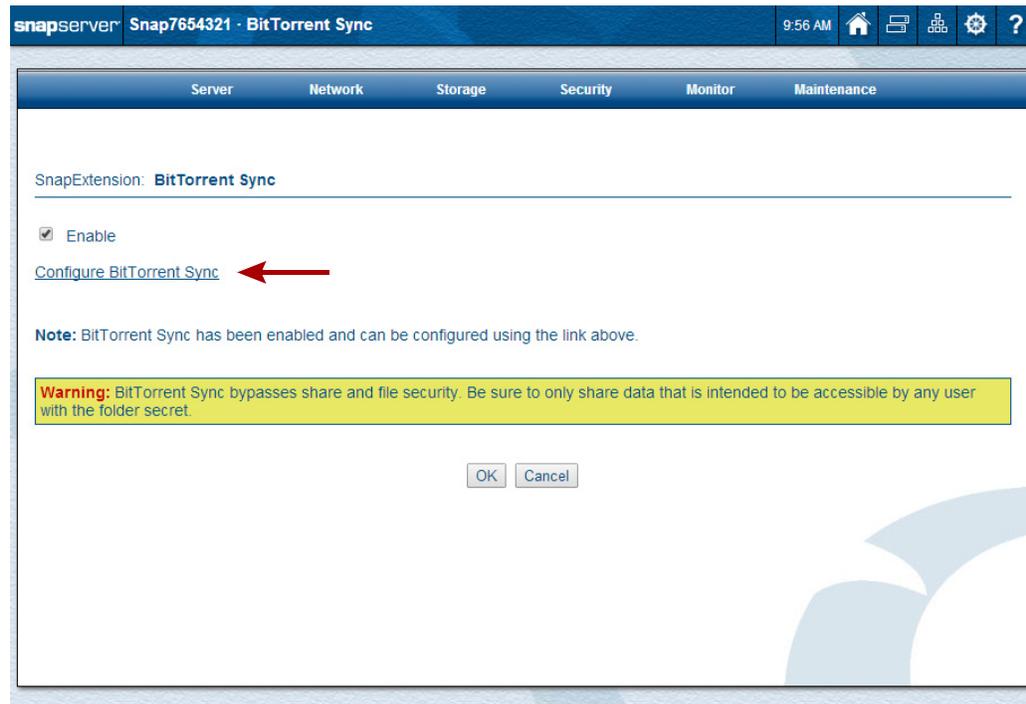
- At the configuration page, check **Enable** and click **OK**.



- At the following page, accept the BitTorrent Privacy Policy, Terms of Use, and End User License Agreement by clicking **I Accept** so you can run BitTorrent Sync.

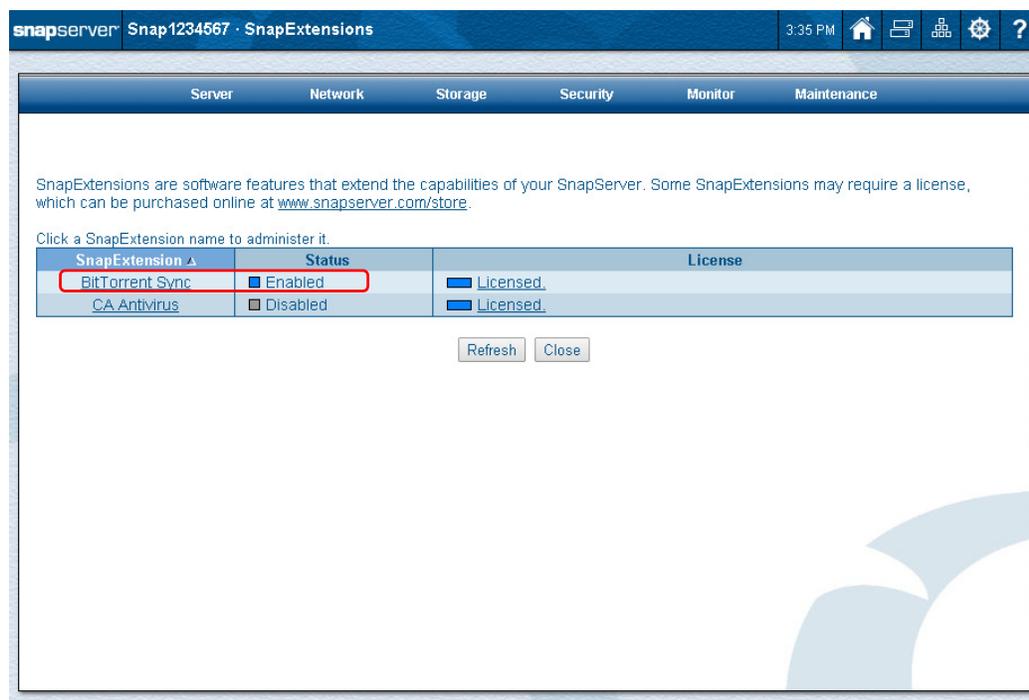


- When returned to the **BitTorrent Sync** page, use the **Configure BitTorrent Sync** link to configure it for your use.



NOTE: Once enabled, you can return at any later time to configure or reconfigure BitTorrent Sync by clicking **BitTorrent Sync** on the **SnapExtensions** page or the **Site Map**.

- When configuration is complete, click **OK**.
You are returned to the **SnapExtensions** page. **BitTorrent Sync** shows **Enabled**.



NOTE: To turn off (disable) the BitTorrent Sync feature, uncheck the BitTorrent Sync box.

BitTorrent Considerations

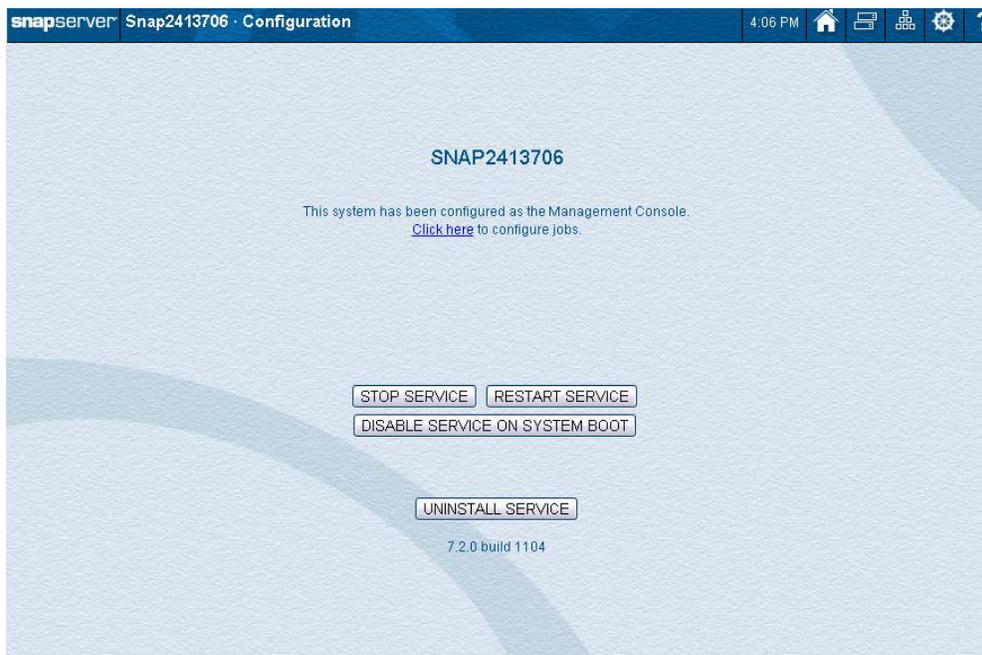
- For the most recent information and details on BitTorrent Sync configuration and use, refer to the documentation available on their web site.
- For sync destinations under a Unix or Mixed security model, BitTorrent Sync creates Unix personality files and directories, and sets the permissions to 644 (files) or 755 (directories) for Unix owner root/admingrp. As a result, all users and groups can read the files but only root can modify them. It is recommended that Windows-only security models are used at the destination, and that Windows permissions are set at the top level destination directory to apply the desired permissions to all files and subdirectories created by BitTorrent Sync.
- BitTorrent Sync cannot be used to replicate snapshots because it requires the ability to write to the sync location and snapshots are read only.
- BitTorrent Sync installs as a hidden directory on a volume (Traditional RAID) or storage pool (DynamicRAID). If the volume or storage pool is deleted or rolled back from a snapshot, the SnapServer attempts to automatically relocate the BitTorrent Sync hidden directory to another volume or storage pool. If there are no more volumes or storage pools, or if none can be found that are large enough, BitTorrent Sync is disabled and cannot be re-enabled until a suitable volume or storage pool becomes available. Once re-enabled, BitTorrent Sync must be completely reconfigured again.
- There are BitTorrent Sync mobile apps that make synced documents available on iOS, Android, Windows Phone 8, and Kindle Fire systems. Refer to their web site for details on configuration and use of BitTorrent Sync on mobile apps.

Snap EDR

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 on the **Configuration** page:



Snap Finder

Snap Finder (🔍) is a powerful tool that lists all the SnapServer appliances, SnapScale clusters, and Uninitialized nodes on your network (and on a remote network segment if so configured), and shows the current status of each. Click the unit name (if you have name resolution) or IP address of a cluster, node, or server to access it through the Web Management Interface.

NOTE: You can sort the columns (ascending or descending order) by clicking the column heading.

snapserver Snap1234567 · Snap Finder 10:29 AM

80 SnapServers. 16 SnapScale Clusters. 14 Uninitialized SnapScale Nodes.
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.
beryl	OK	10.25.3.27	GOS 7.5.048	DX1	2300028	4.34 TB	4.34 TB
BlueBottle	OK	10.25.2.18	GOS 7.6.0.jwinfieldsled11	DX2	2415534	735.27 GB	735.28 GB
bmgos4	OK	10.25.2.40	GOS 7.6.0.briansled11	VirtualSnap	15284780	3.07 GB	4.36 GB
bmgos5	OK	10.25.2.49	GOS 7.5.0.briansled11	VirtualSnap	15459346	5.25 GB	6.09 GB
bmgos8	OK	10.25.2.230	GOS 7.2.130	VirtualSnap	15922671	5.59 GB	6.09 GB
bmros41	Online	10.25.12.230	ROS 4.1.0.briansled11	-	-	25.07 GB	27.81 GB
bobbert	OK	10.25.3.38	GOS 5.0.133	4400	1723986	119.14 GB	280.71 GB
CB-110-SJSE	OK	10.25.3.59	GOS 6.5.029	110	2252267	548.71 GB	549.00 GB
CB-Meadowhawk	OK	10.25.2.108	GOS 7.5.047	DX2	2415100	133.08 GB	449.87 GB
CB-Sundragon	✘	10.25.15.60	GOS 7.5.047	DX2	2413126	1.34 TB	1.46 TB
CCCloudDX2	Online	10.25.17.230	ROS 4.1.075	-	-	31.54 TB	33.18 TB
CCDX1SNAP	OK	10.25.3.88	GOS 7.5.047	DX1	2301028	3.52 TB	6.52 TB
CCMeadowhawk-4G	OK	10.25.2.123	GOS 7.5.045	DX2	2411032	7.22 TB	7.22 TB
CCMeadowhawkEXP	✘	10.25.2.15	GOS 7.6.056	DX2	2414338	4.31 TB	4.31 TB
CCWAVE410	OK	10.25.2.152	GOS 5.2.067	410	2250681	517.50 GB	836.12 GB
CCWave412	OK	10.25.17.87	GOS 6.5.029	410	2277760	2.98 TB	4.36 TB
CCXSD400	OK	192.168.45.19	GOS 7.6.058-kdb	XSD 40	2353002	1.55 TB	2.15 TB
CCXSR120	✘	10.25.2.29	GOS 7.6.056	Unknown	2425096	4.70 TB	5.76 TB
CCXSR120EXP	OK	10.25.3.35	GOS 7.6.058	Unknown	2425094	36.47 TB	37.57 TB
CCXSR400	OK	10.25.2.81	GOS 7.6.057	XSR 40	2312722	4.74 TB	5.79 TB
CLW240	Online	10.25.12.240	ROS 4.1.0.charissa-devel	-	-	2.98 GB	3.99 GB
daedalus	OK	10.25.10.32	GOS 7.2.117	DX1	2302760	1.31 TB	2.00 TB
devqa	✘	10.25.11.2	GOS 6.5.029	N2000	730062	2.52 TB	2.70 TB

Refresh Properties Close

The following table describes the columns in the table:

Identification	Description
Server	Name of the SnapServer appliance, SnapScale cluster, or Uninitialized node. The default name is <i>Snapnnnnnnnn</i> , <i>Scalennnnnnnn</i> , or <i>Nodennnnnnnn</i> , where <i>nnnnnnnn</i> is the number of the server, node originally used to create the cluster, or the Uninitialized node. For example, "Snap23022161."
Status	<ul style="list-style-type: none"> The status of the SnapServer or Uninitialized node (for example, OK or Fan Failure). The status of a SnapScale cluster is always Online.
IP Address	The IP address of the SnapServer, Uninitialized node, or the Management IP address of the SnapScale cluster.
OS Version	The OS version currently installed on the SnapServer, Uninitialized node, or SnapScale cluster.
Model	The hardware model number of the SnapServer or Uninitialized node. This field is not applicable to a SnapScale cluster.
Number	The server or node number derived from the MAC address of the primary Ethernet port, used as part of the default name. This field is not applicable to a SnapScale cluster.
Avail Cap.	The available capacity on the SnapServer or SnapScale cluster. This field is not applicable to an Uninitialized node.
Total Cap.	The total capacity on the SnapServer or SnapScale cluster. This field is not applicable to an Uninitialized node.

To enable remote discovery of clusters, nodes, or servers on a different subnet or to display a warning icon for SnapServers or Uninitialized nodes with an enabled Ethernet port that has no link, click **Properties** at the bottom of the page to open the **Snap Finder Properties** page.

Edit Snap Finder Properties

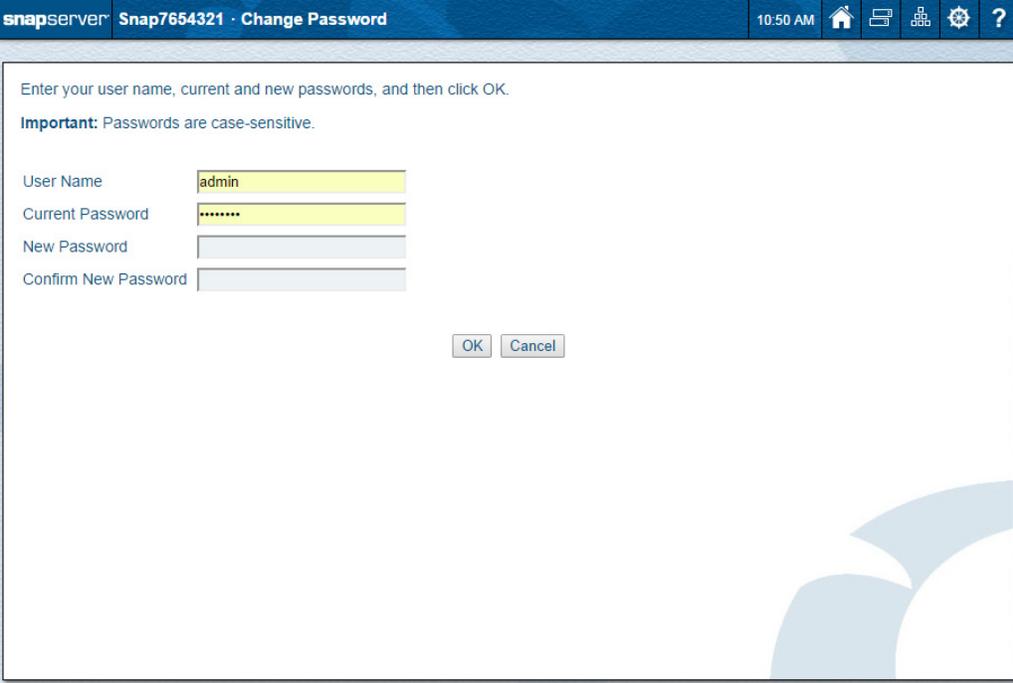
Anyone with administrative privileges can view or edit the Snap Finder properties. Click **Properties** to access the page.

From this page you can choose to display a warning icon for Uninitialized nodes or SnapServers with an enabled Ethernet port that has no link and enable remote discovery of units on a different subnet. Complete the following fields and then click **OK** to return to the **Snap Finder** page:

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 nodes or SnapServers 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 clusters, nodes, or SnapServers on a different subnet.
Add Server	Enter the host name or IP Address of a cluster, node, or server in the field to the right of the Add button, and click Add to incorporate it into the list of Remote Discovery Servers. Remote Discovery Servers send information about themselves as well as all other servers they've discovered on the remote network.
Delete Server	Select a cluster, node, or 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

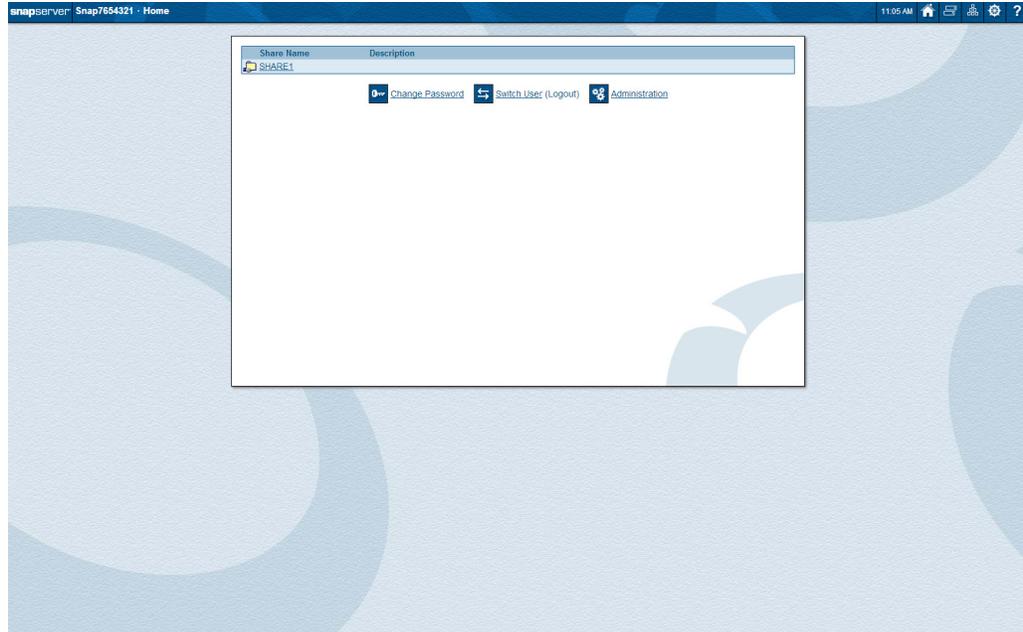
Confirm New Password

Changing Your Password

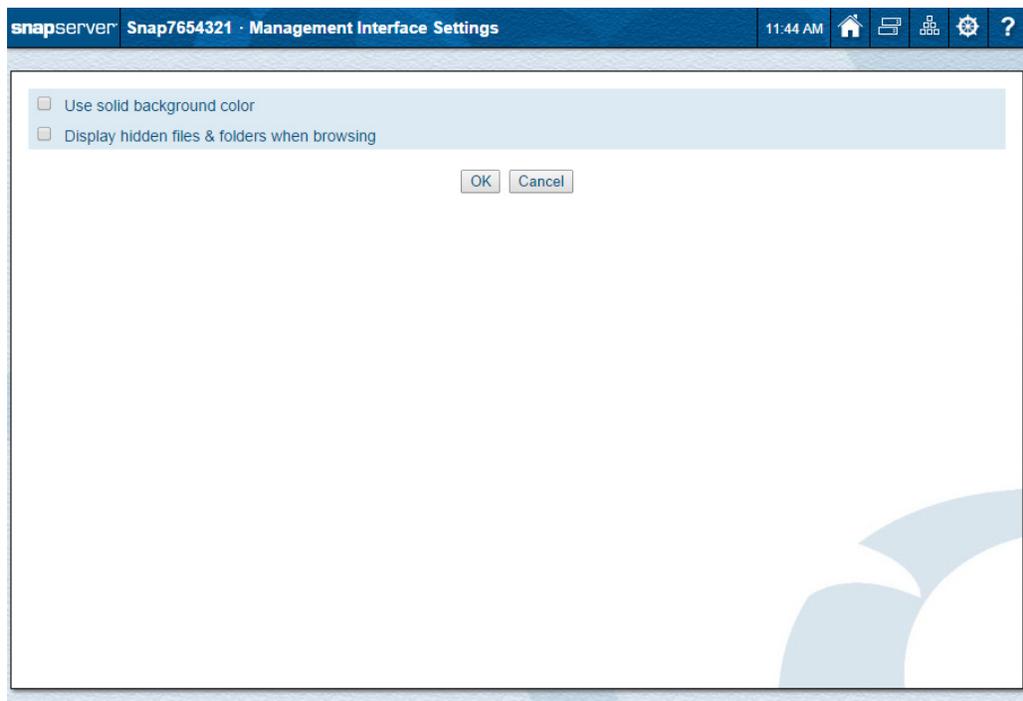
1. On the **Home** page, click the **Change Password** link (🔑).
2. At the **Change Password** page, enter your **User Name** and **Current Password**.
3. Enter and confirm your **new password**.
Passwords are case-sensitive. Use up to 15 alphanumeric characters without spaces.
4. Click **OK**.
5. At the confirmation page, click **OK** again.
You are returned to the **Home** page.

Management Interface Settings

The Web Management Interface default background is light blue with the stylized “O” symbols on a textured blue background:



This can be changed to a solid blue background on the Web Management Interface Settings page by clicking the Site Map icon (⚙️) to access **Management Interface Settings**.



Check the first box to use a solid color background (or clear the box to return to the standard background). Check the second box if you want to display hidden files and folders when browsing volumes for administrative configuration in the Web Management Interface.

This appendix provides a brief description of the supported backup solutions and the Snap Enterprise Data Replicator (Snap EDR) software.

Topics in Backup Solutions:

- [Backup and Replication Solutions](#)
- [Snap Enterprise Data Replicator](#)
- [Backup via SMB, AFP, or NFS](#)
- [Off-the-Shelf Backup Solutions](#)
- [iSCSI Disk Backups](#)

Backup and Replication Solutions

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:

- Data and security metadata backup and replication can be performed using the built-in Snap EDR.
- Backup over network file protocols can be performed using various backup packages that can access the server via SMB, AFP, or NFS.
- Backup from the server or to a tape attached to a server can be performed using supported backup agents and media libraries installed on a server.

Snap Enterprise Data Replicator

Snap EDR provides server-to-server synchronization by moving, copying, or replicating the contents of a share from one cluster or server to another share on one or more different clusters or servers. It comes preinstalled on SnapServers and activates a 45-day free trial if configured as a Management Console.

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 SnapServers, SnapScale clusters, and Windows, Linux, or 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.

Snap EDR Usage

The Snap Enterprise Data Replicator software distribution comes preinstalled on the SnapServer but must first be installed in SnapExtensions and then configured before it's available for use.

All other Snap EDR installations (including another machine running as the Management Console that the server registers to, other Agents that register to a Snap EDR Management Console running on the server, or other Agents replicating to/from the server) need to be able to resolve the SnapServer server name to its IP address in order to interoperate properly with the server. This can be accomplished via a DNS host record, local hosts file entries, or other name resolution services in the environment.

Configuring Snap EDR

To configure the server as a Snap EDR Management Console or an Agent:

1. Click the **SnapExtensions** icon located in the upper right corner of the Web Management Interface.
2. If necessary, install the **software package**:
 - a. Run the **installation routine** from SnapExtensions.
SnapExtensions displays a Snap EDR link and the status **Not Installed**.
 - b. Click the **Snap EDR link** and confirm the installation.
Wait for the installation to complete. The **SnapExtensions** page then displays the **Snap EDR Configuration** link.
3. Click the **link** to launch the **Management Console/Agent** configuration page.
4. Select either **Configure as the Management Console** or **Configure as the Agent**.

NOTE: If you are configuring a server as an Agent, you must provide the server name (for a SnapServer) or cluster management name (for a SnapScale) of the Management Console. The server must be able to resolve the Management Console server name to the correct IP address.

5. Once the server is configured, select the following **options** from the page that appears:

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.

Scheduling Jobs in Snap EDR

To schedule jobs, click the **Snap EDR** link in the **Site Map** (under **Misc.**).

For complete information on scheduling jobs in Snap EDR, see the *Snap EDR Administrator's Guide*.

Backup via SMB, AFP, or NFS

A SnapServer can be backed up via standard file server access.

In this configuration, the backup server is set up to use SMB, AFP, or NFS to connect to the server, examine the file system, and then back up the data onto itself. No special agents or media servers are needed.

Off-the-Shelf Backup Solutions

Special Application Notes for installing the backup agent or media servers can be found on the Overland SnapServer Support website (<http://docs.overlandstorage.com/snapserver>).

NOTE: The backup packages shown in the Application Notes do not support the backup of Windows ACLs. If Windows ACL backup is critical, Overland Storage strongly recommends you create a disaster recovery image before you perform a backup.

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 145](#) 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]
```

- 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
```

- 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

- Close `regedit`.

The SnapServer VSS Provider should now be available to Backup Exec to use for VSS-based backups. Return to **iSCSI Disk Backups** main page.

This appendix provides additional information and configuration options about securing and accessing shares and files on the SnapServer. The GuardianOS supports share-, file-, and directory-level permissions for all local and Windows domain users and groups.

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*.

Files and directories are stored on the server on volumes (or the 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*.

The default security model on newly-created volumes is always Windows/Unix. It can be changed to either a Windows or Unix security model.

Topics in Shares and File Access:

- [Security Model Rules](#)
- [Security Model Directories](#)
- [Security Model Management](#)
- [Special Share Options](#)
- [File and Share Access](#)
- [File-level Security](#)

Security Model Rules

Files and directories created inside security models acquire the security personality and permissions according to the rules of the chosen security model.

Windows/Unix 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).

Windows Security Model:

- All files and directories will have the Windows security personality. Permissions will be inherited according to the ACL of the parent directory.
- The permissions of a file or directory can be changed by any Windows SMB user with sufficient rights to change permissions or ownership. Permissions cannot be changed by NFS, AFP, or FTP clients.
- The personality of files and directories cannot be changed on a Windows security model. All files and directories always have the Windows personality with a Windows ACL. Standard Unix permissions will appear as 777 (rwxrwxrwx), but only the permissions in the Windows ACL will be enforced.

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/Unix 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 as they inherit the security model from the top level.

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. This section covers more details about the special options and features of share security in these subsections:

- [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 one of these two procedures:
 - While creating a share, navigate to **Security > Shares > Create Share > Advanced Share Properties** and check the **Hide this Share** box.
 - Edit a share by selecting the share, clicking to expand **Advanced Share Properties**, and checking the **Hide this Share** box.

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:
 - While creating a share, navigating to **Security > Shares > Create Share > Advanced Share Properties** and enabling/disabling specific protocols.

- Edit a share 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 *jdoe* has Read-Only share access and belongs to the group *sales*, which has Read/Write share access, the result is that the user *jdoe* 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 *Volume1* (DynamicRAID) or *VOL0* (Traditional RAID). 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. 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 **Security > 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 [Chapter 8, Security Options](#).

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. > Home** link in the Site Map, two shares are displayed:

```
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 Home 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 appears to be empty. You can prevent the snapshot share from displaying on this Home 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
2014-02-25.120000
2014-03-01.000100
2014-03-07.020200
```

The *latest* directory always points to the most recent snapshot (in this case, **2014-03-07.020200**, or March 7th, 2014, 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.

Depending on their ability to cross bind mounts, locally-installed backup agents can access the snapshot share in one of two ways:

- via `/shares` (for example, `/shares/SHARE1_SNAP/latest`)
- via `/links` (for example, `/links/SHARE1_SNAP/latest`)

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/Unix 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/Unix 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 Windows security model always have the Windows security personality and permissions can only be set by Windows SMB clients. Files in a Unix security model always have the Unix security personality and permissions can only be set by non-SMB clients.

Windows ACLs

GuardianOS fully supports Windows NTFS-style filesystem ACLs, including configuration, enforcement, and inheritance models. Inside Windows/Unix and Windows security models, 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 Explorer 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) inside a Windows/Unix security model, 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 LDAP or NIS domain. Group membership of GuardianOS 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 is 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**). The ACL is 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 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 the standard Windows Explorer interface. 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 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.

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 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](#)
 - [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 below 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 and thus want RAID 0.

For more information on RAID types, see [Factors in Choosing a RAID Type on page 103](#).

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
A simple, scalable, flexible RAID solution is needed that takes little or no effort to manage.	X	
Be able to easily add more storage capacity as needed.	X	
Be able to change the parity level over time.	X	
Be able to adjust volume size as needed.	X	
Be able to optimize parity based on the number of drives inserted into the system.	X	
The SnapServer needs to be able to configure and manage the RAID array size and parity.	X	
I need to manually configure the RAID array.		X
I need to manually tune my storage system for specific needs, such as RAID levels and/or storage types.		X
User or group quotas are needed.		X
Different drives are mixed in the same chassis, and then these different drives are added together to make a homogeneous RAID.		X
The filesystem must span multiple chassis.		X
Both local and global spares are needed.		X
Snapshot rollbacks are needed.		X

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.
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.

Feature	DynamicRAID	Traditional RAID
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 RAIDs 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

These are the high-level steps to configuring DynamicRAID:

1. During setup, after selecting DynamicRAID, all available **disk drives** on the SnapServer are detected and displayed.
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.

The software wizard configures the SnapServer based on the parity selected.
3. Use the following **options** to fine-tune the configuration:
 - **Storage > Storage Pools** (see [Storage Pools on page 83](#))
 - **Storage > Volumes** (see [Volumes on page 96](#))
 - **Security > Shares** (see [Shares on page 179](#))

DynamicRAID 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 [Edit Volume Properties on page 98](#)).
- The snapshot space may be expanded (see [Adjusting Snapshot Space Size on page 141](#)).
- 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 [Add a Disk Drive to Upgrade Parity on page 95](#)).

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.

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 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>.

GuardianOS Ports

The following table lists the ports used by GuardianOS. The **GOS Feature** column lists access to the feature such as **Storage > iSCSI** (which you can access under the **Storage** tab in the **iSCSI** subsection of the Web Management Interface).

Port #	Layer	GOS Feature	Name	Comment
1	DDP		rtmp	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 > Windows/SMB	snmp	Simple Network Management Protocol (SNMP)

Port #	Layer	GOS Feature	Name	Comment
162	TCP & UDP	Network > Windows/SMB	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 > iSNS	iSNS	iSNS port
3260	TCP	Storage > iSCSI	iSCSI	iSCSI port
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
9050	TCP	Sysbroker		Sysbroker RPC Port

Port #	Layer	GOS Feature	Name	Comment
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	Network > NFS Network > FTP	NFS FTP (Passive)	Dynamically allocated in runtime for user connections

Command Line Interface

GuardianOS includes a command line interface (SnapCLI) that is 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 Web Management Interface or SSM.

 **IMPORTANT:** 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> Specify DynamicRAID mode
		<code>type=Traditional-RAID</code> Specify Traditional RAID mode
	force	<code>yes</code> 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](#)
- [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	Result
?	See an overview of the CLI, with a list of available commands and a description of command syntax.
{command} 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).
{command} 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	<code>get</code>	- <code>get</code> WEB properties
	<code>set [OPTIONS]</code>	- <code>set</code> WEB properties
	- <code>require-webview-auth=(yes no)</code>	- <code>require</code> HTTP/HTTPS clients to authenticate in order to access the server
	- <code>non-secure-http=(yes no)</code>	- <code>enable/disable</code> 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.

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 on page 301](#) 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>'
mkgroup()
{
```

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
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
else
    echo "Group '$2' does not exist."
fi
else
    echo "User '$1' does not exist."
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
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
else
    echo "User '$1' does not exist."
fi
```

```

        if [ $? -ne 0 ]; then
            echo "Deletion of user '$1' failed."
        fi
    return 1
    fi
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
fi
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
fi
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"

```

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.

Address

An address is a data structure or logical convention used to identify a unique entity, such as a particular process or network device.

Algorithm

A sequence of steps designed to solve a problem or execute a process.

ATA

Short for *Advanced Technology Attachment*. A standard interface for connecting storage devices to a PC.

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.

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.

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.

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. It can be used for ATM, token ring, 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).

Channel

A communications path between two computers or devices.

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.

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.

daemon

A process that runs in the background.

default gateway

The router used when there is otherwise no known route to a given subnet.

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.

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.

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.

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.

Failover

A strategy that enables one Ethernet port to assume the role of another port if the first port fails. 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.

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.

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.

FTP

Short for *File Transfer Protocol*. A standard Internet protocol that provides a way to exchange files between computers on the Internet.

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.

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 Home (HTTP/HTTPS), FTP, and AFP protocols. See also [SMB](#).

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 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.

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.

Initiator Device

An iSCSI 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.

I/O (Input/Output)

The operation of transferring data to or from a device, typically through an interface protocol like CIFS, NFS, or HTTP.

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.

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.

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.

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.

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 iSCSI operating systems.

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.

NIC

Short for *Network Interface Card*. A board that provides network communication capabilities to and from a computer.

NIS

Short for *Network Information Service*. It is a client–server directory service protocol for distributing system configuration data such as user and host names between computers on a computer network. Sun Microsystems developed the NIS; the technology is licensed to virtually all other Unix vendors.

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.

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.

PoP

Short for *Proof of Purchase*. The number used to obtain a license key for an upgrade to third-party applications.

Portal

A target's IP address together with its TCP port number used in iSCSI systems.

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.

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.

Quota

A limit on the amount of storage space on a volume that a specific user or NIS group can consume.

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.

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.

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.

SMTP

Short for *Simple Mail Transfer Protocol*. A TCP/IP protocol used for sending and receiving email.

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.

Storage Area Network

See [SAN](#).

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.

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.

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.

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.

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.

Symbols

> (menu flow indicator) 4

A

ACLs

- backing up 239
- resetting to defaults 238
- setting file-level permissions (Windows) 289

Active Directory

- and name resolution servers 60
- joining AD domain 63
- SnapServer interoperability with 61

Active Users page 227

admin password

- changing 276
- default 174
- resetting forgotten 237

Administration page 267

administration password 266

Advanced Share Properties 181, 183

AFP backup 281

AFP terminology 66

AFP, see *Mac OS*

alert definitions 4

alert messages 34

Application Notes 281

Authentication

- default settings 174
- HTTPS/HTTP 77
- Kerberos 61
- LDAP domain 71
- NIS domain 71

automatic disk incorporation 113

automatic shutdown 41

automatic update checking 252

B

background disk scan 113

backup

- coordinating with snapshots 139
- inability to back up iSCSI disks 137, 147
- iSCSI disks 147
- of server and volume settings 239
- off-the-shelf solutions 281

backup.acl 239

backup.qta.groups 239

backup.qta.users 239

BitTorrent Sync 269

C

CA Unicenter TNg 76

change password 266, 276

changing server name 37

CLI connection via SSH 39

client access, configuring

- Apple 65
- FTP 73
- HTTPS/HTTP 77
- NFS 67
- Windows SMB 62

Command Line Interface 299

- running scripts 301
- syntax 299

contact information pop-up 35

conventions, typographical 4

copying RDX media 169

create new share 119

create new volume 119

CUPS server 44

customer support 3

D

- data copying to RDX 169
- data import 244
- date and time settings 38
- defaults
 - admin password 174
 - TCP/IP 50
- directories, home 222
- Disaster Recovery
 - backing up server and volume settings 239
 - creating recovery files 240
- disk drives
 - adding
 - in DynamicRAID 84, 95, 293
 - in Traditional RAID 116
 - automatic incorporation 113
 - 113
 - detecting 21
 - hot swap 158
 - incompatible 84
 - previously configured 84, 293
 - in different system 23, 84
 - in Traditional RAID 117
 - reintegrating orphaned 164
 - replacing 158
- disk icons 95, 156
- Disk is Foreign icon 95
- domain search
 - authentication required 130, 188, 210, 214
- domains
 - joining ADS 176
- drive rotational speeds 107
- drive-size characteristics, DynamicRAID 294
- dual parity 158
- dynamic volumes 292
- DynamicRAID 82
 - compared to Traditional RAID 292
 - drive indicators 294
 - how it works 293
 - implementation 293
 - storage pools 83
 - volumes 96

E

- Email Notification 258
- Email Notification page 30

- Ethernet, see *Gigabit Ethernet*
- Event Log page 233
- Expand Volume button 124
- expansion units
 - configuring initial setup of 25
 - DynamicRAID usage 90
 - integrating orphans 166
 - Traditional RAID usage 112
- exports file, NFS 179

F

- failover, see *Network bonding*
- file/folder security information 266
- files, setting permissions for 288
- format RDX media 170
- FTP
 - connecting via 74

G

- GID 175
- Gigabit Ethernet
 - autonegotiation required 53
- global hot spares 104
- Global Spare 164
- Groups
 - creating local 201
 - file-level access for 288
- GuardianOS specifications 11

H

- hardware information pop-up 35
- Head Unit properties page 90
- home directories 222
- Home page 265
- home pages 265, 267
- hot spares 104
- hot swap
 - automatic incorporation of disks 113
 - disk drive 158
- HP Open View 76
- HTTPS/HTTP, configuring 77

I

- ID mapping 209

- Initial Setup
 - configuring storage
 - expansion units **25**
 - in DynamicRAID **22**
 - in Traditional RAID **24**
 - Initial Setup Wizard **17**
 - integrate expansion units **166**
 - internal temperature, e-mail notification of **259**
 - IP address
 - setting **52**
 - IPP port number
 - Linux **45**
 - Windows **45**
 - iSCSI disks **145**
 - and DynamicRAID **294**
 - backing up **153**
 - configuring iSNS **81**
 - creating **149**
 - LUNs **155**
 - multi-initiator support **147**
 - write cache options **147**
 - iSNS **81**
- ## K
- Kerberos **61**
 - key icon **266**
- ## L
- LDAP
 - configuration **71**
 - domains **71**
 - LEDs
 - in DynamicRAID **294**
 - Link Aggregation (802.3ad) **51**
 - load balancing, configuring server for **52**
 - local groups **200**
 - local hot spares **104**
- ## M
- Mac OS
 - configuring client access **65**
 - Finder **66**
 - Macintosh, supported OS versions **12**
 - maintenance
 - data import **244**
 - disaster recovery **238**
 - factory defaults **237**
 - OS update **250**
 - shutdown and restart **236**
 - support **254**
 - tools **257**
 - managing snapshots **135**
 - mapping, ID **209**
 - menu flow indicator **4**
 - message in Web Management Interface **34**
 - mixed drive capacities **292**
 - mixed drive types **292**
 - monitor
 - active users **227**
 - event log **233**
 - hardware status **226**
 - open files **228**
 - options **225**
 - tape devices **234**
 - monitor network traffic **229**
 - Multihomed configurations **53**
- ## N
- network
 - access **47**
 - bonding, cabling requirements for **54**
 - current settings **48**
 - reset to factory defaults **237**
 - Network Monitor page **229**
 - network monitoring
 - download usage records log **232**
 - graphing options **232**
 - option icons **231**
 - viewing usage **230**
 - Zoom Bar **232**
 - NFS
 - access **67**
 - configuring **67**
 - exports file **179**
 - read-only share access **68**
 - share-level permissions **190**
 - NFS backup **281**
 - NIS
 - configuration **72**
 - domains **71**

O

- Open Files page 228
- operating system 3
- orphaned disk drives 164
- orphaned expansion units 166
- OS update
 - checking for 250
- Overland technical support 3

P

- parity
 - adding disk drives to upgrade 95
 - configuring DynamicRAID 293
 - disk drive failure
 - dual 95
 - single 95
 - increasing protection 95
 - management 94
 - options 292
 - Parity Mode 85, 91
 - Parity Mode change 162
 - removing a drive 161
- parity, dual 158
- password
 - changing 276
 - default for admin account 174
 - unlock 197
- paths
 - connecting via web browser 78
 - for backing up snapshots 140
- permissions
 - share- and file-level interaction 187
 - file-level, default behavior 288
- previously configured drives 113
- Print Server 44
 - adding a printer 45
 - canceling print jobs 46
 - configure the printer 44
 - deleting a printer 46
 - IPP port number, Linux 45
 - IPP port number, Windows 45
 - monitoring print jobs 45
 - pausing the printer 46
- product documentation 3

Q

- Quota
 - Sort drop-down list 132
 - View drop-down list 132
- Quotas
 - adding 131
 - backing up configuration 239
 - defaults 126
 - displaying 131
 - properties 128
 - usage calculation 126

R

- RAID
 - adding disk drives to 116
 - choosing 103
 - creating new 105
 - grouped
 - deleting 112
 - grouping 109
 - multiple RAIDs 112
 - with other grouped RAIDs 112
 - scrubber 113
 - sets
 - creating new 105
 - grouping 109
 - screen 105
 - settings 113
 - Traditional RAID and replacement disks 116
 - type selection 19
- RAINcloudOS specifications 11
- RDX media
 - copying data 169
 - ejecting 171
 - formatting 170
 - properties page 168
 - renaming a volume 172
- RDX media formatting 170
- RDX QuikStor 166
- reboot, setting up alert for 259
- reduced parity with Traditional RAID 116
- refresh RDX properties page 169
- Registration page 29, 255
- remote discovery 275
- replacing disk drives 158
- replication 279

restart **236**
 resynchronization, setting alert for completion of **259**
 rotational speeds of drives **107**

S

Secure Shell (SSH) **39**
 security
 guides **175**
 models **205**
 resetting default ACLs for volumes **238**
 shares **179**
 Windows ACLs **288**
 server
 and volume settings, backing up **239**
 changing server name **37**
 server events notification **258**
 Server Number **226**
 server registration, online **255**
 setting e-mail alerts **259**
 Shares **179**
 backing up configuration **239**
 delete **185**
 edit properties **183**
 shares with Admin rights **266**
 shutdown **236**
 Simple Network Management Protocol, see *SNMP*
 Single-subnet configuration **53**
 site map **35, 264**
 server links **36**
 SMB **59**
 SMB backup **281**
 SMB2 **63, 65**
 SMTP methods supported **258**
 Snap EDR **272, 279**
 Snap Finder **273**
 SnapCLI **299**
 running scripts **301**
 syntax **299**
 SnapDRImage **239**
 SnapExtensions
 BitTorrent **269**
 main page **269**
 Snap EDR **272**
 snapshot
 access **286**
 autobackup of volume settings **240**

combined pools **112**
 coordinating with backup jobs **139**
 estimating storage requirements for **141**
 excluding iSCSI Disks from **147**
 managing **135**
 shares **287**
 ways to adjust pool size **141**
 snapshot share **287**
 SNMP configuration **76**
 SNMP configuration page **76**
 software information pop-up **35**
 software update **250**
 specifications, GuardianOS **11**
 specifications, RAINcloudOS **11**
 speed/duplex options **53**
 standalone **53**
 storage
 guides **103**
 pools **83**
 and incompatible disk drives **84**
 creating initial **22**
 RAID Sets screen **105**
 Volumes screen **118**
 storage pools
 and expansion units **90**
 support **254**
 switch user (logout) **266**
 system monitoring **225**
 System Status page **226**

T

Tape page **234**
 TCP/IP
 configuring **52**
 initial configuration **19**
 options **50**
 technical support **3**
 terminology for AFP **66**
 Tivoli NetView **76**
 tools **257**
 Traditional RAID **102**
 adding disk drives **116**
 and expansion units **112**
 compared to DynamicRAID **292**
 quotas **126**
 RAID sets **103**
 shares **103**

- storage guides **103**
- volumes **103**

typographical conventions **4**

U

UID **175**

Uninterruptible Power Supplies (UPS) **41**

unlock a user password **197**

updates

- overview **250**
- procedure **250**

UPS

- configuring **41**
- enabling support for **41**
- low-power warning **41**

users

- creating local **194**
- file-level access for **288**

V

version of OS **3**

volumes

- and DynamicRAID **96**
- and Traditional RAID **118**
- backing up configuration **239**
- capacity reached alert **259**
- dynamic **292**
- expanding capacity of **123**
- Properties screen **122**
- quotas **293**
- screen **118**
- size limits **293**

W

warranty activation **255**

Web Management Interface

- alert messages **34**

Web Management Interface, overview **32**

Web Root **78**

Web Server **78**

Windows

- enabling guest account access **63, 64**
- guest account access **62**
- name resolution server support **60**
- networking (SMB) **59**

security, joining

- active directory domain **63**

see also *Active Directory*

see also *Authentication*

Windows Active Directory **61**

- setup **176**
- Shares **179**

workgroup environment **61**

workgroup, joining **62**

write cache **122, 147**

Z

Zoom Bar **232**